NetScreen Concepts & Examples

ScreenOS Reference Guide

Version 2.5
P/N 093-0039-000
Rev. A

Copyright Notice

NetScreen Technologies, Inc.
2860 San Tomas Expressway
Santa Clara, CA 95051 U.S.A.
www.netscreen.com

FCC Statement

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 in a light commercial installation. This equipment generates, uses and can radiate radio frequency energy, and, if not installed and used in accordance with the instruction, may cause harmful interference to radio communications. 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.

Product License Agreement

PLEASE READ THIS LICENSE AGREEMENT (“AGREEMENTS”) CAREFULLY BEFORE USING THIS PRODUCT. BY INSTALLING AND OPERATING, YOU INDICATE YOUR ACCEPTANCE OF THE TERMS OF THIS LEGAL AND BINDING AGREEMENT AND ARE CONSENTING TO BE BOUND BY AND ARE BECOMING A PART TO THIS AGREEMENT. IF YOU DO NOT AGREE TO ALL OF THE TERMS OF THIS AGREEMENT, DO NOT START THE INSTALLATION PROCESS.

1. License Grant. This is a license, not a sales agreement, between you, the end user, and NetScreen Technologies, Inc. (“NetScreen”). The term “Firmware” includes all NetScreen and third party Firmware and software provided to you with the NetScreen product, and includes any accompanying documentation, any updates and enhancements of the Firmware and software provided to you by NetScreen, at its option. NetScreen grants to you a non-transferable (except as provided in section 3 (“Transfer”) below, non-exclusive license to use the Firmware and software in accordance with the terms set forth in this License Agreement. The Firmware and software are “in use” on the product when they are loaded into temporary memory (i.e. RAM).

2. Limitation on Use. You may not attempt and if you are a corporation, you will use best efforts to prevent your employees and contractors from attempting to, (a) modify, translate, reverse engineer decompile, disassemble, create, derivative works based on, sublicense, or distribute the Firmware or the accompanying documentation; (b) rent or lease any rights in the Firmware or software or accompanying documentation in any form to any person; or (c) remove any proprietary notice, labels, or marks on the Firmware, software, documentation, and containers.

3. Transfer. You may transfer (not rent or lease) the Firmware or software to the end user on a permanent basis, provided that: (i) the end user receives a copy of this Agreement and agrees in writing to be bound by its terms and conditions, and (ii) you at all times comply with all applicable United States export control laws and regulations.

4. Proprietary Rights. All rights, title, interest, and all copyrights to the Firmware, software, documentation, and any copy made by you remain with NetScreen. You acknowledge that no title to the intellectual property in

the Firmware and software is transferred to you and you will not acquire any rights to the Firmware except for the license as expressly set forth herein.

5. Term and Termination. The term of the license is for the duration of NetScreen's copyright in the Firmware and software. NetScreen may terminate this Agreement immediately without notice if you breach or fail to comply with any of the terms and conditions of this Agreement. You agree that, upon such termination, you will either destroy all copies of the documentation or return all materials to NetScreen. The provisions of this Agreement, other than the license granted in Section 1 (“License Grant”) shall survive termination.

6. Limited Warranty. For a period of one (1) year after delivery to Customer, NetScreen will repair or replace any defective product shipped to Customer, provided it is returned to NetScreen at Customer's expense within that period. For a period of ninety (90) days after the initial delivery of a particular product, NetScreen warrants to Customer that such product will substantially conform with NetScreen's published specifications for that product if properly used in accordance with the procedures described in documentation supplied by NetScreen. NetScreen's exclusive obligation with respect to non-conforming product shall be, at NetScreen's option, to replace the product or use diligent efforts to provide Customer with a correction of the defect, or to refund to customer the purchase price paid for the unit. Defects in the product will be reported to NetScreen in a form and with supporting information reasonably requested by NetScreen to enable it to verify, diagnose, and correct the defect. for returned product, the customer shall notify NetScreen of any nonconforming product during the warranty period, obtain a return authorization for the nonconforming product, from NetScreen, and return the nonconforming product to NetScreen's factory of origin with a statement describing the nonconformance.

NOTWITHSTANDING ANYTHING HEREIN TO THE CONTRARY, THE FOREGOING IS CUSTOMER'S SOLE AND EXCLUSIVE REMEDY FOR BREACH OF WARRANTY BY NETSCREEN WITH RESPECT TO THE PRODUCT.

The warranties set forth above shall not apply to any Product or Hardware which has been modified, repaired or altered, except by NetScreen, or which has not been maintained in accordance with any handling or operating instructions supplied by NetScreen, or which has been subjected to unusual physical or electrical stress, misuse, abuse, negligence or accidents.

