Concepts & Examples
ScreenOS Reference Guide

Volume 3:
Administration

Release 5.4.0, Rev. A

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

Part Number: 530-015770-01, Revision A

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.

Writers: ScreenOS Team

Editor: Lisa Eldridge

Table of Contents

Volume 1 Continued

About This Volume

Juniper Networks security devices provide different ways for you to manage the devices, either locally or remotely. *Volume 3: Administration* contains the following chapters:

- Chapter 1, "Administration," explains the different means available for managing a security device both locally and remotely. This chapter also explains the privileges pertaining to each of the four levels of network administrators that can be defined.
- Chapter 2, "Monitoring Security Devices," explains various monitoring methods and provides guidance in interpreting monitoring output.

Document Conventions

This document uses several types of conventions, which are introduced in the following sections:

CLI Conventions

The following conventions are used to present the syntax of 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

  means "set the management options for the ethernet1, the ethernet2, *or* the ethernet3 interface."

- Variables are in *italic* type:

  set admin user *name1* password *xyz*

In text:

- Commands are in **boldface** type.
- Variables are in *italic* type.

NOTE: When entering a keyword, you only have to type enough letters to identify the word uniquely. For example, typing **set adm u kath j12fmt54** is enough to enter the command **set admin user kathleen j12fmt54**. Although you can use this shortcut when entering commands, all the commands documented here are presented in their entirety.

Illustration Conventions

The following figure shows the basic set of images used in illustrations throughout this manual.

Figure 1: Images in Manual Illustrations

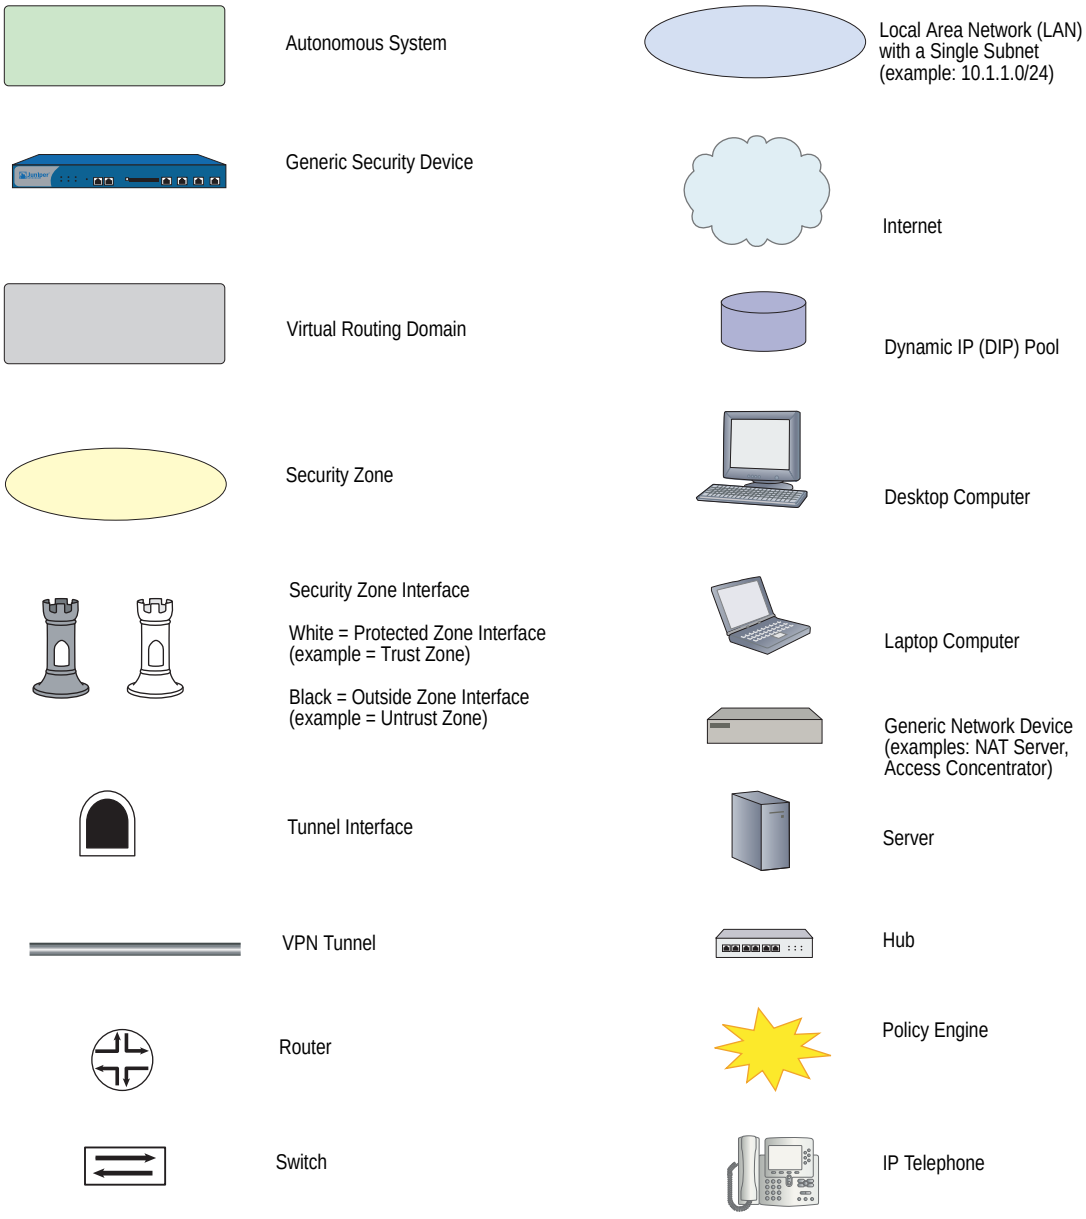

Naming Conventions and Character Types

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 hexadecimals) 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.

WebUI Conventions

A chevron (>) shows the navigational sequence through the WebUI, which you follow by clicking menu options and links. The following figure shows the following path to the address configuration dialog box—Objects > Addresses > List > New:

Figure 2: WebUI Navigation

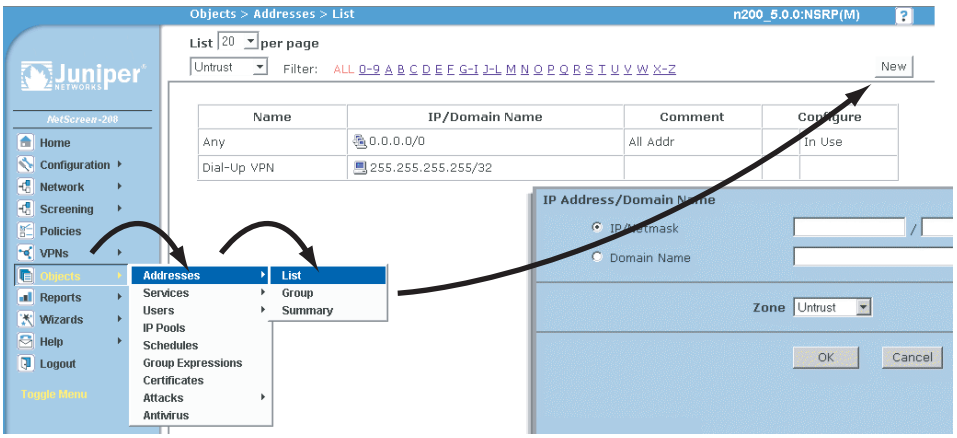

To perform a task with the WebUI, you first navigate to the appropriate dialog box, where you then define objects and set parameters. The set of instructions for each task is divided into navigational path and configuration settings:

The next figure lists the path to the address configuration dialog box with the following sample configuration settings:

Objects > 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

Figure 3: Navigational Path and Configuration Settings

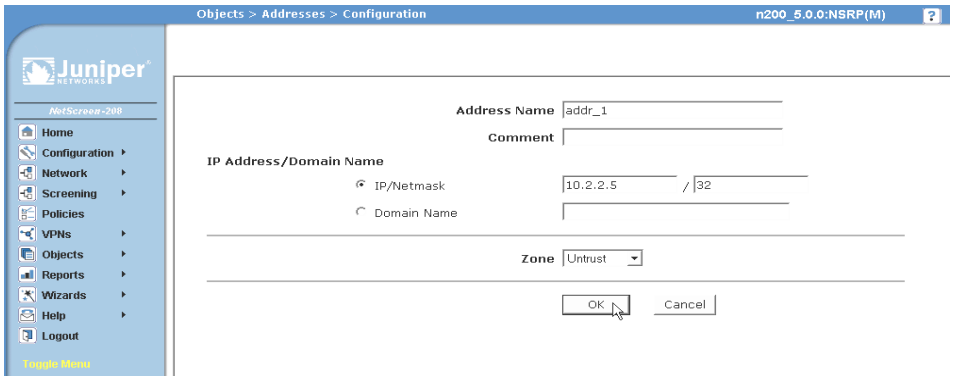

Juniper Networks Documentation

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

For technical support, open a support case using the Case Manager link at http://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (outside the United States).

If you find any errors or omissions in this document, please contact us at the email address below:

techpubs-comments@juniper.net

Chapter 1

Administration

This chapter describes management methods and tools, methods for securing administrative traffic, and the administrative privilege levels that you can assign to admin users. This chapter contains the following sections:

Management via the Web User Interface

You can use the Web User Interface (WebUI) to configure and manage the software for Juniper Networks security devices. Figure 1 shows the WebUI window. The left pane contains the navigation menu, and the right pane displays the navigation window.

Figure 1: WebUI

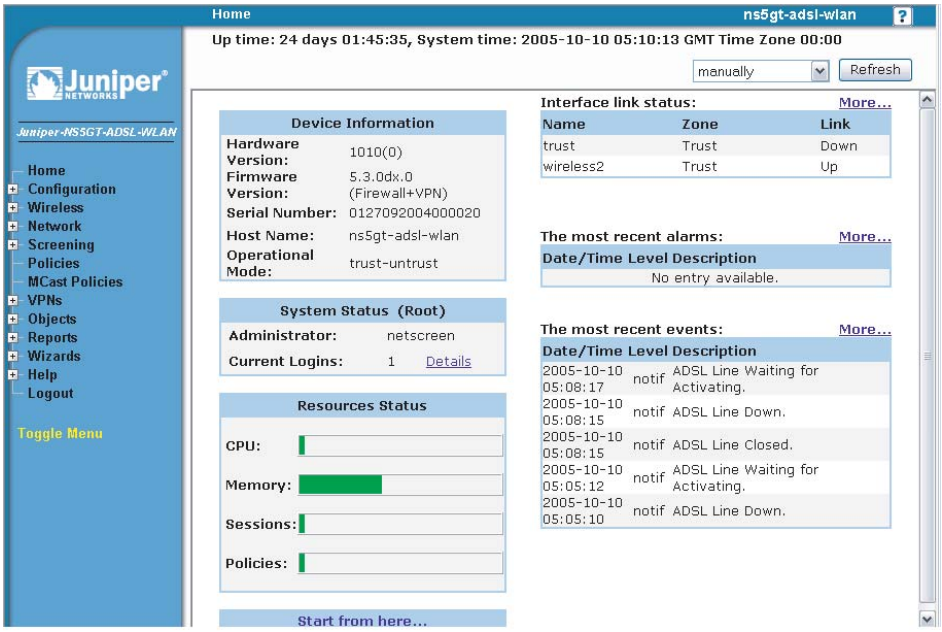

To use the WebUI, you must have the following application and connection:

- Netscape Communicator (version 4.7 or later) or Microsoft Internet Explorer (version 5.5 or later)
- TCP/IP network connection to the security device

WebUI Help

You can view Help files for the WebUI at http://help.juniper.net/help/english/*screenos_version* (for example, http://help.juniper.net/help/english/5.3).

You also have the option of relocating the Help files. You might want to store them locally and point the WebUI to either the administrator's workstation or a secured server on the local network. In case you do not have Internet access, storing the Help files locally provides accessibility to them you otherwise would not have.

Copying the Help Files to a Local Drive

The Help files are available on the documentation CD. You can modify the WebUI to point to the Help files on the CD in your local CD drive. You can also copy the files from the CD to a server on your local network or to another drive on your workstation and configure the WebUI to invoke the Help files from that location.

NOTE: If you want to run the Help files directly from the documentation CD, you can skip this procedure. Proceed to "Pointing the WebUI to the New Help Location" on this page.

---

1. Load the documentation CD in the CD drive of your workstation.
2. Navigate to the CD drive and copy the directory named **help**.
3. Navigate to the location where you want to store the Help directory and paste the Help directory there.

Pointing the WebUI to the New Help Location

You must now redirect the WebUI to point to the new location of the Help directory. Change the default URL to the new file path, where *path* is the specific path to the Help directory from the administrator's workstation.

1. Configuration > Admin > Management: In the Help Link Path field, replace the default URL:

   http://help.juniper.net/help/english/*screenos_version*

   with

   (for local drive) file://*path*…/help

   or

   (for local server) http://*server_name*…/*path*/help

2. Click **Apply**.

   When you click the **help** link in the upper right corner of the WebUI, the device now uses the new path that you specified in the Help Link Path field to locate the appropriate Help file.

HyperText Transfer Protocol

With a standard browser, you can access, monitor, and control your network security configurations remotely using HyperText Transfer Protocol (HTTP).

You can secure HTTP administrative traffic by encapsulating it in a virtual private network (VPN) tunnel or by using the Secure Sockets Layer (SSL) protocol. You can further secure administrative traffic by completely separating it from network user traffic. To do this, you can run all administrative traffic through the MGT interface—available on some security devices—or bind an interface to the MGT zone and devote it exclusively to administrative traffic.

NOTE: For more information, see "Secure Sockets Layer" on page 5, "VPN Tunnels for Administrative Traffic" on page 41, and "MGT and VLAN1 Interfaces" on page 27.

Session ID

The security device assigns each HTTP administrative session a unique session ID. For security devices that support virtual systems (vsys), the ID is globally unique across all systems—root and vsys.

Each session ID is a 39-byte number resulting from the combination of five pseudo-randomly generated numbers. The randomness of the ID generation—versus a simple numerical incremental scheme—makes the ID nearly impossible to predict. Furthermore, the randomness combined with the length of the ID makes accidental duplication of the same ID for two concurrent administrative sessions extremely unlikely.

The following are two benefits that a session ID provides to administrators:

- Figure 2 illustrates how the security device can distinguish concurrent sessions from multiple admins behind a NAT device that assigns the same source IP address to all outbound packets.

Figure 2: Session ID with a NAT device

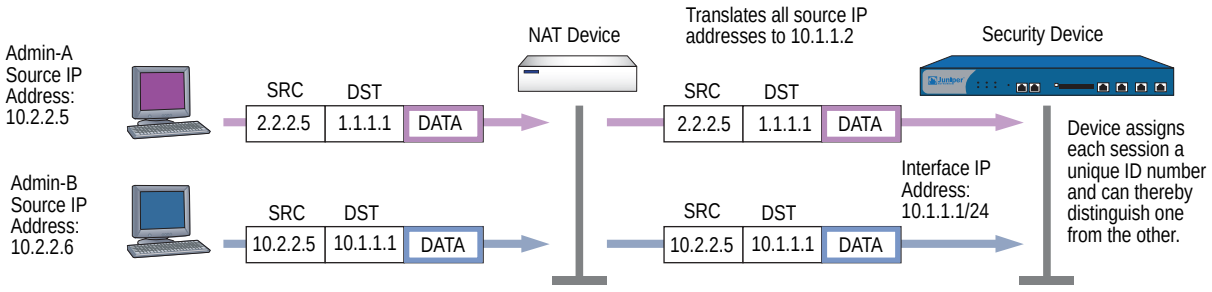

- Figure 3 illustrates how the security device can distinguish concurrent root-level admin sessions from the same source IP address to the root system and from there to different virtual systems.

Figure 3: Session ID with Source IP Address

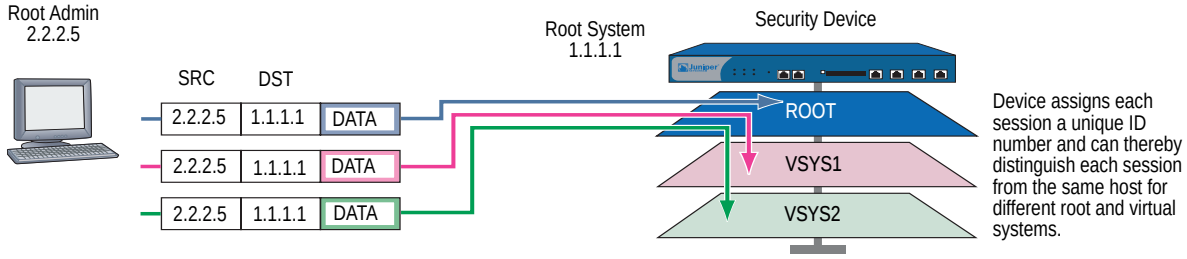

Secure Sockets Layer

Secure Sockets Layer (SSL) is a set of protocols that can provide a secure connection between a web client and a webserver communicating over a TCP/IP network. SSL consists of the SSL Handshake Protocol (SSLHP), which can allow the client and server to authenticate each other and negotiate an encryption method, and the SSL Record Protocol (SSLRP), which provides basic security services to higher-level protocols such as HTTP. These two protocols operate at the following two layers in the Open Systems Interconnection (OSI) Model:

- SSLHP at the Application Layer (Layer 7)
- SSLRP at the Presentation Layer (Layer 6)

Independent of application protocol, SSL uses TCP to provide secure service (see Figure 4). SSL uses certificates to authenticate first the server or both the client and the server, and then encrypt the traffic sent during the session. ScreenOS supports authentication only of the server (security device), not the client (administrator attempting to connect to the security device through SSL).

Figure 4: SSL Client to Server

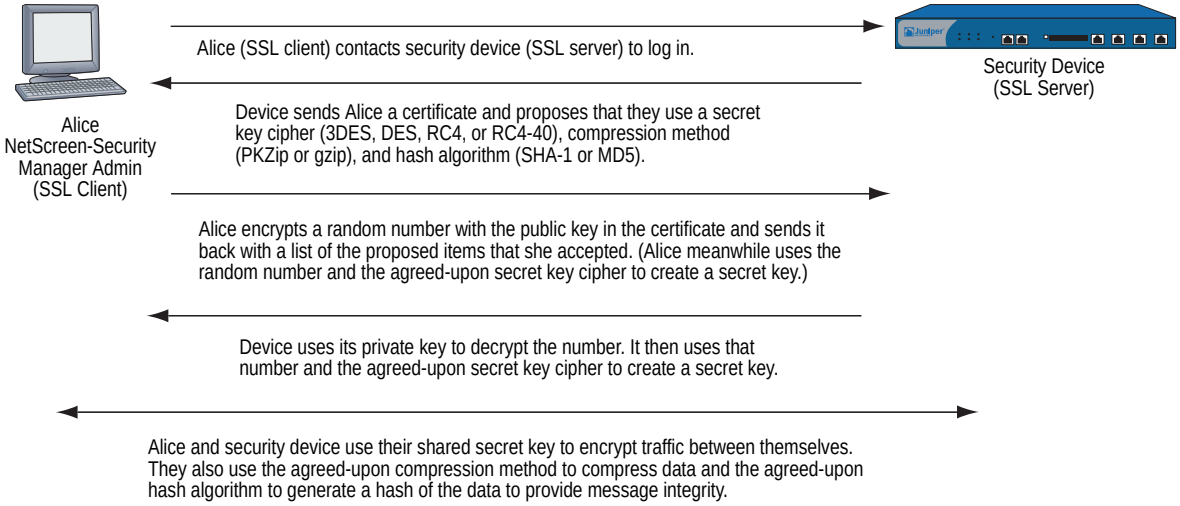

A Juniper Networks security device can redirect administrative traffic using HTTP (default port 80) to SSL (default port 443). The default certificate for SSL is the automatically generated self-signed certificate, although you can later use a different certificate if you want. Because SSL is integrated with PKI key/certificate management, you can select the SSL certificate from any in the certificate list. You can also use the same certificate for an IPSec VPN.

NOTE: For information about redirecting administrative HTTP traffic to SSL, see "Redirecting HTTP to SSL" on page 8. For information about self-signed certificates, see "Self-Signed Certificates" on page **5**-37. For information on obtaining certificates, see "Certificates and CRLs" on page **5**-24.

The ScreenOS implementation of SSL provides the following capabilities, compatibilities, and integration:

- SSL server authentication (not SSL server and client authentication); that is, the security device authenticates itself to the administrator attempting to connect through SSL, but the administrator does not use SSL to authenticate himself to the device
- SSL version 3 compatibility (not version 2)
- Compatibility with Netscape Communicator 4.7x and later and Internet Explorer 5.x later
- Public Key Infrastructure (PKI) key management integration (see "Public Key Cryptography" on page 19)
- The following encryption algorithms for SSL:
  - RC4-40 with 40-bit keys
  - RC4 with 128-bit keys
  - DES: Data Encryption Standard with 56-bit keys
  - 3DES: Triple DES with 168-bit keys
- The same authentication algorithms for SSL as for VPNs:
  - Message Digest version 5 (MD5)—128-bit keys
  - Secure Hash Algorithm version 1 (SHA-1)—160-bit keys

NOTE: The RC4 algorithms are always paired with MD5; DES and 3DES with SHA-1.

SSL Configuration

The basic steps for setting up SSL are as follows:

1. Make use of the self-signed certificate that the security device automatically generates during its initial bootup, or create another self-signed certificate, or obtain a CA-signed certificate and load it on the device.