THE FOREGOING WARRANTIES ARE THE SOLE AND EXCLUSIVE WARRANTIES EXPRESS OR IMPLIED GIVEN BY NETSCREEN IN CONNECTION WITH THE PRODUCT AND HARDWARE, AND NETSCREEN DISCLAIMS ALL IMPLIED WARRANTIES, INCLUDING IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. NETSCREEN DOES NOT PROMISE THAT THE PRODUCT IS ERROR-FREE OR WILL OPERATE WITHOUT INTERRUPTION.

7. Limitation of Liability. IN NO EVENT SHALL NETSCREEN OR ITS LICENSORS BE LIABLE UNDER ANY THEORY FOR ANY INDIRECT, INCIDENTAL, COLLATERAL, EXEMPLARY, CONSEQUENTIAL OR SPECIAL DAMAGES OR LOSSES SUFFERED BY YOU OR ANY THIRD PARTY, INCLUDING WITHOUT LIMITATION LOSS OF USE, PROFITS, GOODWILL, SAVINGS, LOSS OF DATA, DATA FILES OR PROGRAMS THAT MAY HAVE BEEN STORED BY ANY USER OF THE FIRMWARE. IN NO EVENT WILL NETSCREEN'S OR ITS LICENSORS' AGGREGATE LIABILITY CLAIM BY YOU, OR ANYONE CLAIMING THROUGH OR ON BEHALF OF YOU, EXCEED THE ACTUAL AMOUNT PAID BY YOU TO NETSCREEN FOR FIRMWARE.

Some jurisdictions do not allow the exclusions and limitations of incidental, consequential or special damages, so the above exclusions and limitations may not apply to you.

8. Export Law Assurance. You understand that the Firmware is subject to export control laws and regulations.

YOU MAY NOT DOWNLOAD OR OTHERWISE EXPORT OR RE-EXPORT THE FIRMWARE OR ANY UNDERLYING INFORMATION OR TECHNOLOGY EXCEPT IN FULL COMPLIANCE WITH ALL UNITED STATES AND OTHER APPLICABLE LAWS AND REGULATIONS.

9. U.S. Government Restricted Rights. If this Product is being acquired by the U.S. Government, the Product and related documentation is commercial computer Product and documentation developed exclusively at private expense, and (a) if acquired by or on behalf of civilian agency, shall be subject to the terms of this computer Firmware, and (b) if acquired by or on behalf of units of the Department of Defense (“DoD”) shall be subject to terms of this commercial computer Firmware license Supplement and its successors.

10. Tax Liability. You agree to be responsible for the payment of any sales or use taxes imposed at any time whatsoever on this transaction.

11. General. If any provisions of this Agreement are held invalid, the remainder shall continue in full force and effect. The laws of the State of California, excluding the application of its conflicts of law rules shall govern this License Agreement. This Agreement will not be governed by the United Nations Convention on the Contracts for the International Sale of Goods. This Agreement is the entire agreement between the parties as to the subject matter hereof and supersedes any other Technologies, advertisements, or understandings with respect to the Firmware and documentation. This Agreement may not be modified or altered, except by written amendment, which expressly refers to this Agreement and which, is duly executed by both parties.

You acknowledge that you have read this Agreement, understand, it and agree to be bound by its terms and conditions.

Hardware, including technical data, is subject to U.S. export laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Customer agrees to comply strictly with all such regulations and acknowledges that it has the responsibility to obtain licensed to export, re-export, or import hardware.

Table of Contents

Preface

The NetScreen-5, -10, -100 and -1000 are ASIC-based, ICSA-certified[1] Internet security appliances that integrate firewall, virtual private networking (VPN), and traffic-shaping features to provide complete protection of your local area network (LAN) when connecting to the Internet.

- **Firewall:** A firewall screens traffic crossing the boundary between a private LAN and the public network, such as the Internet.
- **VPN:** A VPN provides a secure communications channel between two or more remote network appliances.
- **Traffic Shaping:** Traffic shaping functionality allows administrative monitoring and control of traffic passing across the NetScreen® firewall to maintain a network's quality-of-service (QoS) level.

Note: *The NetScreen-1000 will soon be supported in an upcoming ScreenOS® release. Inclusion of the NetScreen-1000 in this book for ScreenOS 2.5 anticipates that release.*

1. The Internet Computer Security Association (ICSA) is an organization focused on all types of network security for Internet-connected companies. Among its many functions, ICSA provides product certification for several kinds of security products such as virus protection, firewall, PKI, intrusion detection, IPSec, and cryptography. ICSA has certified all NetScreen products for firewall and IPSec.

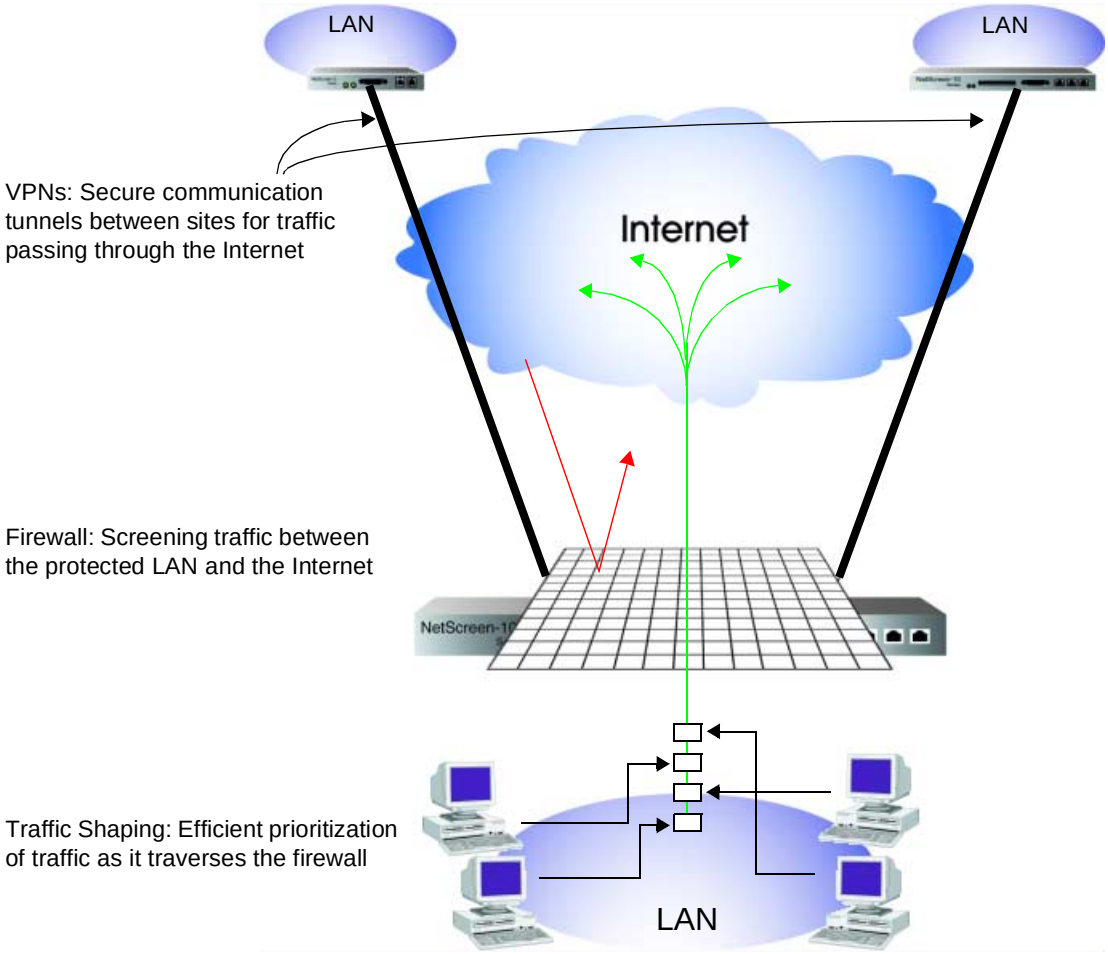

NetScreen ScreenOS version 2.5 is the operating system that provides all the features needed to set up and manage any NetScreen security appliance or system. The *NetScreen Concepts & Examples ScreenOS Reference Guide* provides a useful reference guide for configuring and managing a NetScreen appliance through the ScreenOS.

NETSCREEN DOCUMENTATION

In addition to the *NetScreen Concepts & Examples ScreenOS Reference Guide*, there are other technical publications available from NetScreen. The following are summaries of other available NetScreen technical publications:

NetScreen Concepts & Examples ScreenOS Reference Guide: A guide to the ScreenOS™ used to manage the NetScreen-5, -10, -100, and -1000. This guide presents the concepts behind NetScreen product features, and provides examples to illustrate those concepts in practice.

> ***Note:*** *For a more detailed introduction to this book, see "Concepts & Examples Manual Organization" on page xvi.*

NetScreen WebUI Reference Guide: A thorough examination of the NetScreen Web user interface (WebUI). This guide provides descriptions of all the features on the WebUI for the NetScreen-5, -10, -100, and -1000.

NetScreen CLI Reference Guide: A compendium of all the command line interface (CLI) commands. For each command, the complete syntax is presented, its arguments explained, and examples provided.