NOTE: Check your browser to see how strong the ciphers can be and which ones your browser supports. (Both the security device and your browser must support the same kind and size of ciphers you use for SSL.) In Internet Explorer 5x, click **Help**, **About Internet Explorer**, and read the section about cipher strength. To obtain the advanced security package, click **Update Information**. In Netscape Communicator, click **Help**, **About Communicator**, and read the section about RSA. To change the SSL configuration settings, click **Security Info**, **Navigator**, **Configure SSL v3.**

For more information, see "Self-Signed Certificates" on page **5**-37. For details on requesting and loading a certificate, see "Certificates and CRLs" on page **5**-24.

---

2. Enable SSL management.

NOTE: SSL is enabled by default.

WebUI

Configuration > Admin > Management: Enter the following, then click **Apply**:

SSL: (select)
Port: Use the default port number (443) or change it to another.
Certificate: Select the certificate you intend to use from the drop-down list.
Cipher: Select the cipher you intend to use from the drop-down list.

NOTE: If you change the SSL port number, the admins need to specify the nondefault port number when entering the URL in their browser.

CLI

```
set ssl port num
set ssl cert id_num
set ssl encrypt { { 3des | des } sha-1 | { rc4 | rc4-40 } | md5 }
set ssl enable
save
```

NOTE: To learn the ID number for a certificate, use the following command:
get pki x509 list cert.

---

3. Configure the interface through which you manage the security device to permit SSL management:

WebUI

Network > Interfaces > Edit (for the interface you want to manage): Select the SSL management service checkbox, then click **OK**.

CLI

```
set interface interface manage ssl
save
```

4. Connect to the security device via the SSL port. When you enter the IP address for managing the security device in the browser's URL field, change **http** to **https**, and follow the IP address with a colon and the HTTPS (SSL) port number if you have changed it from the default. For example:

   https://123.45.67.89:1443).

Redirecting HTTP to SSL

The security device can redirect administrative traffic using HTTP (default port 80) to SSL (default port 443), as shown in Figure 5.

During the SSL handshake, the security device sends Alice its certificate. Alice encrypts a random number with the public key contained in the certificate and sends it back to the device, which uses its private key to decrypt the number. Both participants then use the shared random number and a negotiated secret key cipher (3DES, DES, RC4, or RC4-40) to create a shared secret key, which they use to encrypt traffic between themselves. They also use an agreed-upon compression method (PKZip or gzip) to compress data and an agreed-upon hash algorithm (SHA-1 or MD-5) to generate a hash of the data to provide message integrity.

Figure 5: Redirection of HTTP to SSL

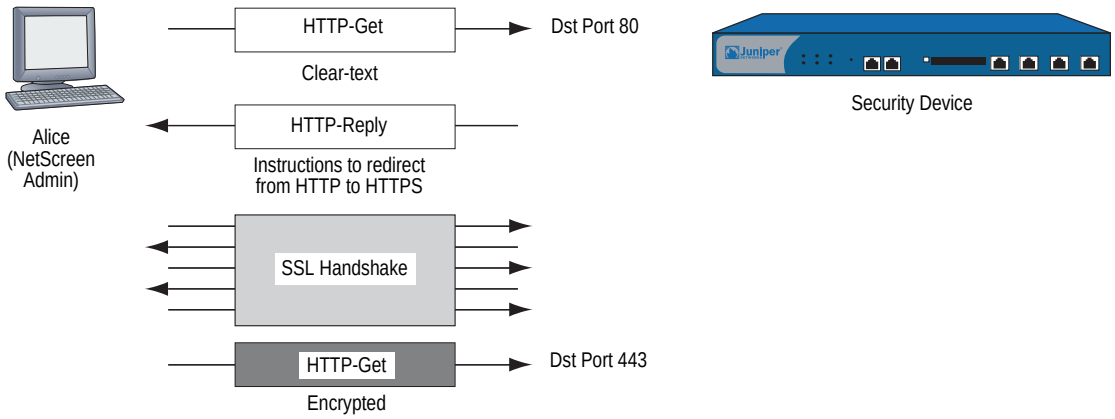

To enable the redirection and use the default automatically generated self-signed certificate for SSL, do either of the following:

WebUI

Configuration > Admin > Management: Enter the following, then click **Apply**:

Redirect HTTP to HTTPS**:** (select)
Certificate: Default – System Self-Signed Cert

CLI

```
set admin http redirect
save
```

NOTE: You do not have to enter a CLI command to apply the automatically generated self-signed certificate for use with SSL because the security device applies it to SSL by default. If you have previously assigned another certificate for use with SSL and you now want to use the default certificate instead, you must unset the other certificate with the **unset ssl cert** *id_num* command, in which *id_num* is the ID number of the previously assigned certificate.

Although HTTP does not provide the security that SSL does, you can configure the security device so that it does not redirect HTTP traffic. To disable the HTTP-to-SSL redirect mechanism, clear the Redirect HTTP to HTTPS option in the WebUI, or enter the **unset admin http redirect** CLI command.

Management via the Command Line Interface

Advanced administrators can attain finer control by using the Command Line Interface (CLI). To configure a security device with the CLI, you can use any software that emulates a VT100 terminal. With a terminal emulator, you can configure the security device using a console from any Windows, UNIX, or Macintosh operating system. For remote administration through the CLI, you can use Telnet or Secure Shell (SSH). With a direct connection through the console port, you can use HyperTerminal.

NOTE: For a complete listing of the ScreenOS CLI commands, refer to the *ScreenOS CLI Reference Guide: IPv4 Command Descriptions*.

Telnet

Telnet is a login and terminal emulation protocol that uses a client/server relationship to connect to and remotely configure network devices over a TCP/IP network. The administrator launches a Telnet client program on the administration workstation and creates a connection with the Telnet server program on the security device. After logging in, the administrator can issue CLI commands, which are sent to the Telnet program on the security device, effectively configuring the device as if operating through a direct connection. Using Telnet to manage security devices requires the following application and connection:

- Telnet software on the administrative workstation
- An Ethernet connection to the security device

Figure 6 illustrates the setup procedure for establishing a Telnet connection.

Figure 6: Establishing a Telnet Connection

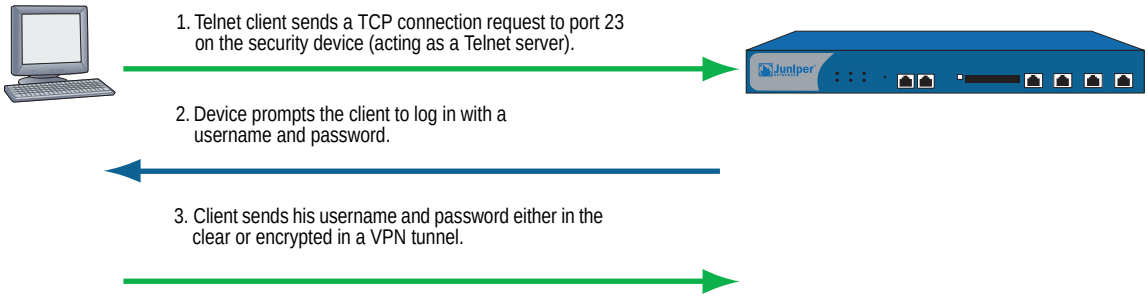

To minimize an unauthorized user's chances of logging into a device, you can limit the number of unsuccessful login attempts allowed before the security device terminates a Telnet session. This restriction also protects against certain types of attacks, such as automated dictionary attacks.

By default, the device allows up to three unsuccessful login attempts before it closes the Telnet session. To change this number, enter the following command:

set admin access attempts *number*

NOTE: You must use the CLI to set this restriction.

Securing Telnet Connections

You can secure Telnet traffic by completely separating it from network user traffic. Depending upon your security device model, you can run all administrative traffic through the MGT interface or devote an interface such as the DMZ entirely to administrative traffic.

In addition, to ensure that admin users use a secure connection when they manage a security device through Telnet, you can require such users to Telnet only through a virtual private network (VPN) tunnel. After you have set this restriction, the device denies access if anyone tries to Telnet without going through a VPN tunnel.

NOTE: For information about VPN tunnels, see *Volume 5: Virtual Private Networks*.

To restrict Telnet access through a VPN, enter the following command:

set admin telnet access tunnel

NOTE: You must use the CLI to set this restriction.

Secure Shell

The built-in Secure Shell (SSH) server on a Juniper Networks security device provides a means by which administrators can remotely manage the device in a secure manner using applications that are SSH aware. SSH allows you to open a remote command shell securely and execute commands. SSH provides protection from IP or DNS spoofing attacks and password or data interception.

You can choose to run either an SSH version 1 (SSHv1) or an SSH version 2 (SSHv2) server on the device. SSHv2 is considered more secure than SSHv1 and is currently being developed as the IETF standard. However, SSHv1 has been widely deployed and is commonly used. Note that SSHv1 and SSHv2 are not compatible with each other. That is, you cannot use an SSHv1 client to connect to an SSHv2 server on the security device, or vice versa. The client console or terminal application must run the same SSH version as the server. Figure 7 illustrates SSH traffic flow.

Figure 7: SSH Traffic Flow

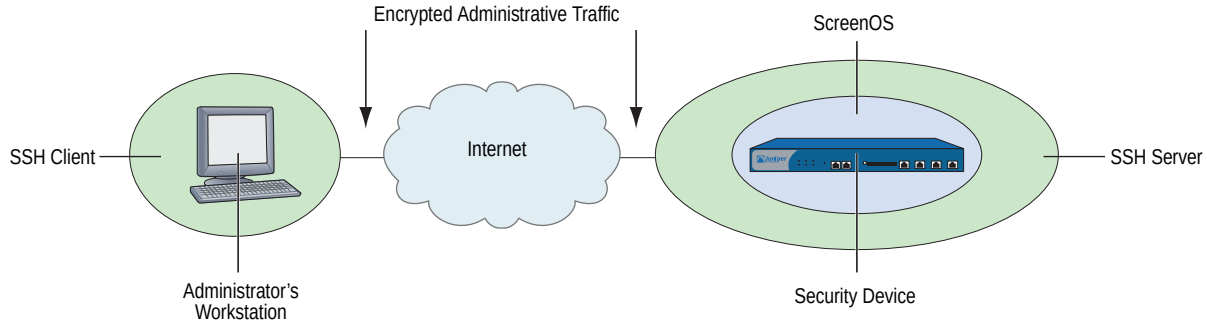

Figure 8 illustrates the basic SSH connection procedure.

Figure 8: SSH Connection

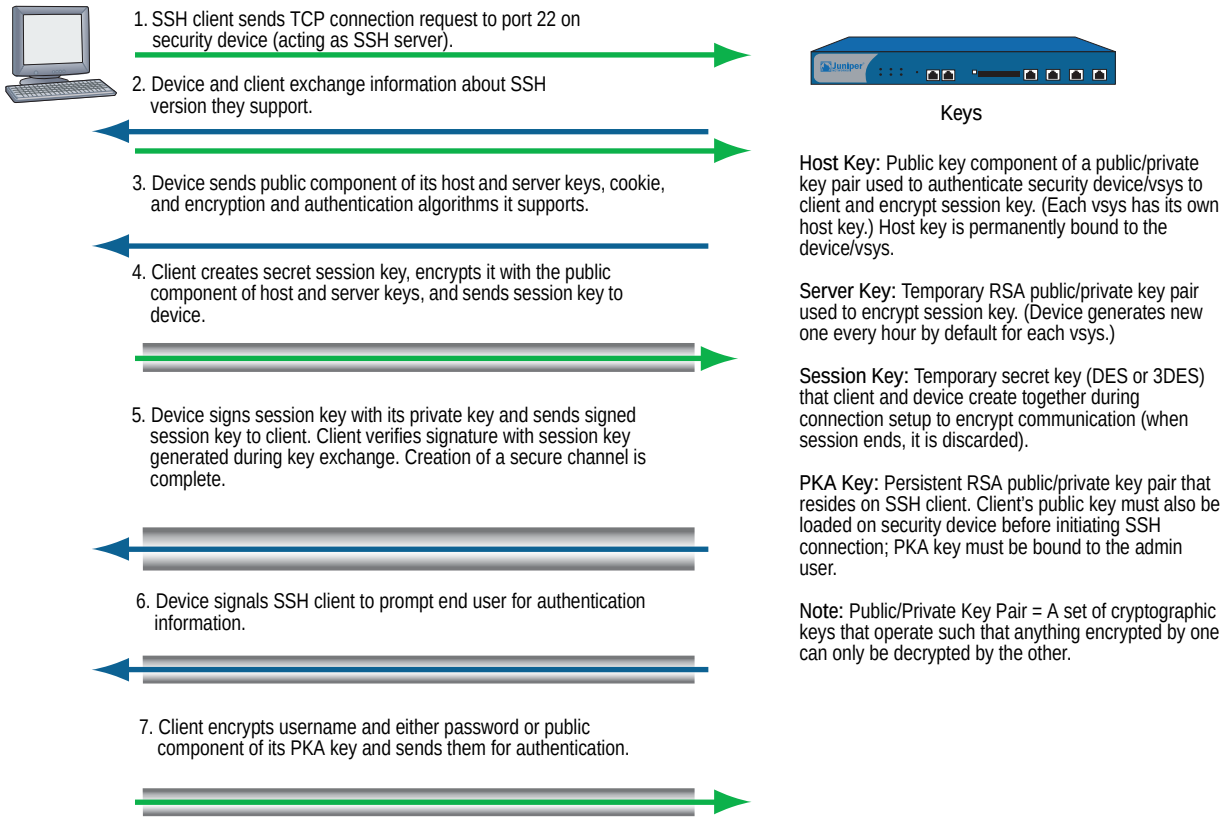

A maximum of five SSH sessions are allowed on a Juniper Networks security device at any one time.

Client Requirements

As mentioned in "Secure Shell" on page 11, the client application must run the same SSH version as the server on the security device. SSHv2 clients must be configured to request the Diffie-Hellman key exchange algorithm and the Digital Signature Algorithm (DSA) for public key device authentication. SSHv1 clients must be configured to request the RSA for public key device authentication.

Basic SSH Configuration on the Device

The following are the basic steps for configuring SSH on a Juniper Networks security device:

1. Determine whether you will use password or Public Key Authentication (PKA) for SSH. If PKA will be used, the PKA keys must be bound to an admin before SSH connections can be made. See "Authentication" on page 14 for more information about using passwords or PKA.

2. Determine which version of SSH you need to enable on the security device. (Remember that the client application and the SSH server on the device must run the same SSH version.) If you enabled SSH on the device in a previous ScreenOS version, SSHv1 runs when you enable SSH now. To see which version of SSH is active but not enabled on the device, enter the **get ssh** CLI command:

   ```
   device> get ssh
   SSH V1 is active
   SSH is not enabled
   SSH is not ready for connections
   Maximum sessions: 8
   Active sessions: 0
   ```

   In the output shown above, SSHv1 is active and runs when you enable SSH. If you want to use a different SSH version, make sure that all keys created with the previous version are removed. For example, to clear SSHv1 keys and to use SSHv2, enter the following CLI commands:

   ```
   device> delete ssh device all
   ```

   The following messages appear:

   ```
   SSH disabled for vsys: 1
   PKA key deleted from device: 0
   Host keys deleted from device: 1
   Execute the 'set ssh version v2' command to activate SSH v2 for the device
   ```

   To use SSHv2, enter the following CLI command:

   ```
   device-> set ssh version v2
   ```

NOTE: Setting the SSH version does not enable SSH on the security device.

3. If you do not want to use port 22 (the default port) for SSH client connections, you can specify a port number between 1024 and 32767.

NOTE: You can also use the WebUI to change the port number and enable SSHv2 and SCP on the Configuration > Admin > Management page.

```
device-> set admin ssh port 1024
```

4. Enable SSH for the root system or for the virtual system. See "SSH and Vsys" on page 16 for additional information about enabling and using SSH on a per-vsys basis.

To enable SSH for the root system:

```
device-> set ssh enable
```

To enable SSH for a vsys, you need to first enter the vsys and then enable SSH:

```
device-> set vsys v1
device(v1)-> set ssh enable
```

5. Enable SSH on the interface on which the SSH client will connect.

```
device-> set interface manage ssh
```

6. Distribute the host key that is generated on the security device to the SSH client. See "Host Key" on page 16 for more information.

Authentication

An administrator can connect to a Juniper Networks security device with SSH using one of two authentication methods:

- **Password Authentication:** This method is used by administrators who need to configure or monitor a security device. The SSH client initiates an SSH connection to the device. If SSH manageability is enabled on the interface receiving the connection request, the device signals the SSH client to prompt the user for a username and password. When the SSH client has this information, it sends it to the device, which compares it with the username and password in the admin user's account. If they match, the device authenticates the user. If they do not match, the device rejects the connection request.

- **Public Key Authentication (PKA):** This method provides increased security over the password authentication method and allows you to run automated scripts. Instead of a username and password, the SSH client sends a username and the public key component of a public/private key pair. The device compares it with up to four public keys that can be bound to an admin. If one of the keys matches, the device authenticates the user. If none of them matches, the device rejects the connection request.

NOTE: The supported authentication algorithms are RSA for SSHv1 and DSA for SSHv2.

Both authentication methods require the establishment of a secure connection before the SSH client logs on. After an SSH client has established an SSH connection with the device, he must authenticate himself either with a username and password or with a username and public key.

Both password authentication and PKA require that you create an account for the admin user on the device and enable SSH manageability on the interface through which you intend to manage the device via an SSH connection. (For information about creating an admin user account, see "Defining Admin Users" on page 33.) The password authentication method does not require any further set up on the SSH client.

On the other hand, to prepare for PKA, you must first perform the following tasks:

1. On the SSH client, generate a public and private key pair using a key generation program. (The key pair is either RSA for SSHv1 or DSA for SSHv2. See the SSH client application documentation for more information.)

NOTE: If you want to use PKA for automated logins, you must also load an agent on the SSH client to decrypt the private key component of the PKA public/private key pair and hold the decrypted version of the private key in memory.

---

2. Move the public key from the local SSH directory to a directory on your TFTP server, and launch the TFTP program.

NOTE: You can also paste the content of the public key file directly into the CLI command **set ssh pka-rsa** [**username** *name_str*] **key** *key_str* (for SSHv1) or **set ssh pka-dsa** [**user-name** *name_str*] **key** *key_str* (for SSHv2), pasting it where indicated by the variable *key_str*, or into the Key field in the WebUI (Configuration > Admin > Administrators > SSH PKA). However, the CLI and WebUI have a size restriction: the public key size cannot exceed 512 bits. This restriction is not present when loading the key via TFTP.

---

3. Log into the device so that you can configure it through the CLI.

4. To load the public key from the TFTP server to the device, enter one of the following CLI commands:

   For SSHv1:

   exec ssh tftp pka-rsa [**username** *name*] **file-name** *name_str* **ip-addr** *tftp_ip_addr*

   For SSHv2:

   exec ssh tftp pka-dsa [**user-name** *name*] **file-name** *name_str* **ip-addr** *tftp_ip_addr*

   The **username** or **user-name** options are only available to the root admin, so that only the root admin can bind an RSA key to another admin. When you—as the root admin or as a read/write admin—enter the command without a username, the device binds the key to your own admin account; that is, it binds the key to the admin that enters the command.

NOTE: The security device supports up to four PKA public keys per admin user.

When an administrator attempts to log in via SSH on an interface that has SSH manageability enabled, the device first checks if a public key is bound to that administrator. If so, the device authenticates the administrator using PKA. If a public key is not bound to the administrator, the device prompts for a username and password. (You can use the following command to force an admin to use only the PKA method: **set admin ssh password disable username** *name_str*.)

Regardless of the authentication method you intend the administrator to use, when you initially define his or her account, you still must include a password, even though when you later bind a public key to this user, the password becomes irrelevant.

SSH and Vsys

For security devices that support vsys, you can enable and configure SSH on a per-vsys basis. Each vsys has its own host key (see "Host Key" on page 16) and maintains and manages a PKA key for the admin of the system.

The maximum number of SSH sessions is a device-wide limit and is between 2 and 24, depending upon the platform. If the maximum number of SSH clients are already logged into the device, no other SSH client can log into the SSH server. The root system and the vsys share the same SSH port number. This means that if you change the SSH port from the default port 22, the port is changed for all vsys as well.

NOTE: When you deploy a large number of virtual systems on a single device, be aware that if many or all vsys admins use SSH, the storage reserved for PKI objects can fill up.

Host Key

The host key allows the security device to identify itself to an SSH client. On devices that support virtual systems (vsys), each vsys has its own host key. When SSH is first enabled on a vsys (for devices that support vsys) or on a device, a host key is generated that is unique to the vsys or device. The host key is permanently bound to the vsys or device and the same host key is used if SSH is disabled and then enabled again.

The host key on the device must be distributed to the SSH client in one of two ways:

- Manually—the root or vsys admin sends the host key to the client admin user via email, telephone, and so on. The receiving admin stores the host key in the appropriate SSH file on the SSH client system. (The SSH client application determines the file location and format.)

- Automatically—When the SSH client connects to the device, the SSH server sends the unencrypted public component of the host key to the client. The SSH client searches its local host key database to see if the received host key is mapped to the address of the device. If the host key is unknown (there is no mapping to the device address in the client's host key database), the Admin user might be able to decide whether to accept the host key. Otherwise, the connection is terminated. (See the appropriate SSH client documentation for information on accepting unknown host keys.)