NetScreen-5, -10/100, and -1000 Installer's Guides: Instructions for connecting a NetScreen device to a network, and doing an initial configuration to set up the device in Transparent mode with an Access Policy permitting outbound traffic only. An overview of the hardware and software is also included.

> ***Note:*** *The Installer's Guides are platform specific; that is, there is a different Installer's Guide for the NetScreen-5 and -1000, and a combined guide for the NetScreen-10/100.*

NetScreen-5, -10, and -100 Getting Started Guides: A full-color glossy card with instructions for connecting the NetScreen-5, -10, or -100 to the network, and doing an initial configuration to set the operational mode of the device and create an Access Policy to permit outbound traffic only.

> ***Note:*** *Like the Installer's Guides, the Getting Started Guides are platform specific, with a different card for the NetScreen-5, -10, and -100.*

NetScreen-Global Manager User's Guide: A manual for installing and using the NetScreen-Global Manager software. NetScreen-Global Manager is software enabling centralized management of NetScreen devices.

NetScreen-Remote Administrator's Guide: A manual for installing and using the NetScreen-Remote software. NetScreen-Remote allows a remote user to connect to a NetScreen security appliance via a virtual private network (VPN) tunnel.

CONCEPTS & EXAMPLES MANUAL ORGANIZATION

The following are summaries of each of the chapters in the *NetScreen Concepts & Examples ScreenOS Reference Guide*:

Chapter 1, "Interfaces and Operational Modes" describes the various physical, logical, and virtual interfaces on NetScreen devices, and explains the concepts behind Transparent, Network Address Translation (NAT), and Route operational modes. In this and all subsequent chapters, each concept is accompanied by an illustrative example.

Chapter 2, "System Parameters" presents the concepts behind firewall settings; route table entries; Domain Name System (DNS) addressing; using Dynamic Host Configuration Protocol (DHCP) to receive and assign TCP/IP settings (NetScreen-5 and -10); URL filtering; downloading and uploading system configurations; and updating software.

Chapter 3, "Administration" explains the different means available for managing a NetScreen device both locally and remotely. This chapter also explains how to secure local and remote administrative traffic. Finally, it explains the privileges pertaining to each of the four levels of network administrators that can be defined.

Chapter 4, "Building Blocks for Access Policies and VPNs" discusses the elements used for creating Access Policies and virtual private networks (VPNs): addresses, users, and services.

Chapter 5, "Access Policies" explores the components and functions of Access Policies, and offers guidance on their creation and application.

Chapter 6, "Virtual Private Networks" explains how to establish Virtual Private Networks (VPN), using LAN-to-LAN, or Client-to-LAN communication for Manual Key and AutoKey IKE (Internet Key Exchange). This chapter contains a discussion of Public Key Infrastructure (PKI), IPSec, Certificates, Certificate Revocation Lists (CRLs), gateways, and IKE proposals. It also explains how to generate certificate requests, submit them to a trusted third party, and load the signed certificates back into your NetScreen device.

Chapter 7, "Traffic Shaping" explains how you can manage bandwidth at the interface and Access Policy levels, prioritize services, and use load balancing among Virtual IP servers to manage network traffic flow.

Chapter 8, “High Availability” explains how to cable, configure, and manage two or more NetScreen-100 or -1000 devices in a redundant group to provide high availability.

Chapter 9, “Monitoring NetScreen Devices” explains various monitoring methods and provides guidance in interpreting monitoring output.

Chapter 10, “Troubleshooting NetScreen Devices” details troubleshooting advice for NetScreen hardware and software.

Appendix A, “Glossary” provides a reference for the terms and acronyms used in the Security and Firewall field.

CONVENTIONS

This book presents two management methods for configuring a NetScreen device: the Web user interface (WebUI) and the command line interface (CLI). The conventions used for both are introduced below.

WebUI Navigation Conventions

Throughout this book, a double chevron (>>) is used to indicate navigation through the WebUI by clicking buttons, tabs, and links.

Example: Policy >> Incoming >> New Policy

To access the Policy Configuration dialog box to create an incoming Access Policy, do the following:

1. Click the **Policy** button in the menu column.
2. Click the **Incoming** tab.
3. Click the **New Policy** link.

   The Policy Configuration dialog box appears.

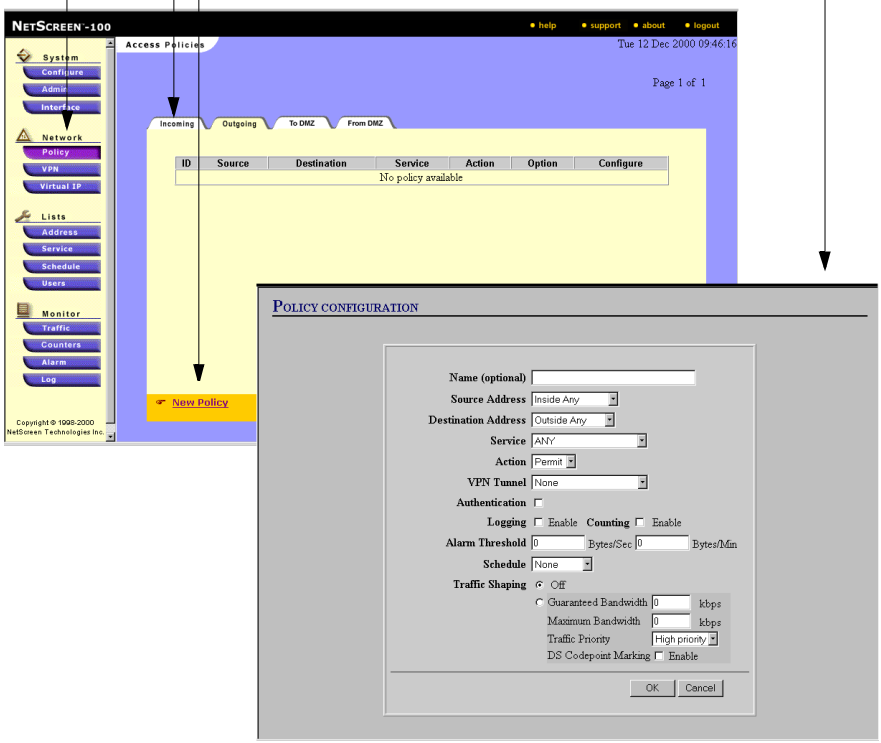

CLI Conventions

The CLI conventions are as follows:

- A parameter inside [ ] (square brackets) is optional.
- A parameter inside { } (braces) is required.
- Anything inside < > is a variable.
- If there is more than one choice for a parameter inside [ ] and { }, they are separated by *a pipe* (|). For example, [auth {md5 | sha-1}] means “choose either MD5 or SHA-1 as your authentication method.”
- IP addresses are represented by <a.b.c.d> and <e.f.g.h>.
- A subnet mask is represented by <A.B.C.D>.

1 Interfaces and Operational Modes

This chapter describes the various physical, logical, and virtual interfaces and the three operational modes supported by NetScreen devices. The chapter is organized into the following sections:

INTERFACES

All NetScreen devices have a Trusted interface and an Untrusted interface. The NetScreen-10 and -100 also have a DMZ interface. These are physical interfaces used for channeling network user traffic. Additionally, on each of the Virtual Systems supported by the NetScreen-1000 there can be one or more Sub interfaces linking a particular Virtual System to one or more virtual LANs (VLANs).

Other interfaces—some physical, some logical, and some virtual—provide exclusive channels for administrative traffic, or for communication among members in a redundant group.

Trusted Interface

The Trusted interface is a physical interface that leads to the network (usually the intranet or corporate network) protected by the NetScreen device.

Untrusted Interface

The Untrusted interface is a physical interface that leads to the network (usually the Internet) against which the NetScreen device defends. If the NetScreen device is in either NAT or Route mode (see “Interface Settings and Operational Modes” on page 1-4), the address for the Untrusted interface can be fixed (all NetScreen devices) or dynamically assigned (NetScreen-5 and -10) via Dynamic Host Configuration Protocol (DHCP). For the NetScreen-5, the address can also be provided by an ISP using Point-to-Point Protocol over Ethernet (PPPoE)—see “PPPoE” on page 2-26.

DMZ Interface

The DMZ interface is a physical interface that leads to a protected network to which access from the Untrusted side is typically granted. The DMZ, which stands for “demilitarized zone,” offers a separate and secure area on your network for receiving incoming traffic from unknown Untrusted sources—unlike the Trusted side, to which access from the Untrusted side is tightly restricted.

Web Management Interface

The Web Management interface is a logical interface that allows network administrators to manage the NetScreen device through an IP address and port number via a Web browser, such as Internet Explorer and Netscape Navigator.

Management Interface

On the NetScreen-1000, you can also manage the device through a separate physical interface—the Management (MGT) interface—moving administrative traffic outside the regular network user traffic. Separating administrative traffic from network user traffic greatly increases security and assures constant management bandwidth.

Note: *For information on configuring the device for administration, see Chapter 3, "Administration" .*

Sub Interface

On the NetScreen-1000, a Sub interface is associated with a VLAN either at the root level or for a particular Virtual System. Each Virtual System can have its own Untrusted interface[1] and one or more Sub interfaces, each Sub interface leading to a different VLAN. In essence, a Sub interface leading to a VLAN is similar to the Trusted interface leading to a protected LAN only you can have more than one Sub interface.