To verify that the SSH client has received the correct host key, the Admin user on the client system can generate the SHA hash of the received host key. The client Admin user can then compare this SHA hash with the SHA hash of the host key on the device. On the device, you can display the SHA hash of the host key by executing the CLI command **get ssh host-key**.

Example: SSHv1 with PKA for Automated Logins

In this example, you (as the root admin) set up SSHv1 public key authentication (PKA) for a remote host that runs an automated script. The sole purpose for this remote host to access the device is to download the configuration file every night. Because authentication is automated, no human intervention is necessary when the SSH client logs into the device.

You define an admin user account named **cfg**, with password **cfg** and read-write privileges. You enable SSH manageability on interface ethernet1, which is bound to the Untrust zone.

You have previously used a key generation program on your SSH client to generate an RSA public/private key pair, moved the public key file, which has the filename "idnt_cfg.pub", to a directory on your TFTP server, and launched the TFTP program. The IP address of the TFTP server is 10.1.1.5.

WebUI

Configuration > Admin > Administrators > New: Enter the following, then click **OK**:

Name: cfg
New Password: cfg
Confirm Password: cfg
Privileges: Read-Write (select)
SSH Password Authentication: (select)

Network > Interfaces > Edit (for ethernet1): Select **SSH** in Service Options, then click **OK**.

NOTE: You can only load a public key file for SSH from a TFTP server via the **exec ssh** command.

CLI

```
set admin user cfg password cfg privilege all
set interface ethernet1 manage ssh
exec ssh tftp pka-rsa username cfg file-name idnt_cfg.pub ip-addr 10.1.1.5
save
```

Secure Copy

Secure Copy (SCP) provides a way for a remote client to transfer files to or from the security device using the SSH protocol. (The SSH protocol provides authentication, encryption, and data integrity to the SCP connection.) The device acts as an SCP server to accept connections from SCP clients on remote hosts.

SCP requires that the remote client be authenticated before file transfer commences. SCP authentication is exactly the same process used to authenticate SSH clients. The SCP client can be authenticated with either a password or a PKA key. Once the SCP client is authenticated, one or more files can be transferred to or from the device. The SCP client application determines the exact method for specifying the source and destination file names; refer to the SCP client application documentation.

SCP is disabled by default on the device. To enable SCP, you must also enable SSH.

WebUI

Configuration > Admin > Management: Select the following, then click **Apply**:

Enable SSH: (select)
Enable SCP: (select)

CLI

```
set ssh enable
set scp enable
save
```

The following is an example of an SCP client command to copy the configuration file from flash memory on a device (administrator name is “juniper” and the IP address is 10.1.1.1) to the file “ns_sys_config_backup” on the client system:

```
scp juniper@10.1.1.1:ns_sys_config ns_sys_config_backup
```

You can also copy a ScreenOS image to and from a device. To save an image named “ns.5.1.0r1” to a device from an SCP client, enter the following SCP client command, in which the administrator's login name is “juniper” and the IP address of the device is 10.1.1.1:

```
scp ns.5.1.0r1 juniper@10.1.1.1:image
```

Then enter the **reset** command to reboot the security device to load and run the new ScreenOS image.

To copy a ScreenOS image from a device to an SCP client and name the saved image “current_image_backup,” enter the following SCP client command:

```
scp juniper@10.1.1.1:image current_image_backup
```

You need to consult your SCP client application documentation for information on how to specify the administrator name, device IP address, source file, and destination file.

Serial Console

You can manage a security device through a direct serial connection from the administrator's workstation to the device via the console port. Although a direct connection is not always possible, this is the most secure method for managing the device provided that the location around the device is secure.

NOTE: To prevent unauthorized users from logging in remotely as the root admin, you can require the root admin to log into the device through the console only. For additional information on this restriction, see "Restricting the Root Admin to Console Access" on page 40.

Depending on your Juniper Networks security device model, creating a serial connection requires one of the following cables:

- A female DB-9 to male DB-25 straight-through serial cable
- A female DB-9 to male DB-9 straight-through serial cable
- A female DB-9 to male MiniDIN-8 serial cable
- A female DB-9 to RJ-45 adapter with an RJ-45 to RJ-45 straight-through Ethernet cable

You will also need HyperTerminal software (or another kind of VT100 terminal emulator) on the management workstation, with the HyperTerminal port settings configured as follows:

- Serial communications 9600 bps
- 8 bit
- No parity
- 1 stop bit
- No flow control

NOTE: For more details on using HyperTerminal, refer to *ScreenOS CLI Reference Guide: IPv4 Command Descriptions* or the documentation for your device.

Modem Port

You can also manage a security device by connecting the administrator's workstation to the modem port on the device. The modem port functions similarly to the console port, except that you cannot define parameters for the modem port or use this connection to upload an image.

To prevent unauthorized users from managing the device through a direct connection to the console or modem port, you can disable both ports by entering the following commands:

```
set console disable
set console aux disable
```

NOTE: On a NetScreen-5XT device, you can use the modem port to connect to an external modem only.

Management via NetScreen-Security Manager

NetScreen-Security Manager is Juniper Networks' enterprise-level management software application that configures and monitors multiple Juniper Networks security devices over a local area network (LAN) or a wide area network (WAN) environment. The NetScreen-Security Manager User Interface (UI) enables network administrators to deploy, configure, and manage multiple devices from central locations.

NetScreen-Security Manager uses three components to enable remote communication with security devices:

- The *NetScreen-Security Manager User Interface (UI)* is a java-based software application that you use to access and configure data on your network with the NetScreen-Security Manager management system. From the UI, you can view, configure, and manage your network.
- The *management system* is a set of services that resides on an external host. These services process, track, and store device management information exchanged between a device and the NetScreen-Security Manager UI. The management system is composed of two components:
  - The *GUI Server* receives and responds to requests and commands from the UI. It manages the system resources and configuration data required to manage your network. It also contains a local data store of information about your managed security devices, administrators, and configurations.
  - The *Device Server* acts as a collection point for all data generated by each of your network devices. It stores this data, primarily traffic logs, in the local data store.
  - *NSM Agent* is a service that resides on each managed security device. NSM Agent receives configuration parameters from the external management system and forwards them to ScreenOS. NSM Agent also monitors each device and transmits reports back to the management system. NSM Agent can download signature packs, certificates, and entitlements between a security device and NetScreen-Security Manager.

Figure 9 shows how NSM Agent communicates with the NetScreen-Security Manager UI.

Figure 9: Security Device with NSM Agent Enabled

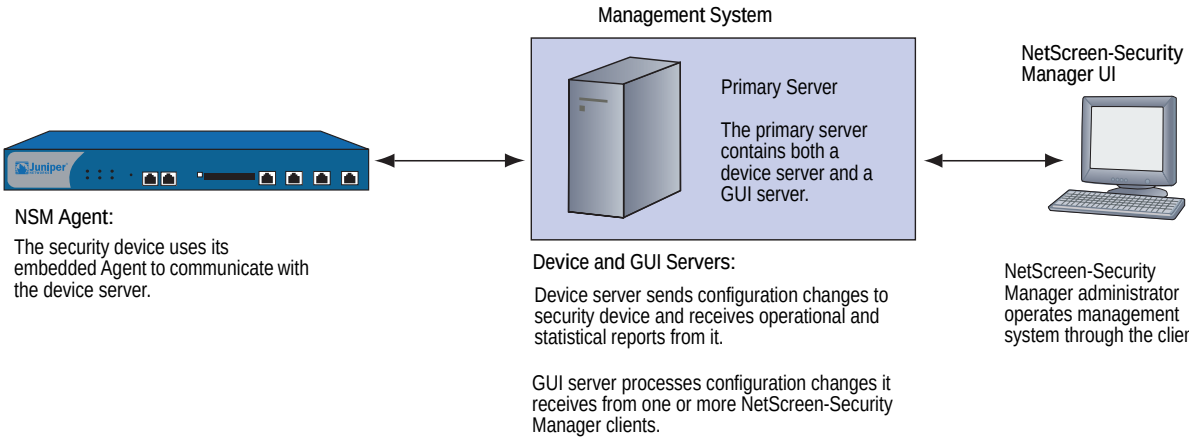

For more information about these and other NetScreen-Security Manager components, refer to the *NetScreen-Security Manager Administrator's Guide*.

Initiating Connectivity Between NSM Agent and the MGT System

Before NetScreen-Security Manager can access and manage a security device, it is necessary to initiate communications between NSM Agent (which resides on the device) and the management system (which resides on an external host). Initialization might require up to two users at two different sites, depending upon the current availability of the security device. These users might include the NetScreen-Security Manager administrator, who uses the NetScreen-Security Manager UI on a client host, and the on-site user, who executes CLI commands on a device via a console session. Possible initialization cases include the following:

- *Case 1:* A device already has a known IP address and is reachable over your network infrastructure.

  In this case, the NetScreen-Security Manager administrator adds the device using the NetScreen-Security Manager UI on the client host. (No on-site user is necessary.) The device automatically connects back to the management system and is ready to send configuration information to the NetScreen-Security Manager database that resides there.

- *Case 2:* The IP address is unreachable.

  In this case, both users perform initialization tasks. The administrator adds the device through the NetScreen-Security Manager UI. The administrator also determines which CLI commands the on-site user needs and delivers them to the user, who then executes them through the console. The device then automatically connects with the management system and is ready to send configuration information to the NetScreen-Security Manager database.

- *Case 3:* The device is a new appliance and contains the factory default settings.

  In this case, both users perform initialization tasks. The on-site user can use an encrypted configuration script, called *Configlet*, which the NetScreen-Security Manager administrator generates. The process is as follows:

  1. The administrator selects the device platform and ScreenOS version, using the Add Device wizard in the NetScreen-Security Manager UI.
  2. The administrator edits the device and enters the desired configuration.
  3. The administrator activates the device.
  4. The administrator generates and delivers the Configlet file (or the necessary CLI commands, as with Case 2) to the on-site user.
  5. The on-site user executes Configlet (or the CLI commands).

For more information, refer to the discussion about adding devices in the *NetScreen-Security Manager Administrator's Guide*.

Enabling, Disabling, and Unsetting NSM Agent

Before a security device can communicate with the management system, you must enable the NetScreen-Security Manager (NSM) Agent residing on the device.

If you want to unset NetScreen-Security Manager, use the **unset nsmgmt all** command. This command sets NSM Agent to its initial defaults, so it acts as though it was never connected to NetScreen-Security Manager. Use the **unset nsmgmt all** command when you want to reconfigure the NetScreen-Security Manager settings.

To enable NSM Agent on the security device, do either of the following:

WebUI

Configuration > Admin > NSM: Select **Enable Communication with NetScreen Security Manager (NSM)**, then click **Apply**.

CLI

```
set nsmgt enable
save
```

To disable NSM Agent on the device, do either of the following:

WebUI

Configuration > Admin > NSM: Clear **Enable Communication with NetScreen Security Manager (NSM)**, then click **Apply**.

CLI

```
unset nsmgt enable
save
```

Setting the Primary Server IP Address of the Management System

The IP address by which NSM Agent identifies the external management system servers is a configurable parameter.

In the following example you set the primary server IP address to 1.1.1.100.

WebUI

Configuration > Admin > NSM: Enter the following, then click **Apply**:

Primary IP Address/Name: 1.1.1.100

CLI

```
set nsmgmt server primary 1.1.1.100
save
```

Setting Alarm and Statistics Reporting

NSM Agent monitors the device events and transmits reports back to the management system. This allows the NetScreen-Security Manager administrator to view the events from the NetScreen-Security Manager UI.

The categories of events tracked by NSM Agent are as follows:

- *Alarms* report potentially dangerous attacks or traffic anomalies, including attacks detected through deep inspection.
- *Log events* report changes in a device's configuration and non-severe changes that occur on a device.
- *Protocol distribution* events report messages generated by the following protocols:
  - Authentication Header (AH)
  - Encapsulating Security Payload (ESP)
  - Generic Routing Encapsulation (GRE)
  - Internet Control Message Protocol (ICMP)
  - Open Shortest Path First (OSPF)
  - Transmission Control Protocol (TCP)
  - User Datagram Protocol (UDP)

- *Statistics* messages report the following statistical information:
  - Attack statistics
  - Ethernet statistics
  - Traffic flow statistics
  - Policy statistics

In the following example, you enable transmission of all Alarm and Statistics messages to the Management System.

WebUI

Configuration > Admin > NSM: Enter the following, then click **Apply**:

Attack Statistics: (select)
Policy Statistics: (select)
Attack Alarms: (select)
Traffic Alarms: (select)
Flow Statistics: (select)
Ethernet Statistics: (select)
Deep Inspection Alarms: (select)
Event Alarms: (select)

CLI

```
set nsmgmt report statistics attack enable
set nsmgmt report statistics policy enable
set nsmgmt report alarm attack enable
set nsmgmt report alarm traffic enable
set nsmgmt report statistics flow enable
set nsmgmt report statistics ethernet enable
set nsmgmt report alarm idp enable
set nsmgmt report alarm other enable
save
```

Configuration Synchronization

If the ScreenOS configuration is changed from the last time it was synchronized with NetScreen-Security Manager, then the security device notifies the NetScreen-Security Manager administrator of the change. For example, the device sends a message when a device administrator uses console, telnet, SSH, or the WebUI to change a security device configuration. Changing the configuration with any application other than NetScreen-Security Manager causes it to be unsynchronized. The NetScreen-Security Manager configuration file must be synchronized with the security device configuration file for NetScreen-Security Manager to work correctly. The synchronization is achieved when you import the configuration file to NetScreen-Security Manager. For information on importing devices, refer to the *NetScreen-Security Manager Administrator's Guide*.

The following example displays the command used to view the configuration status.

Example: Viewing the Configuration State

In the following example, you view the configuration synchronization state of a security device.

WebUI

NOTE: You must use the CLI to retrieve the running configuration state.

CLI

```
get config nsmgmt-dirty
```

NOTE: If applications other than NetScreen-Security Manager applications have not changed the configuration file, then the command returns a blank; otherwise, it returns a "yes."

Example: Retrieving the Configuration Hash

NetScreen-Security Manager uses the configuration hash to verify the configuration synchronization of a security device. In the following example, you retrieve the running configuration hash for a specific virtual system.

WebUI

NOTE: You must use the CLI to retrieve the running configuration hash.

CLI

```
device-> enter vsys vsys1
device(vsys1)-> get config hash
a26a16cd6b8ef40dc79d5b2ec9e1ab4f
device(vsys1)->
device(vsys1)-> exit
```

Retrieving the Configuration Timestamp

A security device provides two configuration timestamps—running-config and saved-config. The running-config timestamp is when the set or unset command was last executed for each virtual system. The saved-config timestamp is when the device configuration was last saved.

In the following example, the security device retrieves the last running and saved configuration timestamps for the vsys1 virtual system:

WebUI

NOTE: You must use the CLI to retrieve the running and saved configuration timestamps.

CLI

```
get config timestamp vsys vsys1
get config saved timestamp
```

NOTE: If you omit **vsys vsys_name** from the command, the security device retrieves the configuration timestamp for the root system. If the timestamp is unavailable, then an "unknown" message is displayed.

Controlling Administrative Traffic

ScreenOS provides the following options for configuring and managing the security device:

- **WebUI:** Selecting this option allows the interface to receive HTTP traffic for management via the Web User Interface (WebUI).
- **Telnet:** A terminal emulation program for TCP/IP networks such as the Internet, Telnet is a common way to remotely control network devices. Selecting this option enables Telnet manageability.
- **SSH:** You can administer the security device from an Ethernet connection or a dial-in modem using Secure Command Shell (SSH). You must have an SSH client that is compatible with Version 1.5 of the SSH protocol. These clients are available for Windows 95 and later, Windows NT, Linux, and UNIX. The security device communicates with the SSH client through its built-in SSH server, which provides device configuration and management services. Selecting this option enables SSH manageability.
- **SNMP:** The security device supports both SNMPv1 and SNMPv2c, and all relevant Management Information Base II (MIB II) groups, as defined in RFC-1213. Selecting this option enables SNMP manageability.
- **SSL:** Selecting this option allows the interface to receive HTTPS traffic for secure management of the security device via the WebUI.
- **NetScreen-Security Manager:** Selecting this option allows the interface to receive NetScreen-Security Manager traffic.
- **Ping:** Selecting this option allows the security device to respond to an ICMP echo request, or ping, which determines whether a specific IP address is accessible over the network.
- **Ident-Reset:** Services like Mail and FTP send identification requests. If they receive no acknowledgement, they send the request again. While the request is processing, there is no user access. By enabling the Ident-reset option, the security device sends a TCP reset announcement in response to an IDENT request to port 113 and restores access that has been blocked by an unacknowledged identification request.

To use these options, you enable them on one or more interfaces, depending on your security and administrative needs.

MGT and VLAN1 Interfaces

Some Juniper Networks security devices have a physical interface—Management (MGT)—dedicated exclusively for management traffic. You can use this interface for management traffic when interfaces are in NAT, Route, or Transparent mode.

In Transparent mode, you can configure all security devices to allow administration through the logical interface, VLAN1. To enable management traffic to reach the VLAN1 interface, you must enable the management options you want both on VLAN1 and on the Layer 2 zones—V1-Trust, V1-Untrust, V1-DMZ, user-defined Layer 2 zone—through which the management traffic passes to reach VLAN1.

To maintain the highest level of security, Juniper Networks recommends that you limit administrative traffic exclusively to the VLAN1 or MGT interface and user traffic to the security zone interfaces. Separating administrative traffic from network user traffic greatly increases administrative security and ensures constant management bandwidth.

Example: Administration Through the MGT Interface

In this example, you set the IP address of the MGT interface to 10.1.1.2/24 and enable the MGT interface to receive web and SSH administrative traffic.

WebUI

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

IP Address/Netmask: 10.1.1.2/24
Management Services: WebUI, SSH: (select)

CLI

```
set interface mgt ip 10.1.1.2/24
set interface mgt manage web
set interface mgt manage ssh
save
```

Example: Administration Through the VLAN1 Interface

In this example, you set the IP address of the VLAN1 interface to 10.1.1.1/24 and enable the VLAN1 interface to receive Telnet and web administrative traffic through the V1-Trust zone.

WebUI

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

IP Address/Netmask: 10.1.1.1/24
Management Services: WebUI, Telnet: (select)

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

Management Services: WebUI, Telnet: (select)

CLI

```
set interface vlan1 ip 10.1.1.1/24
set interface vlan1 manage web
set interface vlan1 manage telnet
set zone v1-trust manage web
set zone v1-trust manage telnet
save
```

Setting Administrative Interface Options

On security devices that have multiple physical interfaces for network traffic, but no physical MGT interface, you might dedicate one physical interface exclusively for administration, separating management traffic completely from network user traffic. For example, you might have local management access the device through an interface bound to the Trust zone and remote management through an interface bound to the Untrust zone.

In this example, you bind ethernet1 to the Trust zone and ethernet3 to the Untrust zone. You assign ethernet1 the IP address 10.1.1.1/24 and give it the Manage IP address 10.1.1.2. (Note that the Manage IP address must be in the same subnet as the security zone interface IP address.) You also allow ethernet1 to receive web and Telnet traffic. You then assign ethernet3 the IP address 1.1.1.1/24 and block all administrative traffic to that interface.

WebUI

Network > Interfaces > Edit (for ethernet1): 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
Management Services:
        WebUI: (select)
        SNMP: (clear)
        Telnet: (select)
        SSL: (clear)
        SSH: (clear)
```

Enter the following, then click **OK**:

```
Interface Mode: NAT
Network > Interfaces > Edit (for ethernet3):
```

Enter the following, then click **OK**:

```
Zone Name: Untrust
Static IP: (select this option when present)
IP Address/Netmask: 1.1.1.1/24
Management Services:
  WebUI: (clear)
  SNMP: (clear)
  Telnet: (clear)
  SSL: (clear)
  SSH: (clear)
```

CLI

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 manage-ip 10.1.1.2
set interface ethernet1 manage web
unset interface ethernet1 manage snmp
set interface ethernet1 manage telnet
unset interface ethernet1 manage ssl
unset interface ethernet1 manage ssh
set interface ethernet1 nat
set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
save
```

NOTE: When you bind an interface to any security zone other than the Trust and V1-Trust zones, all management options are disabled by default. Therefore, in this example, you do not have to disable the management options on ethernet3.

Setting Manage IPs for Multiple Interfaces

Any physical, redundant, or aggregate interface or sub-interface you bind to a security zone can have at least two IP addresses:

- An interface IP address, which connects to a network.
- A logical manage IP address for receiving administrative traffic.

When a security device is a backup unit in a redundant group for high availability (HA), you can access and configure the unit through its manage IP address (or addresses)

NOTE: The manage IP address differs from the VLAN1 address in the following two ways:

When the security device is in Transparent mode, the VLAN1 IP address can be the endpoint of a VPN tunnel, but the manage IP address cannot.

You can define multiple manage IP addresses—one for each network interface—but you can only define one VLAN1 IP address—for the entire system.

If you select the Manageable option on the interface configuration page in the WebUI, you can manage the security device either through the interface IP address or the Manage IP address associated with that interface.

Figure 10 on page 30 illustrates this example in which you bind ethernet2 to the DMZ zone and ethernet3 to the Untrust zone. You set the management options on each interface to provide access for the specific kinds of administrative traffic. You allow HTTP and Telnet access on ethernet2 for a group of local administrators in the DMZ zone, and SNMP access on ethernet3 for central device monitoring from a remote site. Ethernet2 and ethernet3 each have a manage IP address, to which the administrative traffic is directed. You also set a route directing self-generated SNMP traffic out ethernet3 to the external router at 1.1.1.250.

Figure 10: Setting Management IPs for Multiple Interfaces

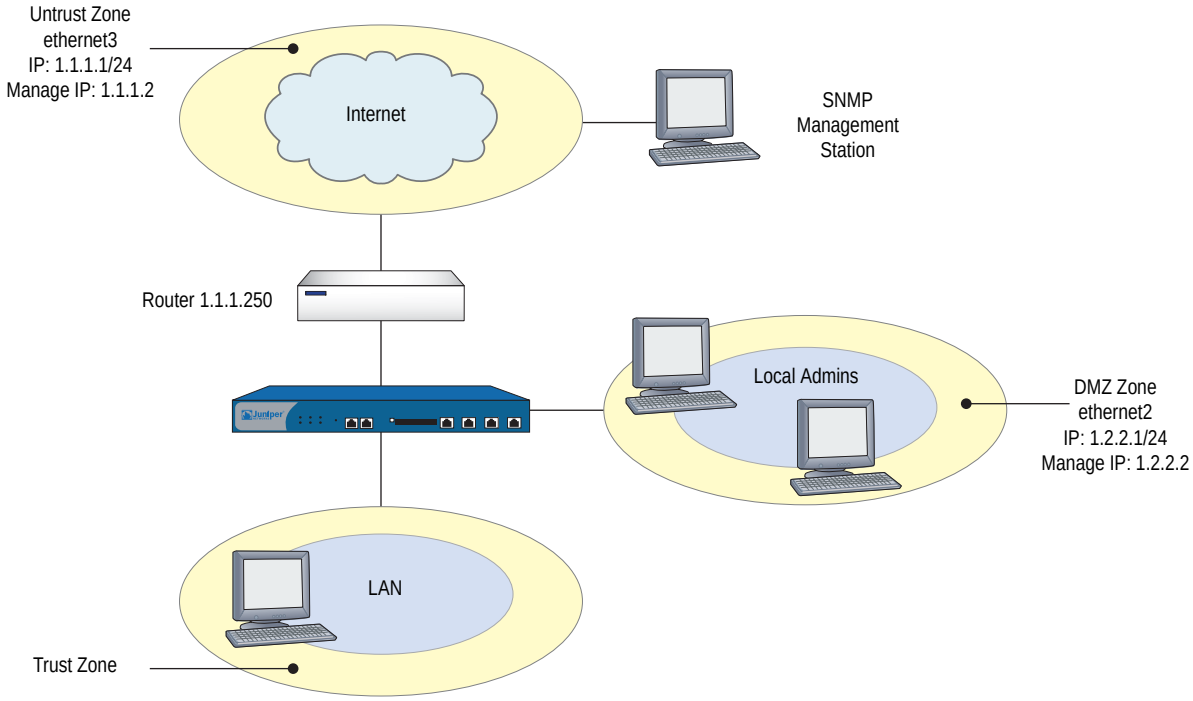

WebUI

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

Zone Name: DMZ
Static IP: (select this option when present)
IP Address/Netmask: 1.2.2.1/24
Manage IP: 1.2.2.2
Management Services:
WebUI: (select)
Telnet: (select)

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

Zone Name: Untrust
Static IP: (select this option when present)
IP Address/Netmask: 1.1.1.1/24
Manage IP: 1.1.1.2
Management Services:
SNMP: (select)

CLI

```
set interface ethernet2 zone dmz
set interface ethernet2 ip 1.2.2.1/24
set interface ethernet2 manage-ip 1.2.2.2
set interface ethernet2 manage web
set interface ethernet2 manage telnet
set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
set interface ethernet3 manage-ip 1.1.1.2
set interface ethernet3 manage snmp
save
```

Levels of Administration

Juniper Networks security devices support multiple administrative users. For any configuration changes made by an administrator, the security device logs the following information:

- The name of the administrator making the change
- The IP address from which the change was made
- The time of the change

There are several levels of administrative user. The availability of some of these levels depends on the model of your Juniper Networks security device. The following sections list all the admin levels and the privileges for each level. These privileges are only accessible to an admin after he or she successfully logs in with a valid username and password.

Root Administrator

The root administrator has complete administrative privileges. There is only one root administrator per security device. The root administrator has the following privileges:

- Manages the root system of the security device
- Adds, removes, and manages all other administrators
- Establishes and manages virtual systems, and assigns physical or logical interfaces to them
- Creates, removes, and manages virtual routers (VRs)
- Adds, removes, and manages security zones
- Assigns interfaces to security zones
- Performs asset recovery
- Sets the device to FIPS mode
- Resets the device to its default settings

- Updates the firmware
- Loads configuration files
- Clears all active sessions of a specified admin or of all active admins

Read/Write Administrator

The read/write administrator has the same privileges as the root administrator, but cannot create, modify, or remove other admin users. The read/write administrator has the following privileges:

- Creates virtual systems and assigns a virtual system administrator for each one
- Monitors any virtual system
- Tracks statistics (a privilege that cannot be delegated to a virtual system administrator)

Read-Only Administrator

The read-only administrator has only viewing privileges using the WebUI, and can only issue the **get** and **ping** CLI commands. The read-only administrator has the following privileges:

- Read-only privileges in the root system, using the following four commands: **enter**, **exit**, **get**, and **ping**
- Read-only privileges in virtual systems

Virtual System Administrator

Some security devices support virtual systems. Each virtual system (vsys) is a unique security domain, which can be managed by virtual system administrators with privileges that apply only to that vsys. Virtual system administrators independently manage virtual systems through the CLI or WebUI. On each vsys, the virtual system administrator has the following privileges:

- Creates and edits auth, IKE, L2TP, XAuth, and Manual Key users
- Creates and edits services
- Creates and edits policies
- Creates and edits addresses
- Creates and edits VPNs
- Modifies the virtual system administrator login password
- Creates and manages security zones
- Adds and removes virtual system read-only administrators

Virtual System Read-Only Administrator

A virtual system read-only administrator has the same set of privileges as a read-only administrator, but only within a specific virtual system. A virtual system read-only administrator has viewing privileges for his particular vsys through the WebUI, and can only issue the **enter**, **exit**, **get**, and **ping** CLI commands within his vsys.

NOTE: For more information on virtual systems, see "Virtual Systems" on page **10**-1.

Defining Admin Users

The root administrator is the only one who can create, modify, and remove admin users. In the following example, the one performing the procedure must be a root administrator.

Example: Adding a Read-Only Admin

In this example, you—as the root admin—add a read-only administrator named Roger with password 2bd21wG7.

WebUI

Configuration > Admin > Administrators > New: Enter the following, then click **OK**:

Name: Roger
New Password: 2bd21wG7
Confirm New Password: 2bd21wG7
Privileges: Read-Only (select)

NOTE: The password can be up to 31 characters long and is case sensitive.

CLI

```
set admin user Roger password 2bd21wG7 privilege read-only
save
```

Example: Modifying an Admin

In this example, you—as the root admin—change Roger's privileges from read-only to read/write.

WebUI

Configuration > Admin > Administrators > Edit (for Roger): Enter the following, then click **OK**:

Name: Roger
New Password: 2bd21wG7
Confirm New Password: 2bd21wG7
Privileges: Read-Write (select)

CLI

```
unset admin user Roger
set admin user Roger password 2bd21wG7 privilege all
save
```

Example: Deleting an Admin

In this example, you—as the root admin—delete the admin user Roger.

WebUI

Configuration > Admin > Administrators: Click **Remove** in the Configure column for Roger.

CLI

```
unset admin user Roger
save
```

Example: Configuring Admin Accounts for Dialup Connections

The NS-5XT and the NS-5GT devices support a modem connection for outbound dial-up disaster recovery situations. You can set up trustee accounts for the interface, for the modem or for both the interface and modem. This section describes the two types of trustees:

- Interface trustee

  An interface trustee only has access to the WebUI and is restricted to the signaling methods and IP address assignment for the primary Untrust interface.

  For devices with ADSL interfaces, an interface trustee has control over the following characteristics:

  - Layer 1 characteristics: VPI/VCI, multiplexing mode, RFC1483 bridged or routed
  - Layer 2 signaling methods (PPPoE or PPPoA, and their parameters)
  - IP address assignment methods (statically defined by an administrator, or dynamically acquired from the circuit through PPPoE or PPPoA).

  For devices with only ethernet interfaces, an interface trustee can control how the interface IP address is assigned (statically defined by administrator, or dynamically acquired from the circuit via DHCP or PPPoE).

- Modem trustee

  A modem trustee only has access to the WebUI, and is restricted to Modem and ISP settings for the serial interface. A modem trustee can created, modify, and delete modem definitions to suit their specific needs, and can create, modify, and delete the settings for ISP1 and ISP2. A modem trustee can view the configurations for ISP3 and ISP4, and can test connectivity for any defined ISP and phone number.

You can view all administrator accounts by entering the **get admin user** command, or you can view only the trustee accounts by entering the **get admin user trustee** command.

In the following example, you configure a Read/Write modem trustee account for Richard Brockie. You set his user name to be *rbrockie* and his password to be *!23fb*.

WebUI

Configuration > Admin > Administrators

CLI

```
Set admin user rbrockie password !23fb privilege all
Set admin user rbrockie trustee modem
```

Example: Clearing an Admin's Sessions

In this example, you—as the root admin—terminate all active sessions of the admin user Roger. When you execute the following command, the security device closes all active sessions and automatically logs off Roger from the system.

WebUI

NOTE: You must use the CLI to clear an admin's sessions.

CLI

```
clear admin name Roger
save
```

Securing Administrative Traffic

To secure the security device during setup, perform the following steps:

1. On the WebUI, change the administrative port.

   See "Changing the Port Number" on page 36.

2. Change the username and password for administration access.

   See "Changing the Admin Login Name and Password" on page 37.

3. Define the management client IP addresses for the admin users.

   See "Restricting Administrative Access" on page 40.

4. Turn off any unnecessary interface management service options.

   See "Controlling Administrative Traffic" on page 26.

5. Disable the ping and ident-reset service options on the interfaces, both of which respond to requests initiated by unknown parties and can reveal information about your network:

WebUI

Network > Interfaces > Edit (for the interface you want to edit): Disable the following service options, then click **OK**:

Ping: Selecting this option allows the security device to respond to an ICMP echo request, or "ping," which determines whether a specific IP address is accessible from the device.

Ident-Reset: When a service such as Mail or FTP sends an identification request and receives no acknowledgment, it sends the request again. While the request is in progress, user access is disabled. With the Ident-Reset checkbox enabled, the security device automatically restores user access.

CLI

```
unset interface interface manage ping
unset interface interface manage ident-reset
```

Changing the Port Number

Changing the port number to which the security device listens for HTTP management traffic improves security. The default setting is port 80, the standard port number for HTTP traffic. After you change the port number, you must then enter the new port number in the URL field in your browser when you next attempt to contact the security device. (In the following example, the administrator needs to enter http://188.30.12.2:15522.)

In this example, the IP address of the interface bound to the Trust zone is 10.1.1.1/24. To manage the security device via the WebUI on this interface, you must use HTTP. To increase the security of the HTTP connection, you change the HTTP port number from 80 (the default) to 15522.

WebUI

Configuration > Admin > Management: In the HTTP Port field, enter **15522**, then click **Apply**.

CLI

```
set admin port 15522
save
```

Changing the Admin Login Name and Password

By default, the initial login name for security devices is **netscreen**. The initial password is also **netscreen**. Because these have been widely published, Juniper Networks recommends you change the login name and password immediately. The login name and password are both case-sensitive. They can contain any character that can be entered from the keyboard with the exception of ? and ". Record the new admin login name and password in a secure manner.

WARNING: Be sure to record your new password. If you forget it, you must reset the security device to its factory settings, and all your configurations will be lost. For more information, see "Resetting the Device to the Factory Default Settings" on page 39.

Admin users for the security device can be authenticated using the internal database or an external auth server. When the admin user logs into the security device, it first checks the local internal database for authentication. If there is no entry present and an external auth server is connected, it then checks for a matching entry in the external auth server database. After an admin user successfully logs into an external auth server, the security device maintains the admin's login status locally.

NOTE: Juniper Networks supports RADIUS, SecurID, and LDAP servers for admin user authentication. (For more information, see "Admin Users" on page **9**-2.) Although the root admin account must be stored on the local database, you can store root-level read/write and root-level read-only admin users on an external auth server. To store root-level and vsys-level admin users on an external auth server and query their privileges, the server must be RADIUS and you must load the netscreen.dct file on it.

For more information about admin user levels, see "Levels of Administration" on page 31. For more about using external auth servers, see "External Authentication Servers" on page **9**-15.

When the root admin changes any attribute of an admin user's profile—username, password, or privilege—any administrative session that the admin currently has open automatically terminates. If the root admin changes any of these attributes for himself, or if a root-level read/write admin or vsys read/write admin changes his own password, all of that user's currently open admin sessions terminate, other than the one in which he made the change.

NOTE: The behavior of an HTTP or HTTPS session using the WebUI is different. Because HTTP does not support a persistent connection, any change that you make to your own user profile automatically logs you out of that and all other open sessions.

Example: Changing an Admin User's Login Name and Password

In this example, you—as the root admin—change a read/write administrator's login name from "John" to "Smith" and his password from xL7s62a1 to 3MAb99j2.

NOTE: Instead of using actual words for passwords, which might be guessed or discovered through a dictionary attack, you can use an apparently random string of letters and numbers. To create such a string that you can easily remember, compose a sentence and use the first letter from each word. For example, "Charles will be 6 years old on November 21" becomes "Cwb6yooN21."

For more information, see "Levels of Administration" on page 31.

WebUI

Configuration > Admin > Administrators > Edit (for John): Enter the following, then click **OK**:

Name: Smith
New Password: 3MAb99j2
Confirm New Password: 3MAb99j2

CLI

```
unset admin user John
set admin user Smith password 3MAb99j2 privilege all
save
```

Example: Changing Your Own Password

Admin users with read/write privileges can change their own administrator password, but not their login name. In this example, an administrator with read/write privileges and the login name "Smith" changes his password from 3MAb99j2 to ru494Vq5.

WebUI

Configuration > Admin > Administrators > Edit (for first entry): Enter the following, then click **OK**:

Name: Smith
New Password: ru494Vq5
Confirm New Password: ru494Vq5

CLI

```
set admin password ru494Vq5
save
```

Setting the Minimum Length of the Root Admin Password

In some corporations, one person might initially configure the device as the root admin, but another person later assumes the role of root admin and manages the device. To prevent the subsequent root admin from using short passwords that are potentially easier to decode, the initial root admin can set a minimum length requirement for the root admin's password to any number from 1 to 31.

You can set the minimum password length only if you are the root admin and your own password meets the minimum length requirement you are attempting to set. Otherwise, the security device displays an error message.

To specify a minimum length for the root admin's password, enter the following CLI command:

set admin password restrict length *number*

NOTE: You must use the CLI to set this restriction.

Resetting the Device to the Factory Default Settings

If the admin password is lost, you can use the following procedure to reset the security device to its default settings. The configurations will be lost, but access to the device will be restored. To perform this operation, you need to make a console connection, which is described in detail in *ScreenOS CLI Reference Guide: IPv4 Command Descriptions* and the documentation for your device.

NOTE: By default, the device recovery feature is enabled. You can disable it by entering the **unset admin device-reset** command. Also, if the security device is in FIPS mode, the recovery feature is automatically disabled.

---

1. At the login prompt, enter the serial number of the device.
2. At the password prompt, enter the serial number again.

   The following message appears:

   !!!! Lost Password Reset !!!! You have initiated a command to reset the device to factory defaults, clearing all current configuration, keys and settings. Would you like to continue? y/n

3. Press the **y** key.

   The following message appears:

   !! Reconfirm Lost Password Reset !! If you continue, the entire configuration of the device will be erased. In addition, a permanent counter will be incremented to signify that this device has been reset. This is your last chance to cancel this command. If you proceed, the device will return to factory default configuration, which is: System IP: 192.168.1.1; username: netscreen; password: netscreen. Would you like to continue? y/n

4. Press the **y** key to reset the device. You can now log in using **netscreen** as the default username and password.

Restricting Administrative Access

You can administer security devices from one or multiple addresses of a subnet. By default, any host on the trusted interface can administer a security device. To restrict this ability to specific workstations, you must configure management client IP addresses.

NOTE: The assignment of a management client IP address takes effect immediately. If you are managing the device via a network connection and your workstation is not included in the assignment, the security device immediately terminates your current session and you are no longer able to manage the device from that workstation.

Example: Restricting Administration to a Single Workstation

In this example, the administrator at the workstation with the IP address 172.16.40.42 is the only administrator specified to manage the security device.

WebUI

Configuration > Admin > Permitted IPs: Enter the following, then click **Add**:

IP Address / Netmask: 172.16.40.42/32

CLI

```
set admin manager-ip 172.16.40.42/32
save
```

Example: Restricting Administration to a Subnet

In this example, the group of administrators with workstations in the 172.16.40.0/24 subnet are specified to manage a security device.

WebUI

Configuration > Admin > Permitted IPs: Enter the following, then click **Add**:

IP Address / Netmask: 172.16.40.0/24

CLI

```
set admin manager-ip 172.16.40.0 255.255.255.0
save
```

Restricting the Root Admin to Console Access

You can also require the root admin to log into the security device through the console only. This restriction requires the root admin to have physical access to the device to log in, thus preventing unauthorized users from logging in remotely as the root admin. After you have set this restriction, the device denies access if anyone tries to log in as the root admin through other means, such as the WebUI, Telnet, or SSH, even if these management options are enabled on the ingress interface.

To restrict the access of the root admin to the console only, enter the following command:

set admin root access console

NOTE: You must use the CLI to set this restriction.

VPN Tunnels for Administrative Traffic

You can use virtual private network (VPN) tunnels to secure remote management of a security device from either a dynamically assigned or fixed IP address. Using a VPN tunnel, you can protect any kind of traffic, such as NetScreen-Security Manager, HTTP, Telnet, or SSH. (For information about creating a VPN tunnel to secure self-initiated traffic such as NetScreen-Security Manager reports, syslog reports, or SNMP traps, see "VPN Tunnels for Self-Initiated Traffic" on page 75.)

Juniper Networks security devices support two types of VPN tunnel configurations:

- **Route-based VPNs**: The security device uses route table entries to direct traffic to tunnel interfaces, which are bound to VPN tunnels.
- **Policy-based VPNs**: The security device uses the VPN tunnel names specifically referenced in policies to direct traffic through VPN tunnels.

For each VPN tunnel configuration type, there are the following types of VPN tunnel:

- **Manual key**: You manually set the three elements that define a Security Association (SA) at both ends of the tunnel: a Security Parameters Index (SPI), an encryption key, and an authentication key. To change any element in the SA, you must manually enter it at both ends of the tunnel.
- **AutoKey IKE with pre-shared key**: One or two pre-shared secrets—one for authentication and one for encryption—function as seed values. Using them, the IKE protocol generates a set of symmetrical keys at both ends of the tunnel; that is, the same key is used to encrypt and decrypt. At predetermined intervals, these keys are automatically regenerated.
- **AutoKey IKE with certificates**: Using the Public Key Infrastructure (PKI), the participants at both ends of the tunnel use a digital certificate (for authentication) and an RSA public/private key pair (for encryption). The encryption is asymmetrical; that is, one key in a pair is used to encrypt and the other to decrypt.

NOTE: For a complete description of VPN tunnels, see *Volume 5: Virtual Private Networks*. For more information on NetScreen-Remote, refer to the *NetScreen-Remote VPN Client Administrator Guide.*

If you use a policy-based VPN configuration, you must create an address book entry with the IP address of an interface in any zone other than the one to which the outgoing interface is bound. You can then use that as the source address in policies referencing the VPN tunnel. This address also serves as the end entity address for the remote IPSec peer. If you are using a route-based VPN configuration, such an address book entry is unnecessary.

Administration Through a Route-Based Manual Key VPN Tunnel

Figure 11 illustrates an example in which you set up a route-based Manual Key VPN tunnel to provide confidentiality for administrative traffic. The tunnel extends from the NetScreen-Remote VPN client running on an admin's workstation at 10.1.1.56 to ethernet1 (10.1.1.1/24). The admin's workstation and ethernet1 are both in the Trust zone. You name the tunnel "tunnel-adm". You create an unnumbered tunnel interface, name it tunnel.1, and bind it to the Trust zone and to the VPN tunnel "tunnel-adm."

The security device uses the internal IP address configured on the NetScreen-Remote client—10.10.10.1—as the destination address to target beyond the peer gateway address of 10.1.1.56. You define a route to 10.10.10.1/32 through tunnel.1. A policy is unnecessary because of the following two reasons:

- The VPN tunnel protects administrative traffic that terminates at the security device itself instead of passing through the device to another security zone.
- This is a route-based VPN, meaning that the route lookup—not a policy lookup—links the destination address to the tunnel interface, which is bound to the appropriate VPN tunnel.

NOTE: Compare this example with "Administration Through a Policy-Based Manual Key VPN Tunnel" on page 45.

NetScreen-Remote uses the IP address of ethernet3—1.1.1.1—as the destination address to target beyond the remote gateway at 10.1.1.1. The NetScreen-Remote configuration specifies the remote party ID type as "IP address" and the protocol as "All."