HA Interface

You can link two or more NetScreen-1000 devices together to form a redundant group, or cluster, through the High Availability (HA) interface. In a redundant group, one unit acts as the Master, performing the network firewall functions, while the other units act as Slaves, basically waiting to take over the firewall functions should the Master unit fail. The HA interface is a physical port used exclusively for HA functions.

Virtual HA Interface

On the NetScreen-100, a Virtual High Availability (HA) interface provides the same functionality as the HA interface on the NetScreen-1000. However, because the NetScreen-100 does not have a separate physical port exclusively used for HA traffic, the Virtual HA interface must be bound to one of the physical ports—Trusted, Untrusted, or DMZ.

1. A Virtual System can have an Untrusted interface if one is defined for it. If an Untrusted interface is not defined for a Virtual System, the Virtual System uses the Untrusted interface at the root level of the NetScreen-1000.

INTERFACE SETTINGS AND OPERATIONAL MODES

The three operational modes are Transparent, Network Address Translation (NAT), and Route. The configuration of the Trusted, Untrusted, and (on the NetScreen-10 and -100) DMZ interfaces of a NetScreen device defines which mode is in operation. Each mode offers distinct advantages.

Transparent Mode

In Transparent mode, the NetScreen device filters packets traversing the firewall without modifying any of the source or destination information in the IP packet header. All interfaces behave as though they are part of the same network, with the NetScreen device acting much like a layer-2 switch or bridge. Because it does not translate addresses, the IP addresses on the protected network must be valid, routable addresses on the Untrusted network[2], which might be the Internet. In Transparent mode, the IP addresses for the Trusted and Untrusted interfaces are set at 0.0.0.0, making the presence of the NetScreen device invisible, or "transparent," to users.

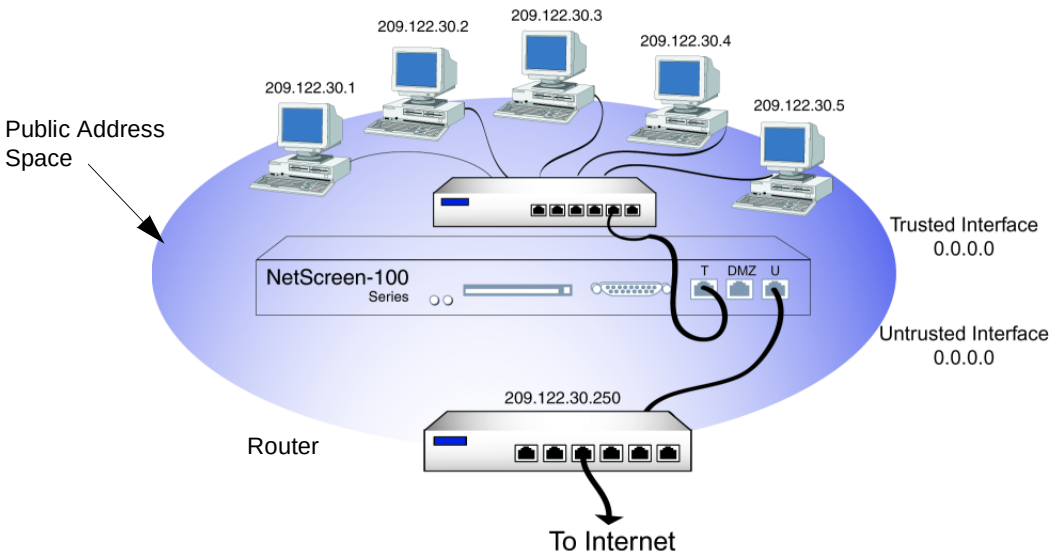

2. If the router on the Untrusted side performs NAT, then the addresses on the Trusted side can be private IP addresses.

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

- No need to reconfigure the IP settings of routers or protected servers
- No need to create Mapped or Virtual IP addresses for incoming traffic to reach protected servers
- (NetScreen-100) Because port numbers are not translated when the NetScreen-100 is operating in Transparent mode, there can be twice as many concurrent outgoing sessions (from ~64,000 to ~128,000 sessions) than when it is operating in NAT mode. The maximum number of sessions—outgoing and incoming—remains the same (~128,000) in either mode, but the maximum number of outgoing sessions is not limited to 64,000 in Transparent mode because the limit imposed by port translation is not involved.