Figure 11: Administration Through a Route-Based Manual Key VPN Tunnel

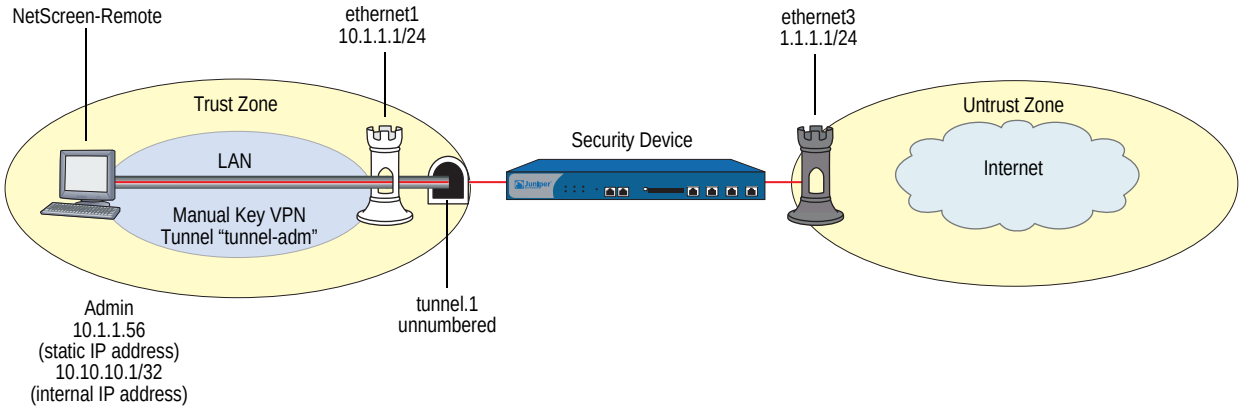

WebUI

1. **Interfaces**

   Network > Interfaces > Edit (ethernet1): 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 (ethernet3): Enter the following, then click **OK**:

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

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

   Tunnel Interface Name: Tunnel.1
   Zone (VR): Trust (trust-vr)
   Unnumbered: (select)
   Interface: ethernet1(trust-vr)

NOTE: The unnumbered tunnel interface borrows the IP address of the specified security zone interface.

---

2. **VPN**

   VPNs > Manual Key > New: Enter the following, then click **OK**:

   VPN Tunnel Name: tunnel-adm
   Gateway IP: 10.1.1.56
   Security Index (HEX Number): 5555 (Local) 5555 (Remote)
   Outgoing Interface: ethernet1
   ESP-CBC: (select)
   Encryption Algorithm: DES-CBC
   Generate Key by Password: netscreen1
   Authentication Algorithm: MD5
   Generate Key by Password: netscreen2

   > Advanced: Enter the following, then click **Return** to set the advanced options and return to the basic configuration page:

   Bind to Tunnel Interface: (select), Tunnel.1

NOTE: Because NetScreen-Remote processes passwords into keys differently than do other Juniper Networks products, after you configure the tunnel you need to do the following: (1) Return to the Manual Key Configuration dialog box (click **Edit** in the Configure column for "tunnel-adm"); (2) copy the generated hexadecimal keys; (3) use those hexadecimal keys when configuring the NetScreen-Remote end of the tunnel.

---

3. Route

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

Network Address/Netmask: 10.10.10.1/32
Gateway: (select)
Interface: Tunnel.1
Gateway IP Address: 0.0.0.0

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat
set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
set interface tunnel.1 zone trust
set interface tunnel.1 ip unnumbered interface ethernet1
```

NOTE: The unnumbered tunnel interface borrows the IP address of the specified security zone interface.

2. VPN

```
set vpn tunnel-adm manual 5555 5555 gateway 10.1.1.56 outgoing ethernet1
    esp des password netscreen1 auth md5 password netscreen2
set vpn tunnel-adm bind interface tunnel.1
```

NOTE: Because NetScreen-Remote processes passwords into keys differently than do other Juniper Networks products, after you configure the tunnel you need to do the following: (1) Enter **get vpn admin-tun**; (2) copy the hexadecimal keys generated by "netscreen1" and "netscreen2"; (3) use those hexadecimal keys when configuring the NetScreen-Remote end of the tunnel.

3. Route

```
set vrouter trust-vr route 10.10.10.1/32 interface tunnel.1
save
```

NetScreen-Remote Security Policy Editor

1. Click **Options > Global Policy Settings**, and select the Allow to Specify Internal Network Address checkbox.
2. Click **Options > Secure > Specified Connections**.
3. Click **Add a new connection**, and enter **Admin** next to the new connection icon that appears.
4. Configure the connection options:

Connection Security: Secure
Remote Party Identity and Addressing:
ID Type: IP Address, 1.1.1.1
Protocol: All
Connect using Secure Gateway Tunnel: (select)
ID Type: IP Address, 10.1.1.1

5. Click the **PLUS** symbol, located to the left of the unix icon, to expand the connection policy.

6. Click **My Identity**, in the Select Certificate dropdown list, choose **None**, and in the Internal Network IP Address, enter **10.10.10.1**.

7. Click **Security Policy**, and select **Use Manual Keys**.

8. Click the **PLUS** symbol, located to the left of the Security Policy icon, and then click the **PLUS** symbol to the left of Key Exchange (Phase 2) to expand the policy further.

9. Click **Proposal 1**, and select the following IPSec Protocols:

   Encapsulation Protocol (ESP): (select)
   Encrypt Alg: DES
   Hash Alg: MD5
   Encapsulation: Tunnel

10. Click **Inbound Keys**, and in the Security Parameters Index field, enter **5555**.

11. Click **Enter Key**, enter the following, then click **OK**:

    Choose key format: Binary
    ESP Encryption Key: dccbee96c7e546bc
    ESP Authentication Key: dccbe9e6c7e546bcb0b667794ab7290c

NOTE: These are the two generated keys that you copied after configuring the security device.

---

12. Click **Outbound Keys**, and, in the Security Parameters Index field, enter **5555**.

13. Click **Enter Key**, enter the following, then click **OK**:

    Choose key format: Binary
    ESP Encryption Key: dccbee96c7e546bc
    ESP Authentication Key: dccbe9e6c7e546bcb0b667794ab7290c

14. Click **Save**.

Administration Through a Policy-Based Manual Key VPN Tunnel

Figure 12 illustrates an example in which you set up a policy-based Manual Key VPN tunnel for administrative traffic. The tunnel extends from the NetScreen-Remote VPN client running on an admin's workstation at 10.1.1.56 to ethernet1 (10.1.1.1/24). The admin's workstation and ethernet1 are both in the Trust zone. You name the tunnel "tunnel-adm" and bind it to the Trust zone.

The security device uses the internal IP address configured on the NetScreen-Remote—10.10.10.1—as the destination address to target beyond the peer gateway address of 10.1.1.56. You define a Trust zone address book entry specifying 10.10.10.1/32, and an Untrust zone address book entry specifying the IP address of ethernet3. Although the address of the ethernet3 interface is 1.1.1.1/24, the address you create has a 32-bit netmask: 1.1.1.1/32. You use this address and

the internal address of the admin's workstation in the policy you create referencing the tunnel "tunnel-adm". A policy is necessary because this is a policy-based VPN, meaning that the policy lookup—not a route lookup—links the destination address to the appropriate VPN tunnel.

You must also define a route to 10.10.10.1/32 through ethernet1.

NOTE: Compare this example with "Administration Through a Route-Based Manual Key VPN Tunnel" on page 42.

NetScreen-Remote uses the IP address 1.1.1.1 as the destination address to target beyond the remote gateway at 10.1.1.1. The NetScreen-Remote tunnel configuration specifies the remote party ID type as IP address and the protocol as "All."

Figure 12: Administration Through a Policy-Based Manual Key VPN Tunnel

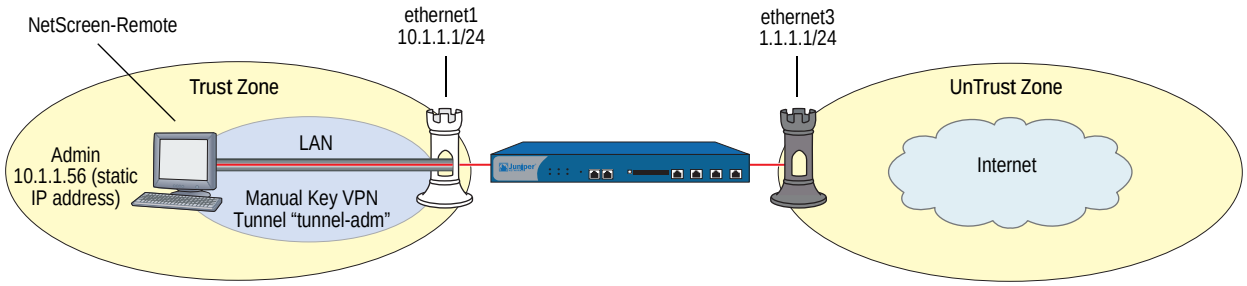

WebUI

1. **Interfaces**
   Network > Interfaces > Edit (ethernet1): 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 (ethernet3): 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**
   Objects > Addresses > List > New: Enter the following, then click **OK**:

   Address Name: Untrust-IF
   IP Address/Domain Name:
   IP/Netmask: (select), 1.1.1.1/32
   Zone: Untrust

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

Address Name: admin
IP Address/Domain Name:
    IP/Netmask: (select), 10.10.10.1/32
Zone: Trust

3. **VPN**
   VPNs > Manual Key > New: Enter the following, then click **OK**:

   VPN Tunnel Name: tunnel-adm
   Gateway IP: 10.1.1.56
   Security Index (HEX Number): 5555 (Local) 5555 (Remote)
   Outgoing Interface: ethernet1
   ESP-CBC: (select)
   Encryption Algorithm: DES-CBC
   Generate Key by Password: netscreen1
   Authentication Algorithm: MD5
   Generate Key by Password: netscreen2

NOTE: Because NetScreen-Remote processes passwords into keys differently than do other Juniper Networks products, after you configure the tunnel you need to do the following: (1) Return to the Manual Key Configuration dialog box (click **Edit** in the Configure column for "tunnel-adm"); (2) copy the generated hexadecimal keys; (3) use those hexadecimal keys when configuring the NetScreen-Remote end of the tunnel.

---

4. **Route**
   Network > Routing > Routing Entries > trust-vr New: Enter the following, then click **OK**:

   Network Address/Netmask: 10.10.10.1/32
   Gateway: (select)
       Interface: ethernet1
       Gateway IP Address: 0.0.0.0

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

   Source Address:
       Address Book Entry: (select), admin
   Destination Address:
       Address Book Entry: (select), Untrust-IF
   Service: Any
   Action: Tunnel
   Tunnel:
       VPN: tunnel-adm
   Modify matching bidirectional VPN policy: (select)
   Position at Top: (select)

CLI

1. **Interfaces**
   set interface ethernet1 zone trust
   set interface ethernet1 ip 10.1.1.1/24
   set interface ethernet1 nat
   set interface ethernet3 zone untrust
   set interface ethernet3 ip 1.1.1.1/24

2. **Addresses**
   set address trust admin 10.10.10.1/32
   set address untrust Untrust-IF 1.1.1.1/32

3. **VPN**
   set vpn tunnel-adm manual 5555 5555 gateway 10.1.1.56 outgoing ethernet1
   esp des password netscreen1 auth md5 password netscreen2

NOTE: Because NetScreen-Remote processes passwords into keys differently than do other Juniper Networks products, after you configure the tunnel you need to do the following: (1) Enter **get vpn admin-tun**; (2) copy the hexadecimal keys generated by "netscreen1" and "netscreen2"; (3) use those hexadecimal keys when configuring the NetScreen-Remote end of the tunnel.

---

4. **Route**
   set vrouter trust-vr route 10.10.10.1/32 interface ethernet1

5. **Policies**
   set policy top from trust to untrust admin Untrust-IF any tunnel vpn tunnel-adm
   set policy top from untrust to trust Untrust-IF admin any tunnel vpn tunnel-adm
   save

NetScreen-Remote Security Policy Editor

1. Click **Options > Secure > Specified Connections**.

2. Click **Add a new connection**, and enter **Admin** next to the new connection icon that appears.

3. Configure the connection options:

   Connection Security: Secure
   Remote Party Identity and Addressing:
   ID Type: IP Address, 1.1.1.1
   Protocol: All
   Connect using Secure Gateway Tunnel: (select)
   ID Type: IP Address, 10.1.1.1

4. Click the **PLUS** symbol, located to the left of the Unix icon, to expand the connection policy.

5. Click **My Identity**, and in the **Select Certificate** drop-down list, choose **None**.

6. Click **Security Policy**, and select **Use Manual Keys**.

7. Click the **PLUS** symbol, located to the left of the Security Policy icon, and then the **PLUS** symbol to the left of Key Exchange (Phase 2) to expand the policy further.

8. Click **Proposal 1**, and select the following IPSec Protocols:

   Encapsulation Protocol (ESP): (select)
   Encrypt Alg: DES
   Hash Alg: MD5
   Encapsulation: Tunnel

9. Click **Inbound Keys**, and in the Security Parameters Index field, enter **5555**.

10. Click **Enter Key**, enter the following, and click **OK**:

    Choose key format: Binary
    ESP Encryption Key: dccbee96c7e546bc
    ESP Authentication Key: dccbe9e6c7e546bcb0b667794ab7290c

NOTE: These are the two generated keys that you copied after configuring the security device.

11. Click **Outbound Keys**, and in the Security Parameters Index field, enter **5555**.

12. Click **Enter Key**, enter the following, then click **OK**:

    Choose key format: Binary
    ESP Encryption Key: dccbee96c7e546bc
    ESP Authentication Key: dccbe9e6c7e546bcb0b667794ab7290c

13. Click **Save**.

Password Policy

The password policy feature allows you to enforce a minimum length and a complexity scheme for administrator (admin) and authenticated (auth) user passwords. The password policy feature is intended for use in a local database, and therefore is useful in environments where the Windows directory or RADIUS are not available to provide centralized password policy enforcement.

Setting a Password Policy

You can create a password policy to require that admin and auth passwords fulfill one or both of the following:

- Minimum length
- Complexity

The range for password minimum length is 1 to 32 characters. Use the following command to create a password policy requiring a minimum length of 8 characters for admin passwords:

```
set password-policy user-type admin minimum-length 8
```

Password complexity means passwords must include at least two uppercase letters, two lowercase letters, and two alphanumeric and two non-alphanumeric characters; for example: AAbb12@#. To require that passwords contain complexity, you set complexity to 1. To unset the complexity requirement, set complexity to 0. Use the following command to create a password policy requiring that auth passwords contain complexity:

```
set password-policy user-type auth complexity 1
```

In the following example, you create a password policy for admin and auth accounts requiring complexity and a minimum length of 8 characters:

CLI

```
set password-policy user-type admin minimum-length 8
set password-policy user-type admin complexity 1
set password-policy user-type auth minimum-length 8
set password-policy user-type auth complexity 1
save
```

NOTE: You can configure a password policy only from the command line interface (CLI).

Removing a Password Policy

Use the **unset password-policy** command to delete a password policy. When you remove a password policy, the password requirement for the account reverts to the default settings. In the following example, you remove the minimum length requirement for auth passwords.

CLI

```
unset password-policy user-type auth minimum-length
```

Viewing a Password Policy

Use the **get password-policy** command to display the password policy for admin and auth users.

Recovering from a Rejected Default Admin Password

When you delete (unset) the root admin account on a device on which you have a password policy configured, you might need to set a new admin password before logging off the system. This is because ScreenOS reverts to the default password (*netscreen*) when you delete the root admin account. If you have a password policy requiring complexity, or a minimum length greater than 9 characters, your next login attempt will fail. If this happens, use the asset recovery procedure to gain access to the device. Refer to the user's guide for your platform for details.

In the following example, you delete the admin account named **admin2005**, then display the current password policy. As shown, the policy specifies that passwords must have a minimum length of 8 characters, and use complexity (a minimum of two uppercase, two lowercase, two alphanumeric, and two non-alphanumeric characters). You then create a new admin account named **admin2006** and set a password for it that fulfills the minimum length and complexity requirements of the password policy.

CLI

```
unset admin admin2005
get password-policy

user-type: admin
password minimum length: 8
password complexity scheme: 1

user-type: auth
password minimum length: 8
password complexity scheme: 1

set admin admin2006 password AAbb12@#
save
```

NOTE: You can configure an admin account only from the command line interface (CLI).

Chapter 2
Monitoring Security Devices

This chapter discusses the following topics about monitoring Juniper Networks security devices:

Storing Log Information

All Juniper Networks security devices allow you to store event and traffic log data internally (in flash storage) and externally (in a number of locations). Although storing log information internally is convenient, the amount of device memory is limited. When the internal storage space completely fills up, the security device begins overwriting the oldest log entries with the latest ones. If this first-in-first-out (FIFO) mechanism occurs before you save the logged information, you can lose data. To mitigate such data loss, you can store event and traffic logs externally in a syslog or WebTrends server or in the NetScreen-Global PRO database. The security device sends new event and traffic log entries to an external storage location every second.

The following list provides the possible destinations for logged data:

- **Console:** A destination for all log entries to appear when you are troubleshooting a security device through the console. Optionally, you might elect to have only alarm messages (critical, alert, emergency) appear here to alert you immediately if you happen to be using the console at the time an alarm is triggered.
- **Internal:** Allows you store a limited number of log entries.
- **Email:** A method for sending event and traffic logs to remote administrators.
- **SNMP:** In addition to the transmission of SNMP traps, a security device can also send alarm messages (critical, alert, emergency) from its event log to an SNMP community.
- **Syslog:** All event and traffic log entries that a security device can store internally, it can also send to a syslog server. Because syslog servers have a much greater storage capacity than the internal flash storage on a security device, sending data to a syslog server can mitigate data loss that might occur when log entries exceed the maximum internal storage space. Syslog stores alert- and emergency-level events in the security facility that you specify, and all other events (including traffic data) in the facility you specify.
- **WebTrends:** Allows you to view log data for critical-, alert-, and emergency-level events in a more graphical format than syslog, which is a text-based tool.
- **CompactFlash (PCMCIA):** Allows you to store data on a CompactFlash card.

Event Log

ScreenOS provides an event log for monitoring system events such as admin-generated configuration changes, and self-generated messages and alarms regarding operational behavior and attacks. The security device categorizes system events by the following severity levels:

- **Emergency:** Messages on SYN attacks, Tear Drop attacks, and Ping of Death attacks. For more information on these types of attacks, see *Volume 4: Attack Detection and Defense Mechanisms*.
- **Alert:** Messages about conditions that require immediate attention, such as firewall attacks and the expiration of license keys.
- **Critical:** Messages about conditions that probably affect the functionality of the device, such as high availability (HA) status changes.
- **Error:** Messages about error conditions that probably affect the functionality of the device, such as a failure in antivirus scanning or in communicating with SSH servers.
- **Warning:** Messages about conditions that could affect the functionality of the device, such as a failure to connect to email servers or authentication failures, timeouts, and successes.

- **Notification:** Messages about normal events, including configuration changes initiated by an admin.
- **Information:** Messages that provide general information about system operations.
- **Debugging:** Messages that provide detailed information used for debugging purposes.

The event log displays the date, time, level and description of each system event. You can view system events for each category stored in flash storage on the security device through the WebUI or the CLI. You can also open or save the file to the location you specify, and then use an ASCII text editor (such as Notepad or WordPad) to view the file. Alternatively, you can send them to an external storage space (see "Storing Log Information" on page 53).

NOTE: For detailed information about the messages that appear in the event log, refer to the *ScreenOS Message Log Reference Guide*.

Viewing the Event Log by Severity Level and Keyword

You can view the event log stored in the device by using the CLI or the WebUI. You can display log entries by severity level and search the event log by keyword in both the WebUI and CLI.

To display the event log by severity level, do either of the following:

WebUI

Reports > System Log > Event: Select a severity level from the Log Level drop-down list.

CLI

```
get event level { emergency | alert | critical | error | warning | notification |
    information | debugging }
```

To search the event log by keyword, do either of the following:

WebUI

Reports > System Log > Event: Enter a word or phrase up to 15 characters in length in the search field, then click **Search**.

CLI

```
get event include word_string
```

In this example, you view event log entries with a "warning" severity level and do a search for the keyword AV.

WebUI

Reports > System Log > Event:

Log Level: Warning (select)

Search: Enter AV, then click **Search**.

CLI

get event level warning include av

Date	Time	Module Level	Level	Type	Description
2003-05-16	15:56:20	system warn	00547	AV	scanman is removed.
2003-05-16	09:45:52	system warn	00547	AV	test1 is removed.
Total entries matched = 2					

Sorting and Filtering the Event Log

Additionally, you can use the CLI to sort or filter the event log based on the following criteria:

- **Source or Destination IP Address:** Only certain events contain a source or destination IP address, such as land attacks or ping flood attacks. When you sort event logs by source or destination IP address, the device sorts and displays only the event logs that contain source or destination IP addresses. It ignores all event logs with no source or destination IP address.

  When you filter the event log by specifying a source or destination IP address or range of addresses, the device displays the log entries for the specified source or destination IP address, or range of addresses.

- **Date:** You can sort the event log by date only, or by date and time. When you sort log entries by date and time, the device lists the log entries in descending order by date and time.

  You can also filter event log entries by specifying a start date, an end date, or a date range. When you specify a start date, the device displays log entries with date/time stamps after the start date. When you specify an end date, the device displays log entries with date/time stamps before the end date.

- **Time:** When you sort logs by time, the device displays the log entries in descending order by time, regardless of the date. When you specify a start time, the device displays log entries with time stamps after the specified start time, regardless of the date. When you specify an end time, the device displays log entries with time stamps before the specified end time, regardless of the date. When you specify both a start and end time, the device displays log entries with time stamps within the specified time period.

- **Message Type ID Number:** You can display event log entries for a specific message type ID number, or you can display log entries with message type ID numbers within a specified range. The device displays log entries with the message type ID number(s) you specified, in descending order by date and time.

In this example you view event log entries that contain source IP addresses within the range 10.100.0.0 to 10.200.0.0. The log entries are also sorted by source IP address.

WebUI

NOTE: You must use the CLI to sort the event log by address entries.

CLI

```
get event sort-by src-ip 10.100.0.0-10.200.0.0
```

Downloading the Event Log

You can open or save the event log to the location you specify, and then use an ASCII text editor (such as Notepad or WordPad) to view the file. Alternatively, you can send the log entries to an external storage space (see "Storing Log Information" on page 53). You can download the entire event log through the WebUI. You can download the event log by severity level through the CLI.

Example: Downloading the Entire Event Log

In this example, you download the event log to the local directory. Using the WebUI, you download it to *C:\netscreen\logs*. Using the CLI, you download it to the root directory of a TFTP server at the IP address 10.1.1.5. You name the file "evnt07-02.txt."

WebUI

1. Reports > System Log > Event: Click **Save**.

   The File Download dialog box prompts you to open the file (using an ASCII editor) or save it to disk.

2. Select the **Save** option, then click **OK**.

   The File Download dialog box prompts you to choose a directory.

3. Specify **C:\netscreen\logs**, name the file *evnt07-02.txt*, then click **Save**.

CLI

```
get event > tftp 10.1.1.5 evnt07-02.txt
```

Example: Downloading the Event Log for Critical Events

In this example, you download the critical events entered in the event log to the root directory of a TFTP server at the IP address 10.1.1.5. You name the file *crt_evnt07-02.txt*.

WebUI

NOTE: You must use the CLI to download entries by severity level.

CLI

```
get event level critical > tftp 10.1.1.5 crt_evnt07-02.txt
```

Traffic Log

The Juniper Networks security device can monitor and record traffic that it permits or denies based on previously configured policies. You can enable the logging option for each policy that you configure. When you enable the logging option for a policy that permits traffic, the device records the traffic allowed by that policy. When you enable the logging option for a policy that denies traffic, the device records traffic that attempted to pass through the device, but was dropped because of that policy.

A traffic log notes the following elements for each session:

- Date and time that the connection started
- Duration
- Source address and port number
- Translated source address and port number
- Destination address and port number
- The duration of the session
- The service used in the session

To log all traffic that a security device receives, you must enable the logging option for all policies. To log specific traffic, enable logging only on policies that apply to that traffic. To enable the logging option on a policy, do either of the following:

WebUI

Policies > (From: *src_zone*, To: *dst_zone*) New: Select **Logging** and then click **OK**.

CLI

set policy from *src_zone* to *dst_zone src_addr dst_addr service action* log

In addition to logging traffic for a policy, the device can also maintain a count in bytes of all network traffic to which the policy was applied. When you enable the counting option, the device includes the following information when it displays traffic log entries

- Bytes transmitted from a source to a destination
- Bytes transmitted from a destination to a source

You can enable counting on a policy from the WebUI and from the Command Line Interface (CLI).

WebUI

Policies > (From: *src_zone*, To: *dst_zone*) New > Advanced: Select **Counting**, click **Return**, then click **OK**.

CLI

set policy from *src_zone* to *dst_zone src_addr dst_addr service action* log count

Viewing the Traffic Log

You can view traffic log entries stored in flash storage on the security device using either the WebUI or the CLI.

WebUI

Policies > Logging (for policy ID *number*)

or

Reports > Policies > Logging (for policy ID *number*)

CLI

get log traffic policy *number*

Example: Viewing Traffic Log Entries

In this example, you view the traffic log details of a policy with ID number 3, and for which you have previously enabled logging:

WebUI

Policies: Click the Logging icon for the policy with ID number 3.

The following information appears:

- Date/Time: 2003-01-09 21:33:43
- Duration: 1800 sec.
- Source IP Address/Port: 1.1.1.1:1046
- Destination IP Address/Port: 10.1.1.5:80
- Service: HTTP
- Reason for Close: Age out
- Translated Source IP Address/Port: 1.1.1.1:1046
- Translated Destination IP Address/Port: 10.1.1.5:80
- Policy ID number: 3

CLI

get log traffic policy 3

Sorting and Filtering the Traffic Log

Similar to the event log, when you use the CLI to view the traffic log, you can sort or filter the log entries according to the following criteria:

- **Source or Destination IP Address:** You can sort the traffic log by source or destination IP address. You can also filter the traffic log by specifying a source or destination IP address or range of addresses.

- **Date:** You can sort the traffic log by date only, or by date and time. The device lists the log entries in descending order by date and time.

  You can also filter event log entries by specifying a start date, an end date, or a date range. When you specify a start date, the device displays log entries with date/time stamps after the start date. When you specify an end date, the device displays log entries with date/time stamps before the end date.

- **Time:** When you sort the traffic log by time, the device displays the log entries in descending order by time, regardless of the date. When you specify a start time, the device displays log entries with time stamps after the specified start time, regardless of the date. When you specify an end time, the device displays log entries with time stamps before the specified end time, regardless of the date. When you specify both a start and end time, the device displays log entries with time stamps within the specified time period.

Example: Sorting the Traffic Log by Time

In this example you view the traffic log sorted by time with a time stamp after 1:00 a.m.

WebUI

NOTE: The ability to sort the traffic log by time is available only through the CLI.

CLI

```
get log traffic sort-by time start-time 01:00:00
```

Downloading the Traffic Log

You can also open or save the log to the location you specify, and then use an ASCII text editor (such as Notepad or WordPad) to view the file.

Alternatively, you can send traffic log entries to an external storage space (see “Storing Log Information” on page 53). The security device makes an entry in the traffic log when a session terminates. When you enable the security device to send traffic log entries to an external storage location, it sends new entries every second. Because the security device makes a traffic log entry when a session closes, the security device sends traffic log entries for all sessions that have closed within the past second. You can also include traffic log entries with event log entries sent by email to an admin.

In this example, you download the traffic log for a policy with ID number 12. For the WebUI, you download it to the local directory “C:\netscreen\logs”. For the CLI, you download it to the root directory of a TFTP server at the IP address 10.10.20.200. You name the file “traf_log11-21-02.txt.”

WebUI

1. Reports > Policies > Logging (for policy ID 12): Click **Save**.

   The File Download dialog box prompts you to open the file (using an ASCII editor) or save it to disk.

2. Select the **Save** option, then click **OK**.

   The File Download dialog box prompts you to choose a directory.

3. Specify C:\netscreen\logs, name the file traf_log11-21-02.txt, then click **Save**.

CLI

```
get log traffic policy 12 > tftp 10.10.20.200 traf_log11-21-02.txt
```

Removing the Reason for Close Field

By default ScreenOS records and displays the reason for session close so that you can differentiate session creation messages from session close messages. If you do not want the reason to display, you can explicitly configure the device not to display the field.

Table 1 lists the reasons for session close that ScreenOS identifies. Any session that cannot be identified is labeled OTHER.

Table 1: Reason Codes for Session Close

Logged Reason	Meaning
TCP FIN	TCP connection torn down due to FIN packet.
TCP RST	TCP connection torn down due to RST packet.
RESP	Special sessions, such as PING and DNS, close when response is received.
ICMP	ICMP error received.
AGE OUT	Connection aged out normally.
ALG	ALG forced session close either due to error or other reason specific to that ALG.
NSRP	NSRP session close message received.
AUTH	Auth failure.
IDP	Closed by IDP.
SYN PROXY FAIL	SYN Proxy failure.
SYN PROXY LIMIT	System limit for SYN proxy sessions reached.
TENT2NORM CONV	Failure of tentative to normal session conversion.
PARENT CLOSED	Parent session closed.
CLI	User command closed.
OTHER	Reason for close not identified.

Sample traffic log with reason for close listed:

```
ns-> get log traffic
PID 1, from Trust to Untrust, src Any, dst Any, service ANY, action Permit
Total traffic entries matched under this policy = 2300
================================================================================
Date Time Duration Source IP Port Destination IP Port Service
Reason Xlated Src IP Port Xlated Dst IP Port ID
================================================================================
2001-10-25 07:08:51 0:00:59 10.251.10.25 137 172.24.16.10 137 NETBIOS (NS)
Close - AGE OUT 172.24.76.127 8946 172.24.16.10 137
2001-10-25 07:08:51 0:00:59 10.251.10.25 137 172.24.244.10 137 NETBIOS (NS)
Close - AGE OUT 172.24.76.127 8947 172.24.244.10 137
2001-10-25 07:07:53 0:00:01 10.251.10.25 1028 172.24.16.10 53 DNS
Close - RESP 172.24.76.127 8945 172.24.16.10 53
2001-10-25 07:06:29 0:01:00 10.251.10.25 138 172.24.244.10 138 NETBIOS (DGM)
Close - AGE OUT 172.24.76.127 8933 172.24.244.10 138
2001-10-25 07:06:11 0:03:16 10.251.10.25 2699 172.24.60.32 1357 TCP PORT 1357
Close - TCP FIN 172.24.76.127 8921 172.24.60.32 1357
```

Sample traffic log without reason for close listed:

```
ns-> get log traffic
PID 1, from Trust to Untrust, src Any, dst Any, service HTTP, action Permit
Total traffic entries matched under this policy = 1538
================================================================================
Date Time Duration Source IP Port Destination IP Port Service
Xlated Src IP Port Xlated Dst IP Port ID
================================================================================
2002-07-19 15:53:11 0:01:33 10.251.10.25 2712 207.17.137.108 80 HTTP
10.251.10.25 2712 207.17.137.108 80
2002-07-19 15:51:33 0:00:12 10.251.10.25 2711 66.163.175.128 80 HTTP
10.251.10.25 2711 66.163.175.128 80
2002-07-19 15:41:33 0:00:12 10.251.10.25 2688 66.163.175.128 80 HTTP
10.251.10.25 2688 66.163.175.128 80
2002-07-19 15:31:39 0:00:18 10.251.10.25 2678 66.163.175.128 80 HTTP
10.251.10.25 2678 66.163.175.128 80
```

In the following example, you configure the device to not display the reason for closing sessions because it interferes with a script that you want to run on the traffic log. You must use the command line interface to change the log output style.

WebUI

Not available.

CLI

```
set log traffic detail 0
save
```

Self Log

ScreenOS provides a self log to monitor and record all packets terminated at the security device. For example, if you disable some management options on an interface—such as WebUI, SNMP, and ping—and HTTP, SNMP, or ICMP traffic is sent to that interface, entries appear in the self log for each dropped packet.

To activate the self log, do one of the following:

WebUI

Configuration > Report Settings > Log Settings: Select the **Log Packets Terminated to Self** checkbox, then click **Apply**.

CLI

```
set firewall log-self
```

When you enable the self log, the security device logs the entries in two places: the self log and the traffic log. Similar to the traffic log, the self log displays the date, time, source address/port, destination address/port, duration, and service for each dropped packet terminating at the security device. Self log entries typically have a source zone of Null and a destination zone of "self."

Viewing the Self Log

You can view the self log, which is stored in flash storage on the security device, through either the CLI or WebUI.

WebUI

Reports > System Log > Self

CLI

```
get log self
```

Sorting and Filtering the Self Log

Similar to the event and traffic logs, when you use the CLI to view the self log, you can sort or filter the log entries according to the following criteria:

- **Source or Destination IP Address:** You can sort the self log by source or destination IP address. YOu can also filter the self log by specifying a source or destination IP address or range of addresses.

- **Date:** You can sort the self log by date only, or by date and time. The device lists the log entries in descending order by date and time.

  You can also filter self log entries by specifying a start date, an end date, or a date range. When you specify a start date, the device displays log entries with date/time stamps after the start date. When you specify an end date, the device displays log entries with date/time stamps before the end date.

- **Time:** When you sort the self log by time, the security device displays the log entries in descending order by time, regardless of the date. When you specify a start time, the device displays log entries with time stamps after the specified start time, regardless of the date. When you specify an end time, the device displays log entries with time stamps before the specified end time, regardless of the date. When you specify both a start and end time, the device displays log entries with time stamps within the specified time period.

Example: Filtering the Self Log by Time

In this example, you filter self log entries by the end time. The security device displays log entries with time stamps before the specified end time:

WebUI

NOTE: The ability to filter the self log by time is available only through the CLI.

CLI

```
get log self end-time 16:32:57
```

Date	Time	Duration	Source IP	Port	Destination IP	Port	Service
2003-08-21	16:32:57	0:00:00	10.100.25.1	0	224.0.0.5	0	OSPF
2003-08-21	16:32:47	0:00:00	10.100.25.1	0	224.0.0.5	0	OSPF
Total entries matched = 2							

Downloading the Self Log

You can also save the log as a text file to a location you specify, and then use an ASCII text editor (such as Notepad or WordPad) to view it.

In this example, you download a self log to the local directory "C:\netscreen\logs" (WebUI) or to the root directory of a TFTP server at the IP address 10.1.1.5 (CLI). You name the file "self_log07-03-02.txt."

WebUI

1. Reports > System Log > Self: Click **Save**.

   The File Download dialog box prompts you to open the file (using an ASCII editor) or save it to disk.

2. Select the **Save** option, then click **OK**.

   The File Download dialog box prompts you to choose a directory.

3. Specify C:\netscreen\logs, name the file self_log07-03-02.txt, then click **Save**.

CLI

```
get log self > tftp 10.1.1.5 self_log07-03-02.txt
```

Downloading the Asset Recovery Log

A Juniper Networks security device provides an asset recovery log to display information about each time the device is returned to its default settings using the asset recovery procedure (see “Resetting the Device to the Factory Default Settings” on page 39). In addition to viewing the asset recovery log through the WebUI or CLI, you can also open or save the file to the location you specify. Use an ASCII text editor (such as Notepad) to view the file.

In this example, you download the asset recovery log to the local directory “C:\netscreen\logs” (WebUI) or to the root directory of a TFTP server at the IP address 10.1.1.5 (CLI). You name the file “sys_rst.txt,”

WebUI

1. Reports > System Log > Asset Recovery: Click **Save**.

   The File Download dialog box prompts you to open the file (using an ASCII editor) or save it to disk.

2. Select the **Save** option, then click **OK**.

   The File Download dialog box prompts you to choose a directory.

3. Specify C:\netscreen\logs, name the file sys_rst.txt, then click **Save**.

CLI

```
get log asset-recovery > tftp 10.1.1.5 sys_rst.txt
```

Traffic Alarms

The security device supports traffic alarms when traffic exceeds thresholds that you have defined in policies. You can configure the security device to alert you through one or more of the following methods whenever the security device generates a traffic alarm:

- Console
- Internal (Event Log)
- Email
- SNMP
- Syslog
- WebTrends
- NetScreen-Global PRO

You set alarm thresholds to detect anomalous activity. To know what constitutes anomalous activity, you must first establish a baseline of normal activity. To create such a baseline for network traffic, you must observe traffic patterns over a period of time. Then, after you have determined the amount of traffic that you consider as

normal, you can set alarm thresholds above that amount. Traffic exceeding that threshold triggers an alarm to call your attention to a deviation from the baseline. You can then evaluate the situation to determine what caused the deviation and whether you need to take action in response.

You can also use traffic alarms to provide policy-based intrusion detection and notification of a compromised system. Examples of the use of traffic alarms for these purposes are provided below.

Example: Policy-Based Intrusion Detection

In this example, there is a webserver with IP address 211.20.1.5 (and name "web1") in the DMZ zone. You want to detect any attempts from the Untrust zone to access this webserver via Telnet. To accomplish this, you create a policy denying Telnet traffic from any address in the Untrust zone destined to the webserver named web1 in the DMZ zone, and you set a traffic alarm threshold at 64 bytes. Because the smallest size of IP packet is 64 bytes, even one Telnet packet attempting to reach the webserver from the Untrust zone will trigger an alarm.

WebUI

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

Address Name: web1
IP Address/Domain Name:
IP/Netmask: (select), 211.20.1.5/32
Zone: DMZ

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), web1
Service: Telnet
Action: Deny

> Advanced: Enter the following, then click **Return** to set the advanced options and return to the basic configuration page:

Counting: (select)
Alarm Threshold: 64 Bytes/Sec, 0 Kbytes/Min

CLI

```
set address dmz web1 211.20.1.5/32
set policy from untrust to dmz any web1 telnet deny count alarm 64 0
save
```

Example: Compromised System Notification

In this example, you use traffic alarms to provide notification of a compromised system. You have an FTP server with IP address 211.20.1.10 (and name ftp1) in the DMZ zone. You want to allow FTP-get traffic to reach this server. You don't want traffic of any kind to originate from the FTP server. The occurrence of such traffic would indicate that the system has been compromised, perhaps by a virus similar to the NIMDA virus. You define an address for the FTP server in the Global zone, so that you can then create two global policies.

WebUI

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

Address Name: ftp1
IP Address/Domain Name:
IP/Netmask: (select), 211.20.1.10/32
Zone: Global

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), ftp1
Service: FTP-Get
Action: Permit

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

Source Address:
Address Book Entry: (select), ftp1
Destination Address:
Address Book Entry: (select), Any
Service: ANY
Action: Deny

> Advanced: Enter the following, then click **Return** to set the advanced options and return to the basic configuration page:

Counting: (select)
Alarm Threshold: 64 Bytes/Sec, 0 Kbytes/Min

CLI

```
set address global ftp1 211.20.1.10/32
set policy global any ftp1 ftp-get permit
set policy global ftp1 any deny count alarm 64 0
save
```

Example: Sending E-mail Alerts

In this example, you set up notification by email alerts when there is an alarm. The mail server is at 172.16.10.254, the first email address to be notified is jharker@juniper.net, and the second address is driggs@juniper.net. The security device includes traffic logs with event logs sent via email.

WebUI

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

```
Enable E-Mail Notification for Alarms: (select)
Include Traffic Log: (select)
    SMTP Server Name: 172.16.10.254
E-Mail Address 1: jharker@juniper.net
E-Mail Address 2: driggs@juniper.net
```

NOTE: If you have DNS enabled, you can also use a host name for the mail server, such as mail.juniper.net.

CLI

```
set admin mail alert
set admin mail mail-addr1 jharker@juniper.net
set admin mail mail-addr2 driggs@juniper.net
set admin mail server-name 172.16.10.254
set admin mail traffic-log
save
```

Syslog

A security device can generate syslog messages for system events at predefined severity levels (see the list of severity levels in "Event Log" on page 54), and optionally for traffic that policies permit across a firewall. It sends these messages to up to four designated syslog hosts running on UNIX/Linux systems. For each syslog host, you can specify the following:

- Whether the security device includes traffic log entries, event log entries, or both traffic and event log entries.
- Whether to send traffic through a VPN tunnel to the syslog server and—if through a VPN tunnel—which interface to use as the source interface (for examples, see "Example: Self-Generated Traffic Through a Route-Based Tunnel" on page 76 and "Example: Self-Generated Traffic Through a Policy-Based Tunnel" on page 83).
- The port to which the security device sends syslog messages.
- The security facility, which classifies and sends emergency and alert level messages to the Syslog host; and the regular facility, which classifies and sends all other messages for events unrelated to security.

By default, the security device sends messages to syslog hosts via UDP (port 514). To increase the reliability of the message delivery, you can change the transport protocol for each syslog host to TCP.

You can use syslog messages to create email alerts for the system administrator, or to display messages on the console of the designated host using UNIX syslog conventions.

NOTE: On UNIX/Linux platforms, modify the /etc/rc.d/init.d/syslog file so that syslog retrieves information from the remote source (syslog -r).

Example: Enabling Multiple Syslog Servers

In this example, you configure the security device to send event and traffic logs via TCP to three syslog servers at the following IP addresses/port numbers: 1.1.1.1/1514, 2.2.2.1/2514, and 3.3.3.1/3514. You set both the security and facility levels to Local0.

WebUI

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

Enable syslog messages: Select this option to send logs to the specified syslog servers.
No.: Select 1, 2, and 3 to indicate you are adding 3 syslog servers.
IP/Hostname: 1.1.1.1, 2.2.2.1, 3.3.3.1
Port: 1514, 2514, 3514
Security Facility: Local0, Local0, Local0
Facility: Local0, Local0, Local0
Event Log: (select)
Traffic Log: (select)
TCP: (select)

CLI