Packet Flow Sequence

The packet flow initiating a session from a host on the Trusted side of a NetScreen device in Transparent mode to a host on the Untrusted side progresses as follows:

1. Host A, on the Trusted side of the NetScreen device, sends an IP packet to Host B, which is located on the Untrusted side.
2. The NetScreen device receives the IP packet and checks if there is an Access Policy allowing outbound TCP/IP traffic from Host A to Host B of the specified service.

3. If there is an Access Policy, the NetScreen device creates a new session in its session table.
4. The NetScreen device forwards the IP packet.

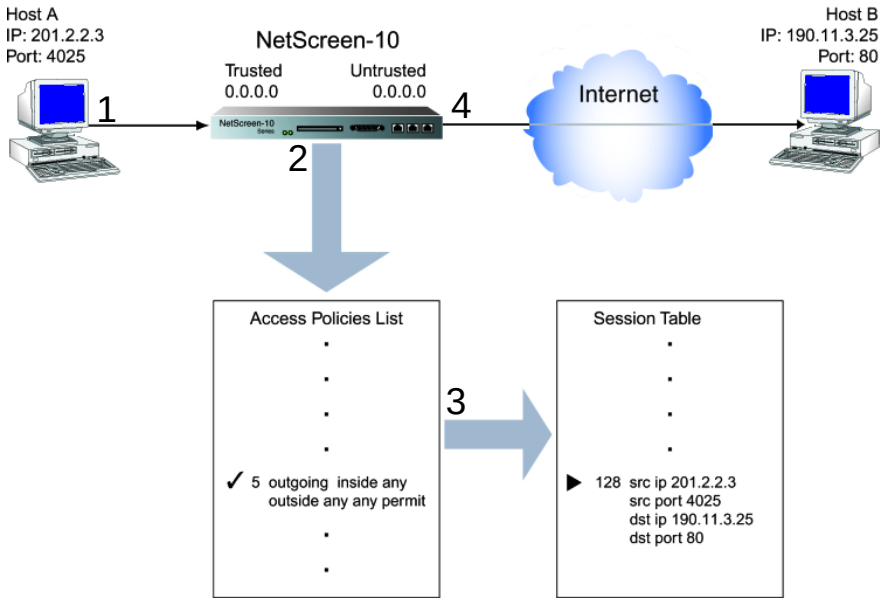

5. When the NetScreen device receives a responding IP packet from Host B, it inspects the address information in the packet header. If it matches the addressing information stored in the session table, it forwards the packet to Host A.

   The connection is established. Host B knows Host A's actual IP address and port number.

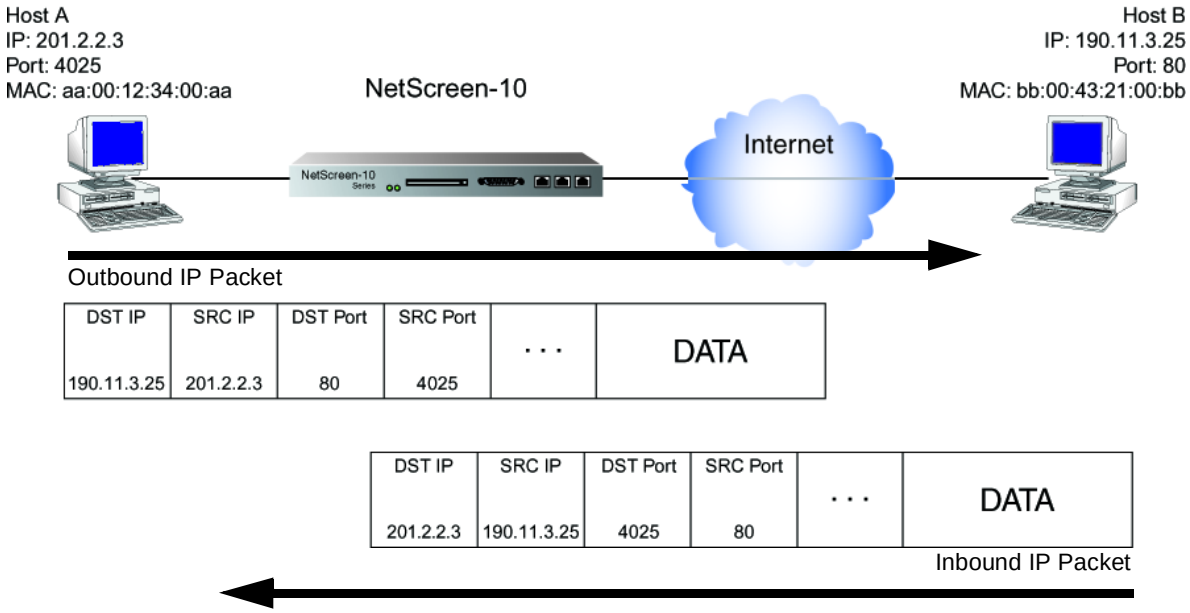

Note: *The flow sequence for any session requiring a packet to traverse the NetScreen firewall proceeds similarly; that is, when the NetScreen device receives a packet originating from any interface (Trusted, Untrusted, DMZ) and destined for any other interface, it performs the following three actions:*

1. *Checks the Access Policies list*
2. *Finding permission for the passage granted, creates an entry in the session table*
3. *Forwards the packet*

6. When Hosts A and B close their connection, the NetScreen device removes the entry from its session table. Host B can no longer send traffic to Host A.