```
set syslog config 1.1.1.1 port 1514
set syslog config 1.1.1.1 log all
set syslog config 1.1.1.1 facilities local0 local0
set syslog config 1.1.1.1 transport tcp
set syslog config 2.2.2.1 port 2514
set syslog config 2.2.2.1 log all
set syslog config 2.2.2.1 facilities local0 local0
set syslog config 2.2.2.1 transport tcp
set syslog config 3.3.3.1 port 3514
set syslog config 3.3.3.1 log all
set syslog config 3.3.3.1 facilities local0 local0
set syslog config 2.2.2.1 transport tcp
set syslog enable
save
```

Enabling WebTrends for Notification Events

NetIQ offers a product called the WebTrends Firewall Suite that allows you to create customized reports based on the logs generated by a security device. WebTrends analyzes the log files and displays the information you need in a graphical format. You can create reports on all events and severity levels or focus on an area such as firewall attacks. (For additional information on WebTrends, refer to the WebTrends product documentation.)

You can also send WebTrends messages through a VPN tunnel. In the WebUI, use the **Use Trust Zone Interface as Source IP for VPN** option. In the CLI, use the **set webtrends vpn** command.

In the following example, you send notification messages to the WebTrends host (172.10.16.25).

WebUI

1. **WebTrends Settings**
   Configuration > Report Settings > WebTrends: Enter the following, then click **Apply**:

   Enable WebTrends Messages: (select)
   WebTrends Host Name/Port: 172.10.16.25/514

2. **Severity Levels**
   Configuration > Report Settings > Log Settings: Enter the following, then click **Apply**:

   WebTrends Notification: (select)

NOTE: When you enable WebTrends on a security device running in Transparent mode, you must set up a static route. See "Static Routing" on page **7**-1.

CLI

3. **WebTrends Settings**
   ```
   set webtrends host-name 172.10.16.25
   set webtrends port 514
   set webtrends enable
   ```

4. **Severity Levels**
   ```
   set log module system level notification destination webtrends
   save
   ```

Simple Network Management Protocol

The Simple Network Management Protocol (SNMP) agent for the Juniper Networks security device provides network administrators with a way to view statistical data about the network and the devices on it, and to receive notification of system events of interest.

Juniper Networks security devices support the SNMPv1 protocol, described in RFC-1157, *A Simple Network Management Protocol*, and the SNMPv2c protocol, described in the following RFCs:

- RFC-1901, *Introduction to Community-based SNMPv2*
- RFC-1905, *Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)*
- RFC-1906, *Transport Mappings for Version 2 of the Simple Network Management Protocol (SNMPv2)*

Security devices also support all relevant Management Information Base II (MIB II) groups defined in RFC-1213, *Management Information Base for Network Management of TCP/IP-based Internets: MIB-II*. The devices also have private enterprise MIB files, which you can load into an SNMP MIB browser.

NOTE: Using SNMP MIB Browser, you can check the cpu usage, memory usage, and session usage counts on both the ScreenOS and Intrusion Detection Protection (IDP) security modules.

The Juniper Networks SNMP agent generates the following traps, or notifications, when specified events or conditions occur:

- **Cold Start Trap:** The security device generates a cold start trap when it becomes operational after you power it on.

- **Trap for SNMP Authentication Failure:** The SNMP agent in the security device triggers the authentication failure trap if someone attempts to connect to it using an incorrect SNMP community string or if the IP address of the host attempting the connection is not defined in an SNMP community. (This option is enabled by default.)

- **Traps for System Alarms:** security device error conditions and firewall conditions trigger system alarms. Three enterprise traps are defined to cover alarms related to hardware, security, and software. (For more information on firewall settings and alarms, see "ICMP Fragments" on page **4**-228 and "Traffic Alarms" on page 65.)

- **Traps for Traffic Alarms:** Traffic alarms are triggered when traffic exceeds the alarm thresholds set in policies. (For more information on configuring policies, see "Policies" on page **2**-171.)

"Trap Alarm Types" on page 72 lists possible alarm types and their associated trap number:

Table 2: Trap Alarm Types

Trap Enterprise ID	Description
100	Hardware problems
200	Firewall problems
300	Software problems
400	Traffic problems
500	VPN problems
600	NSRP problems
800	DRP problems
900	Interface failover problems
1000	Firewall attacks

NOTE: The network administrator must have an SNMP manager application such as HP OpenView or SunNet Manager to browse the SNMP MIB II data and to receive traps from either the trusted or untrusted interface. Shareware and freeware SNMP manager applications available from the Internet.

Security devices do not ship with a default configuration for the SNMP manager. To configure your security device for SNMP, you must first create communities, define their associated hosts, and assign permissions (read/write or read-only).

When you create an SNMP community, you can specify whether the community supports SNMPv1, SNMPv2c, or both SNMP versions, as required by the SNMP management stations. (For backward compatibility with earlier ScreenOS releases that only support SNMPv1, security devices support SNMPv1 by default.) If an SNMP community supports both SNMP versions, you must specify a trap version for each community member.

For security reasons, an SNMP community member with read/write privileges can change only the following variables on a security device:

- **sysContact** - Contact information for the admin of the security device in case the SNMP admin needs to contact him or her. This can be the security admin's name, email address, telephone number, office location, or a combination of such information.
- **sysLocation** - The physical location of the security device. This can be the name of a country, city, building, or its exact location on a rack in a network operations center (NOC).
- **sysName** - The name that SNMP administrators use for the security device. By convention, this is a fully-qualified domain name (FQDN), but it can be something else.
- **snmpEnableAuthenTraps** - This enables or disables the SNMP agent in the security device to generate a trap whenever someone attempts to contact the SNMP agent with an incorrect SNMP community name.

- **ipDefaultTTL** - The default value inserted into the time-to-live (TTL) field in the IP header of datagrams originating from the security device whenever the Transport Layer protocol does not supply a TTL value.

- **ipForwarding** - This indicates whether or not the security device forwards traffic—other than that destined for the security device itself. By default, the security device indicates that it does not forward traffic.

Implementation Overview

Juniper Networks has implemented SNMP in its devices in the following ways:

- SNMP administrators are grouped in SNMP communities. A device can support up to three communities, with up to eight members in each community.

- A community member can be either a single host or a subnet of hosts, depending on the netmask you use when defining the member. By default, the security device assigns an SNMP community member with a 32-bit netmask (255.255.255.255), which defines it as a single host.

- If you define an SNMP community member as a subnet, any device on that subnet can poll the security device for SNMP MIB information. However, the security device cannot send an SNMP trap to a subnet, only to an individual host.

- Each community has either read-only or read-write permission for the MIB II data.

- Each community can support SNMPv1, SNMPv2c, or both. If a community supports both versions of SNMP, you must specify a trap version for each community member.

- You can allow or deny each community from receiving traps.

- You can access the MIB II data and traps through any physical interface.

- Each system alarm (a system event classified with a severity level of critical, alert, or emergency) generates a single enterprise SNMP trap to each of the hosts in each community that is set to receive traps.

- The security device sends Cold Start / Link Up / Link Down traps to all hosts in communities that you set to receive traps.

- If you specify trap-on for a community, you also have the option to allow traffic alarms.

- You can send SNMP messages through a route-based or policy-based VPN tunnel. For more information, see “VPN Tunnels for Self-Initiated Traffic” on page 75.

Defining a Read/Write SNMP Community

In this example, you create an SNMP community, named *MAge11*. You assign it read/write privileges and enable its members to receive MIB II data and traps. It has the following two members 1.1.1.5/32 and 1.1.1.6/32. Each of these members has an SNMP manager application running a different version of SNMP: SNMPv1 and SNMPv2c. The community name functions as a password and needs to be protected.

You provide contact information for the local admin of the security device in case an SNMP community member needs to contact him—name: al_baker@mage.com. You also provide the location of the security device—location: 3-15-2. These numbers indicate that the device is on the third floor, in the fifteenth row, and in the second position in that row.

You also enable the SNMP agent to generate traps whenever someone illegally attempts an SNMP connection to the security device. Authentication failure traps is a global setting that applies to all SNMP communities and is disabled by default.

Finally, you enable SNMP manageability on ethernet1, an interface that you have previously bound to the Trust zone. This is the interface through which the SNMP manager application communicates with the SNMP agent in the security device.

WebUI

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

System Contact: al_baker@mage.com
Location: 3-15-2
Enable Authentication Fail Trap: (select)

Configuration > Report Settings > SNMP > New Community: Enter the following settings, then click **OK**:

Community Name: MAge11
Permissions:
    Write: (select)
    Trap: (select)
    Including Traffic Alarms: (clear)
Version: ANY (select)
Hosts IP Address/Netmask and Trap Version:
  1.1.1.5/32 v1
  1.1.1.6/32 v2c

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

Service Options:
    Management Services: SNMP

CLI

```
set snmp contact al_baker@mage.com
set snmp location 3-15-2
set snmp auth-trap enable
set snmp community MAge11 read-write trap-on version any
set snmp host Mage 1.1.1.5/32 trap v1
set snmp host Mage 1.1.1.6/32 trap v2
set interface ethernet1 manage snmp
save
```

VPN Tunnels for Self-Initiated Traffic

You can use virtual private network (VPN) tunnels to secure remote monitoring of a security device from a fixed IP address. Using a VPN tunnel, you can protect traffic addressed to and initiated from a security device. Types of traffic initiated from a security device can include NetScreen-Global PRO reports, event log entries sent to syslog and WebTrends servers, and SNMP MIB traps.

Juniper Networks security devices support two types of VPN tunnel configurations:

- **Route-Based VPNs**: The security device uses route table entries to direct traffic to tunnel interfaces, which are bound to VPN tunnels.

  To send traffic such as event log entries, NetScreen-Global PRO reports, or SNMP traps generated by the security device through a route-based VPN tunnel, you must manually enter a route to the proper destination. The route must point to the tunnel interface that is bound to the VPN tunnel through which you want the security device to direct the traffic. No policy is required.

- **Policy-Based VPNs**: The security device uses the VPN tunnel names specifically referenced in policies to direct traffic through VPN tunnels.

  To send self-initiated traffic through a policy-based VPN tunnel, you must include the source and destination addresses in the policy. For the source address, use the IP address of an interface on the security device. For the destination address, use the IP address of the storage server or SNMP community member's workstation, if it is located behind a remote security device. If the remote SNMP community member uses the NetScreen-Remote VPN client to make VPN connections to the local security device, use an internal IP address defined on the NetScreen-Remote as the destination address.

If either the remote gateway or the end entity has a dynamically assigned IP address, then the security device cannot initiate the formation of a VPN tunnel because these addresses cannot be predetermined, and thus you cannot define routes to them. In such cases, the remote host must initiate the VPN connection. After either a policy-based or route-based VPN tunnel is established, both ends of the tunnel can initiate traffic if policies permit it. Also, for a route-based VPN, there must be a route to the end entity through a tunnel interface bound to the VPN tunnel—either because you manually entered the route or because the local security device received the route through the exchange of dynamic routing messages after a tunnel was established. (For information about dynamic routing

protocols, see *Volume 7: Routing*.) You can also use VPN monitoring with the rekey option or IKE heartbeats to ensure that once the tunnel is established, it remains up regardless of VPN activity. (For more information about these options, see "VPN Monitoring" on page **5**-238 and "Monitoring Mechanisms" on page **5**-289.)

For each VPN tunnel configuration type, you can use any of the following types of VPN tunnel:

- **Manual Key**: You manually set the three elements that define a Security Association (SA) at both ends of the tunnel: a Security Parameters Index (SPI), an encryption key, and an authentication key. To change any element in the SA, you must manually enter it at both ends of the tunnel.
- **AutoKey IKE with Pre-shared Key**: One or two pre-shared secrets—one for authentication and one for encryption—function as seed values. Using them, the IKE protocol generates a set of symmetrical keys at both ends of the tunnel; that is, the same key is used to encrypt and decrypt. At predetermined intervals, these keys are automatically regenerated.
- **AutoKey IKE with Certificates**: Using the Public Key Infrastructure (PKI), the participants at both ends of the tunnel use a digital certificate (for authentication) and an RSA public/private key pair (for encryption). The encryption is asymmetrical; that is, one key in a pair is used to encrypt and the other to decrypt.

NOTE: For a complete description of VPN tunnels, see *Volume 5: Virtual Private Networks*. For more information on NetScreen-Remote, refer to the *NetScreen-Remote VPN Client Administrator Guide*.

Example: Self-Generated Traffic Through a Route-Based Tunnel

Figure 13 illustrates an example in which you configure a local security device (Device-A) to send SNMPv1 MIB traps and syslog reports through a route-based AutoKey IKE VPN tunnel to an SNMP community member behind a remote security device (Device-B). The tunnel uses a pre-shared key (Ci5y0a1aAG) for data origin authentication and the security level predefined as "Compatible" for both Phase 1 and Phase 2 proposals. You, as the local admin for Device-A, create the tunnel.1 interface and bind it to vpn1. You and the admin for Device-B define the proxy IDs as shown in Table 3:

Table 3: Proxy IDs for Route-Based Tunnel

Device-A		Device-B	
Local IP	10.1.1.1/32	Local IP	10.2.2.2/32
Remote IP	10.2.2.2/32	Remote IP	10.1.1.1/32
Service	Any	Service	Any

You bind ethernet1 to the Trust zone and ethernet3 to the Untrust zone. The default gateway IP address is 1.1.1.250. All zones are in the trust-vr routing domain.

Figure 13: Traffic Through a Route-Based Tunnel

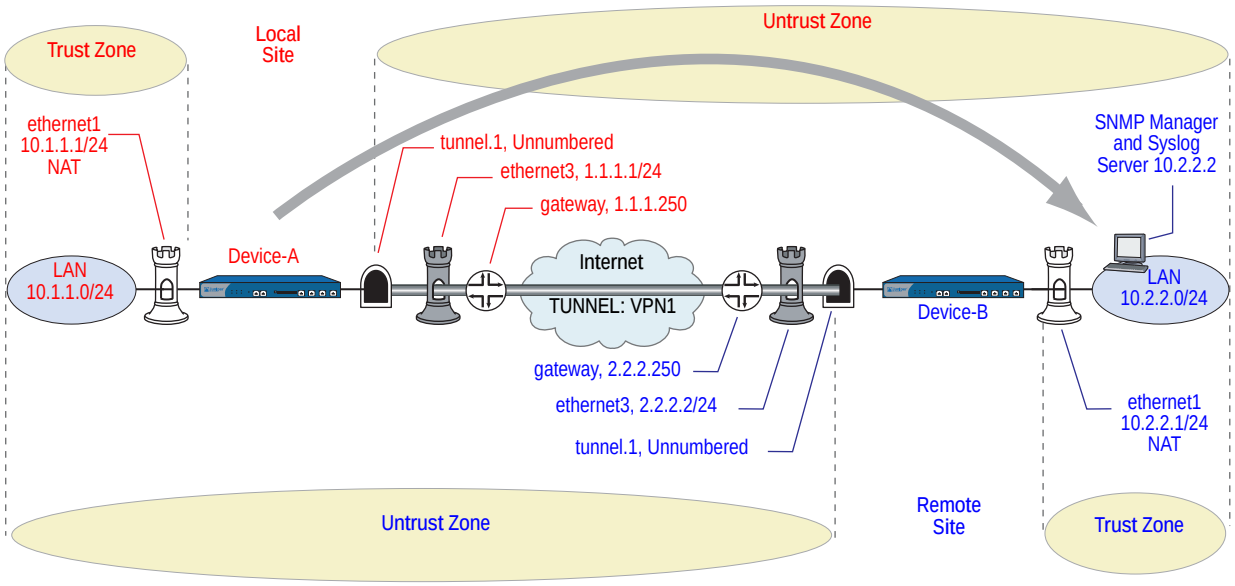

The remote admin for Device-B uses similar settings to define that end of the AutoKey IKE VPN tunnel so that the pre-shared key, proposals, and proxy IDs match.

You also configure an SNMP community named “remote_admin” with read/write privileges, and you enable the community to receive traps. You define the host at 10.2.2.2/32 as a community member.

NOTE: This example assumes that the remote admin has already set up the syslog server and SNMP manager application that supports SNMPv1. When the remote admin sets up the VPN tunnel on his security device, he uses 1.1.1.1 as the remote gateway and 10.1.1.1 as the destination address.

WebUI (Device-A)

1. **Interfaces**
   Network > Interfaces > Edit (for ethernet1): 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 (select)

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.

When the remote admin configures the SNMP manager, he must enter **10.1.1.1** in the Remote SNMP Agent field. This is the address to which the SNMP manager sends queries.

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

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

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

Tunnel Interface Name: tunnel.1
Zone (VR): Untrust (trust-vr)
Unnumbered: (select)
Interface: ethernet1(trust-vr)

2. Syslog and SNMP

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

Enable Syslog Messages: (select)
No.: Select 1 to indicate you are adding 1 syslog server.
IP / Hostname: 10.2.2.2
Port: 514
Security Facility: auth/sec
Facility: Local0

Configuration > Report Settings > SNMP > New Community: Enter the following, then click **OK**:

Community Name: remote_admin
Permissions:
Write: (select)
Trap: (select)
Including Traffic Alarms: (clear)
Version: V1
Hosts IP Address/Netmask:
10.2.2.2/32 V1
Trap Version:
V1

3. **VPN**

   VPNs > AutoKey IKE > New: Enter the following, then click **OK**:

   VPN Name: vpn1
   Security Level: Compatible
   Remote Gateway: Create a Simple Gateway: (select)
   Gateway Name: to_admin
   Type: Static IP, Address/Hostname: 2.2.2.2
   Preshared Key: Ci5y0a1aAG
   Security Level: Compatible
   Outgoing interface: ethernet3

   > Advanced: Enter the following advanced settings, then click **Return** to return to the basic AutoKey IKE configuration page:

   Bind to: Tunnel Interface: (select), tunnel.1
   Proxy-ID: (select)
   Local IP/Netmask: 10.1.1.1/32
   Remote IP/Netmask: 10.2.2.2/32
   Service: ANY

4. **Routes**

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

   Network Address/Netmask: 10.2.2.2/32
   Gateway: (select)
   Interface: tunnel.1
   Gateway IP Address: 0.0.0.0

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

   Network Address/Netmask: 0.0.0.0/0
   Gateway: (select)
   Interface: ethernet3
   Gateway IP Address: (select) 1.1.1.250

CLI (Device-A)

1. **Interfaces**

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat
set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
set interface ethernet3 manage snmp
set interface tunnel.1 zone untrust
set interface tunnel.1 ip unnumbered interface ethernet1
```

NOTE: When the remote admin configures the SNMP manager, he must enter **10.1.1.1** in the Remote SNMP Agent field. This is the address to which the SNMP manager sends queries.

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.

---

2. **VPN**

```
set ike gateway to_admin address 2.2.2.2 outgoing-interface ethernet3 preshare
    Ci5y0a1aAG sec-level compatible
set vpn vpn1 gateway to_admin sec-level compatible
set vpn vpn1 bind interface tunnel.1
set vpn vpn1 proxy-id local-ip 10.1.1.1/32 remote-ip 10.2.2.2/32 any
```

3. **Syslog and SNMP**

```
set syslog config 10.2.2.2 auth/sec local0
set syslog enable
set snmp community remote_admin read-write trap-on version v1
set snmp host remote_admin 10.2.2.2/32
```

4. **Routes**

```
set vrouter trust-vr route 10.2.2.2/32 interface tunnel.1
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
save
```

WebUI (Device-B)

1. **Interfaces**

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

Zone Name: Trust
Static IP: (select this option when present)
IP Address/Netmask: 10.2.2.1/24

Select the following, then click **OK**:

Interface Mode: NAT

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

Zone Name: Untrust
Static IP: (select this option when present)
IP Address/Netmask: 2.2.2.2/24

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

Tunnel Interface Name: tunnel.1
Zone (VR): Untrust (trust-vr)
Unnumbered: (select)
Interface: ethernet1(trust-vr)

2. **Addresses**

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

Address Name: addr1
IP Address/Domain Name: IP/Netmask: 10.2.2.2/32
Zone: Trust

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

Address Name: ns-a
IP Address/Domain Name: IP/Netmask: 10.1.1.1/32
Zone: Untrust

3. **Service Group**

   Objects > Services > Groups > New: Enter the following group name, move the following services, then click **OK**:

   Group Name: s-grp1

   Select **Syslog** and use the < < button to move the service from the Available Members column to the Group Members column.

   Select **SNMP** and use the < < button to move the service from the Available Members column to the Group Members column.

4. **VPN**

   VPNs > AutoKey IKE > New: Enter the following, then click **OK**:

   VPN Name: vpn1
   Security Level: Compatible
   Remote Gateway: Create a Simple Gateway: (select)
   Gateway Name: to_admin
   Type: Static IP, Address/Hostname: 1.1.1.1
   Preshared Key: Ci5y0a1aAG
   Security Level: Compatible
   Outgoing Interface: ethernet3

   > Advanced: Enter the following advanced settings, then click **Return** to return to the basic AutoKey IKE configuration page:

   Bind to: Tunnel Interface: (select), tunnel.1
   Proxy-ID: (select)
   Local IP/Netmask: 10.2.2.2/32
   Remote IP/Netmask: 10.1.1.1/32
   Service: Any

5. **Routes**

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

   Network Address/Netmask:10.1.1.1/32
   Gateway: (select)
   Interface: tunnel.1
   Gateway IP Address: 0.0.0.0

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

   Network Address/Netmask: 0.0.0.0/0
   Gateway: (select)
   Interface: ethernet3
   Gateway IP Address: (select) 2.2.2.250

6. **Policies**

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

   Source Address:
   Address Book Entry: (select), addr1
   Destination Address:
   Address Book Entry: (select), ns-a
   Service: s-grp1
   Action: Permit
   Position at Top: (select)

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

   Source Address:
   Address Book Entry: (select), ns-a
   Destination Address:
   Address Book Entry: (select), addr1
   Service: s-grp1
   Action: Permit
   Position at Top: (select)

CLI (Device-B)

1. **Interfaces**

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.2.2.1/24
set interface ethernet1 nat
set interface ethernet3 zone untrust
set interface ethernet3 ip 2.2.2.2/24
set interface tunnel.1 zone untrust
set interface tunnel.1 ip unnumbered interface ethernet1
```

2. **Addresses**

```
set address trust addr1 10.2.2.2/32
set address untrust ns-a 10.1.1.1/32
```

3. **Service Group**

```
set group service s-grp1
set group service s-grp1 add syslog
set group service s-grp1 add snmp
```

4. **VPN**

```
set ike gateway to_admin address 1.1.1.1 outgoing-interface ethernet3 preshare
    Ci5y0a1aAG sec-level compatible
set vpn vpn1 gateway to_admin sec-level compatible
set vpn vpn1 bind interface tunnel.1
set vpn vpn1 proxy-id local-ip 10.2.2.2/32 remote-ip 10.1.1.1/32 any
```

5. **Routes**

```
set vrouter trust-vr route 10.1.1.1/32 interface tunnel.1
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 2.2.2.250
```

6. **Policies**

```
set policy top from trust to untrust addr1 ns-a s-grp1 permit
set policy top from untrust to trust ns-a addr1 s-grp1 permit
save
```

Example: Self-Generated Traffic Through a Policy-Based Tunnel

In this example (illustrated in Figure 14), you configure a local security device (Device-A) to send SNMPv2c MIB traps and syslog reports through a policy-based AutoKey IKE VPN tunnel (vpn1) to an SNMP community member behind a remote security device (Device-B). The tunnel uses a preshared key (Ci5y0a1aAG) for data origin authentication and the security level predefined as "Compatible" for both Phase 1 and Phase 2 proposals.

NOTE: This example assumes that the remote admin has already set up the syslog server and an SNMP manager application that supports SNMPv2c. When the remote admin sets up the VPN tunnel on his security device, he uses 1.1.1.1 as the remote gateway and 10.1.1.1 as the destination address.

Both you and the remote admin bind ethernet1 to the Trust zone, and ethernet3 to the Untrust zone on Device-A and Device-B. The default gateway IP address for Device-A is 1.1.1.250. The default gateway IP address for Device-B is 2.2.2.250. All zones are in the trust-vr routing domain.

Figure 14: Traffic Through a Policy-Based Tunnel

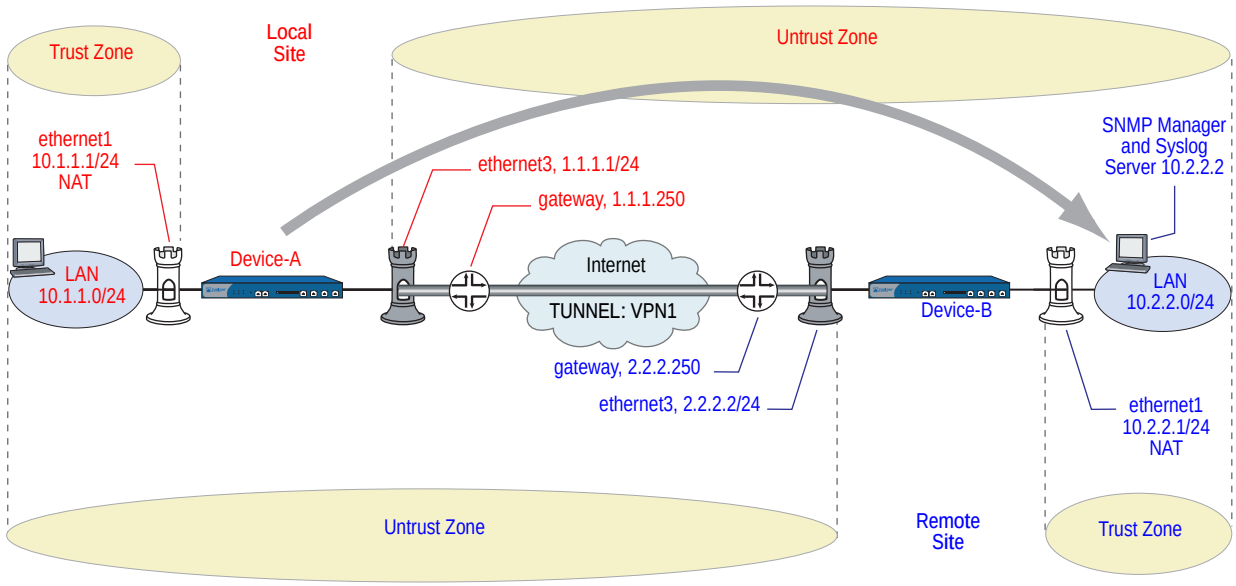

You also configure an SNMP community named "remote_admin" with read/write privileges, and you enable the community to receive traps. You define the host at 10.2.2.2/32 as a community member.

The inbound and outbound policies on Device-A match the outbound and inbound policies on Device-B. The addresses and service used in the policies are as follows:

- 10.1.1.1/32, the address of the Trust zone interface on Device-A
- 10.2.2.2/32, the address of the host for the SNMP community member and syslog server
- Service group named "s-grp1," which contains SNMP and syslog services

From the policies that you and the admin for Device-B create, the two security devices derive the following proxy IDs for vpn1:

Table 4: Proxy IDs for Policy-Based Tunnel

Device-A		Device-B	
Local IP	10.1.1.1/32	Local IP	10.2.2.2/32
Remote IP	10.2.2.2/32	Remote IP	10.1.1.1/32
Service	Any	Service	Any

NOTE: The security device treats a service group as "any" in proxy IDs.

WebUI (Device-A)

1. **Interfaces—Security Zones**
   Network > Interfaces > Edit (for ethernet1): Enter the following, then click **OK**:

   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 (select)

NOTE: When the remote admin configures the SNMP manager, he must enter **10.1.1.1** in the Remote SNMP Agent field. This is the address to which the SNMP manager sends queries.

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 ethernet3): Enter the following, then click **OK**:

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

2. **Addresses**
   Objects > Addresses > List > New: Enter the following, then click **OK**:

   Address Name: trust_int
   IP Address/Domain Name:
   IP/Netmask: 10.1.1.1/32
   Zone: Trust

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

Address Name: remote_admin
IP Address/Domain Name:
IP/Netmask: 10.2.2.2/32
Zone: Untrust

3. **Service Group**
Objects > Services > Groups > New: Enter the following group name, move the following services, then click **OK**:

Group Name: s-grp1
Select **Syslog** and use the < < button to move the service from the Available Members column to the Group Members column.

Select **SNMP** and use the < < button to move the service from the Available Members column to the Group Members column.

4. **VPN**
VPNs > AutoKey IKE > New: Enter the following, then click **OK**:

VPN Name: vpn1
Security Level: Compatible
Remote Gateway: Create a Simple Gateway: (select)
Gateway Name: to_admin
Type: Static IP, Address/Hostname: 2.2.2.2
Preshared Key: Ci5y0a1aAG
Security Level: Compatible
Outgoing Interface: ethernet3

5. **Syslog and SNMP**
Configuration > Report Settings > Syslog: Enter the following, then click **Apply**:

Enable Syslog Messages: (select)
Source Interface: ethernet1
No.: Select 1 to indicate you are adding 1 syslog server.
IP/Hostname: 10.2.2.2
Port: 514
Security Facility: auth/sec
Facility: Local0

Configuration > Report Settings > SNMP > New Community: Enter the following, then click **OK**:

Community Name: remote_admin
Permissions:
    Write: (select)
    Trap: (select)
    Including Traffic Alarms: (clear)
Version: V2C
Hosts IP Address/Netmask:
    10.2.2.2/32 V2C
Trap Version:
    V2C
Source Interface:
    ethernet1 (select)

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

Enable Authentication Fail Trap: (select)

6. **Route**

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

Network Address/Netmask: 0.0.0.0/0
Gateway: (select)
    Interface: ethernet3
    Gateway IP Address: 1.1.1.250

7. **Policies**

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

Source Address:
    Address Book Entry: (select), trust_int
Destination Address:
    Address Book Entry: (select), remote_admin
Service: s-grp1
Action: Tunnel
Tunnel VPN: vpn1
Modify matching outgoing VPN policy: (select)
Position at Top: (select)

CLI (Device-A)

1. **Interfaces—Security Zones**

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat
set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
set interface ethernet3 manage snmp
```

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.

2. **Addresses**
   set address trust trust_int 10.1.1.1/32
   set address untrust remote_admin 10.2.2.2/32

3. **Service Group**
   set group service s-grp1
   set group service s-grp1 add syslog
   set group service s-grp1 add snmp

4. **VPN**
   set ike gateway to_admin address 2.2.2.2 outgoing-interface ethernet3 preshare Ci5y0a1aAG sec-level compatible
   set vpn vpn1 gateway to_admin sec-level compatible

5. **Syslog and SNMP**
   set syslog config 10.2.2.2 auth/sec local0
   set syslog src-interface ethernet1
   set syslog enable
   set snmp community remote_admin read-write trap-on version v2c
   set snmp host remote_admin 10.2.2.2/32 src-interface ethernet1

6. **Route**
   set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250

7. **Policies**
   set policy top from trust to untrust trust_int remote_admin s-grp1 tunnel vpn vpn1
   set policy top from untrust to trust remote_admin trust_int s-grp1 tunnel vpn vpn1
   save

WebUI (Device-B)

1. **Interfaces—Security Zones**
   Network > Interfaces > Edit (for ethernet1): Enter the following, then click **Apply**:

   Zone Name: Trust
   Static IP: (select this option when present)
   IP Address/Netmask: 10.2.2.1/24

   Select the following, then click **OK**:

   Interface Mode: NAT

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

   Zone Name: Untrust
   Static IP: (select this option when present)
   IP Address/Netmask: 2.2.2.2/24

2. **Addresses**
   Objects > Addresses > List > New: Enter the following, then click **OK**:

   Address Name: addr1
   IP Address/Domain Name:
   IP/Netmask: 10.2.2.2/32
   Zone: Trust

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

Address Name: ns-a
IP Address/Domain Name:
IP/Netmask: 10.1.1.1/32
Zone: Untrust

3. **Service Group**
Objects > Services > Group: Enter the following group name, move the following services, then click **OK**:

Group Name: s-grp1

Select **Syslog** and use the < < button to move the service from the Available Members column to the Group Members column.

Select **SNMP** and use the < < button to move the service from the Available Members column to the Group Members column.

4. **VPN**
VPNs > AutoKey IKE > New: Enter the following, then click **OK**:

VPN Name: vpn1
Security Level: Compatible
Remote Gateway: Create a Simple Gateway: (select)
Gateway Name: to_admin
Type: Static IP, IP Address: 1.1.1.1
Preshared Key: Ci5y0a1aAG
Security Level: Compatible
Outgoing interface: ethernet3

5. **Route**
Network > Routing > Routing Table > trust-vr New: Enter the following, then click **OK**:

Network Address/Netmask: 0.0.0.0/0
Gateway: (select)
Interface: ethernet3
Gateway IP Address: (select) 2.2.2.250

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

Source Address:
Address Book Entry: (select), addr1
Destination Address:
Address Book Entry: (select), ns-a
Service: s-grp1
Action: Tunnel
Tunnel VPN: vpn1
Modify matching outgoing VPN policy: (select)
Position at Top: (select)

CLI (Device-B)

1. **Interfaces—Security Zones**
   set interface ethernet1 zone trust
   set interface ethernet1 ip 10.2.2.1/24
   set interface ethernet1 nat
   set interface ethernet3 zone untrust
   set interface ethernet3 ip 2.2.2.2/24

2. **Addresses**
   set address trust addr1 10.2.2.2/32
   set address untrust ns-a 10.1.1.1/32

3. **Service Group**
   set group service s-grp1
   set group service s-grp1 add syslog
   set group service s-grp1 add snmp

4. **VPN**
   set ike gateway to_admin address 1.1.1.1 outgoing-interface ethernet3 preshare Ci5y0a1sec-level compatible
   set vpn vpn1 gateway to_admin sec-level compatible

5. **Route**
   set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 2.2.2.250

6. **Policies**
   set policy top from trust to untrust addr1 ns-a s-grp1 tunnel vpn vpn1
   set policy top from untrust to trust ns-a addr1 s-grp1 tunnel vpn vpn1
   save

Viewing Screen Counters

Juniper Networks security devices provide screen, hardware, and flow counters for monitoring traffic. Counters give processing information for specified zones and interfaces, and help you to verify configurations for desired policies.

Table 5 shows the screen counters for monitoring general firewall behavior and for viewing the amount of traffic affected by specified policies.

Table 5: Screen Counters

Counter	Description
Bad IP Option Protection	Number of frames discarded because of malformed or incomplete IP options
Dst IP-based session limiting	Number of sessions dropped after the session threshold was reached
FIN bit with no ACK bit	Number of packets detected and dropped with an illegal combination of flags
Fragmented packet protection	Number of blocked IP packet fragments
HTTP Component Blocked	Number of blocked packets with HTTP components
HTTP Component Blocking for ActiveX controls	Number of ActiveX components blocked
HTTP Component Blocking for .exe files	Number of blocked HTTP packets with .exe files
HTTP Component Blocking for Java applets	Number of blocked Java components

Counter	Description
HTTP Component Blocking for .zip files	Number of blocked HTTP packets with .zip files
ICMP Flood Protection	Number of ICMP packets blocked as part of an ICMP flood
ICMP Fragment	Number of ICMP frames with the More Fragments flag set, or with offset indicated in the Offset field
IP Spoofing Attack Protection	Number of IP addresses blocked as part of an IP spoofing attack
IP Sweep Protection	Number of IP sweep attack packets detected and blocked
Land Attack Protection	Number of packets blocked as part of a suspected land attack
Large ICMP Packet	Number of ICMP frames detected with an IP length greater than 1024
Limit Session	Number of undeliverable packets because the session limit had been reached
Loose Src Route IP Option	Number of IP packets detected with the Loose Source Route option enabled
Malicious URL Protection	Number of suspected malicious URLs blocked
Ping-of-Death Protection	Number of suspected and rejected ICMP packets that are oversized or of an irregular size
Port Scan Protection	Number of port scans detected and blocked
Record Route IP Option	Number of frames detected with the Record Route option enabled
Security IP Option	Number of frames discarded with the IP Security option set
Src IP-based session limiting	Number of sessions dropped after the session threshold was reached
Source Route IP Option Filter	Number of IP source routes filtered
Stream IP Option	Number of packets discarded with the IP Stream identifier set
Strict Src Route IP Option	Number of packets detected with the Strict Source Route option enabled
SYN-ACK-ACK-Proxy DoS	Number of blocked packets because of the SYN-ACK-ACK-proxy DoS SCREEN option
SYN and FIN bits set	Number of packets detected with an illegal combination of flags
SYN Flood Protection	Number of SYN packets detected as part of a suspected SYN flood
SYN Fragment Detection	Number of packet fragments dropped as part of a suspected SYN fragments attack
Timestamp IP Option	Number of IP packets discarded with the Internet Timestamp option set
TCP Packet without Flag	Number of illegal packets dropped with missing or malformed flags field
Teardrop Attack Protection	Number of packets blocked as part of a Teardrop attack
UDP Flood Protection	Number of UDP packets dropped as part of a suspected UDP flood
Unknown Protocol Protection	Number of packets blocked as part of an unknown protocol
WinNuke Attack Protection	Number of packets detected as part of a suspected WinNuke attack

Table 6 on page 91 shows the hardware counters for monitoring hardware performance and packets with errors.

Table 6: Hardware Counters

Counter	Description
drop vlan	Number of packets dropped because of missing VLAN tags, an undefined sub-interface, or because VLAN trunking was not enabled when the security device was in Transparent mode
early frame	Number of counters used in an Ethernet driver buffer descriptor management
in align err	Number of incoming packets with an alignment error in the bit stream
in bytes	Number of bytes received
in coll err	Number of incoming collision packets
in crc err	Number of incoming packets with a cyclic redundancy check (CRC) error
in dma err	Number of incoming packets with a Direct Memory Access (DMA) error
in misc err	Number of incoming packets with a miscellaneous error
in no buffer	Number of unreceived packets because of unavailable buffers
in overrun	Number of transmitted overrun packets
in packets	Number of packets received
in short frame	Number of incoming packets with an Ethernet frame shorter than 64 bytes (including the frame checksum)
in underrun	Number of transmitted underrun packets
late frame	Number of counters used in an Ethernet driver buffer descriptor management
out bs pak	Number of packets held in back store while searching for an unknown MAC address When the security device forwards a packet, it first checks if the destination MAC address is in the ARP table. If it cannot find the destination MAC in the ARP table, the security device sends an ARP request to the network. If the security device receives another packet with the same destination MAC address before it receives a reply to the first ARP request, it increases the out bs pak counter by one.
out bytes	Number of bytes sent
out coll err	Number of outgoing collision packets
out cs lost	Number of dropped outgoing packets because the Carrier Sense Multiple Access/Collision Detect (CSMA/CD) protocol lost the signal
out defer	Number of deferred outgoing packets
out discard	Number of discarded outgoing packets
out heartbeat	Number of outgoing heartbeat packets
out misc err	Number of outgoing packets with a miscellaneous error
out no buffer	Number of unsent packets because of unavailable buffers
out packets	Number of packets sent
re xmt limit	Number of dropped packets when the retransmission limit was exceeded while an interface was operating at half-duplex

Table 7 on page 92 shows the flow counters for monitoring the number of packets inspected at the flow level.

Table 7: Flow Counters

Counter	Description
address spoof	Number of suspected address spoofing attack packets received
auth deny	Number of times user authentication was denied
auth fail	Number of times user authentication failed
big bkstr	Number of packets that are too big to buffer in the ARP back store while waiting for MAC-to-IP address resolution
connections	Number of sessions established since the last boot
encrypt fail	Number of failed Point-to-Point Tunneling Protocol (PPTP) packets
***icmp broadcast**	Number of ICMP broadcasts received
icmp flood	Number of ICMP packets that are counted toward the ICMP flood threshold
illegal pak	Number of packets dropped because they do not conform to the protocol standards
in arp req	Number of incoming arp request packets
in arp resp	Number of outgoing arp request packets
in bytes	Number of bytes received
in icmp	Number of Internet Control Message Protocol (ICMP) packets received
in other	Number of incoming packets that are of a different Ethernet type
in packets	Number of packets received
in self	Number of packets addressed to the Management IP address
***in un auth**	Number of unauthorized incoming TCP, UDP, and ICMP packets
***in unk prot**	Number of incoming packets using an unknown Ethernet protocol
in vlan	Number of incoming vlan packets
in vpn	Number of IPSec packets received
invalid zone	Number of packets destined for an invalid security zone
ip sweep	Number of packets received and discarded beyond the specified ip sweep threshold
land attack	Number of suspected land attack packets received
loopback drop	Number of packets dropped because they cannot be looped back through the security device. An example of a loopback session is when a host in the Trust zone sends traffic to a MIP or VIP address that is mapped to a server that is also in the Trust zone. The security device creates a loopback session that directs such traffic from the host to the MIP or VIP server.
mac relearn	Number of times that the MAC address learning table had to relearn the interface associated with a MAC address because the location of the MAC address changed
mac tbl full	Number of times that the MAC address learning table completely filled up
mal url	Number of blocked packets destined for a URL determined to be malicious
***misc prot**	Number of packets using a protocol other than TCP, UDP, or ICMP
mp fail	Number of times a problem occurred when sending a PCI message between the master processor module and the processor module
no conn	Number of packets dropped because of unavailable Network Address Translation (NAT) connections
no dip	Number of packets dropped because of unavailable Dynamic IP (DIP) addresses
no frag netpak	Number of times that the available space in the netpak buffer fell below 70 %
***no frag sess**	The number of times that fragmented sessions were greater than half of the maximum number of NAT sessions
no g-parent	Number of packets dropped because the parent connection could not be found

Counter	Description
no gate	Number of packets dropped because no gate was available
no gate sess	Number of terminated sessions because there were no gates in the firewall for them
no map	Number of packets dropped because there was no map to the trusted side
no nat vector	Number of packets dropped because the Network Address Translation (NAT) connection was unavailable for the gate
***no nsp tunnel**	Number of dropped packets sent to a tunnel interface to which no VPN tunnel is bound
no route	Number of unroutable packets received
no sa	The number of packets dropped because no Security Associations (SA) was defined
no sa policy	Number of packets dropped because no policy was associated with an SA
***no xmit vpnf**	Number of dropped VPN packets due to fragmentation
null zone	Number of dropped packets erroneously sent to an interface bound to the Null zone
nvec err	Number of packets dropped because of NAT vector error
out bytes	Number of bytes sent
out packets	Number of packets sent
out vlan	Number of outgoing vlan packets
ping of death	Number of suspected Ping of Death attack packets received
policy deny	Number of packets denied by a defined policy
port scan	Number of packets that are counted as a port scan attempt
proc sess	Number of times that the total number of sessions on a processor module exceeded the maximum threshold
sa inactive	Number of packets dropped because of an inactive SA
sa policy deny	Number of packets denied by an SA policy
sessn thresh	the threshold for the maximum number of sessions
***slow mac**	Number of frames whose MAC addresses were slow to resolve
src route	Number of packets dropped because of the filter source route option
syn frag	Number of dropped SYN packets because of a fragmentation
tcp out of seq	Number of TCP segments received whose sequence number is outside the acceptable range
tcp proxy	Number of packets dropped from using a TCP proxy such as the SYN flood protection option or user authentication
tear drop	Number of packets blocked as part of a suspected Tear Drop attack
tiny frag	Number of tiny fragmented packets received
trmn drop	Number of packets dropped by traffic management
trmng queue	Number of packets waiting in the queue
udp flood	Number of UDP packets that are counted toward the UDP flood threshold
url block	Number of HTTP requests that were blocked
winnuke	Number of WinNuke attack packets received
wrong intf	Number of session creation messages sent from a processor module to the master processor module
wrong slot	Number of packets erroneously sent to the wrong processor module

NOTE: For more information about the Carrier Sense Multiple Access/Collision Detect (CSMA/CD) protocol, see the IEEE 802.3 standard available at http://standards.ieee.org.

In this example, you view the device screen counters for the Trust zone.

WebUI

Reports > Counters > Zone Screen: Select **Trust** from the Zone drop-down list.

CLI

```
get counter screen zone trust
```

Index

M

N

P

R

S

T

U

V

W