Interface Settings

For Transparent mode, define the following interface settings, where <a.b.c.d> and <e.f.g.h> represent numbers in an IP address, <A.B.C.D> represents the numbers in a subnet mask, and <number> represents the bandwidth size in kbps:

Trusted	IP: 0.0.0.0 Subnet Mask: 0.0.0.0 Default Gateway: 0.0.0.0 Manage IP: <a.b.c.d> Traffic Bandwidth*: <number>
Untrusted	IP: 0.0.0.0 Subnet Mask: 0.0.0.0 Default Gateway: 0.0.0.0 Manage IP: <a.b.c.d> Traffic Bandwidth*: <number>
DMZ (NetScreen-10 and -100)	IP: 0.0.0.0 Subnet Mask: 0.0.0.0 Default Gateway: 0.0.0.0 Manage IP: <a.b.c.d> Traffic Bandwidth*: <number>
Web Management	System IP: <a.b.c.d> Port: <port_number>†
MGT (NetScreen-1000)	IP: <a.b.c.d> Subnet Mask: <A.B.C.D> Default Gateway: <e.f.g.h>

* Optional setting for traffic shaping

† The default port number is 80. Changing this to any number between 1024 and 32,767 is advised for discouraging unauthorized access and modifications to the configuration.

Note: *For managing the devices, you can use the System IP address, the Manage IP addresses, or the MGT IP address (NetScreen-1000).*

Example: Transparent Mode

The following example illustrates a basic configuration for a single LAN protected by a NetScreen-100 in Transparent mode. Access Policies permit outgoing traffic for all four Trusted hosts, incoming mail for the mail server, and incoming FTP for the FTP server. The device is managed through its System IP address.

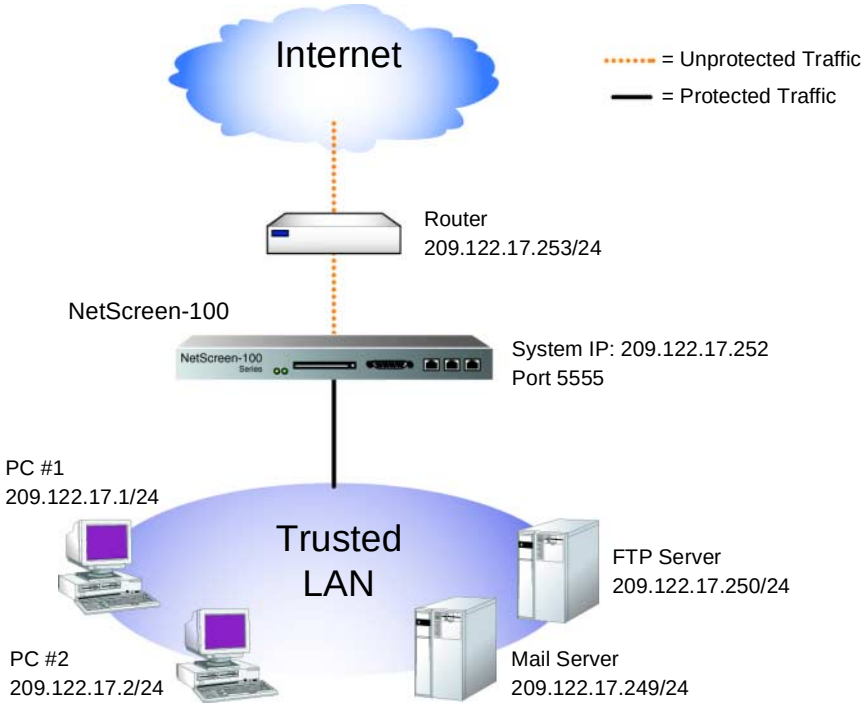

WebUI

1. Admin >> Settings: Enter the following, and then click **Apply**:

   System IP Address: 209.122.17.252

2. Admin >> Web: Enter the following, and then click **Apply**:

   Port: 5555[3]

3. When logging in to manage the device later, enter the following in the URL field of your Web browser: http://172.16.10.40:5555.

3. Interface >> Trusted >> Edit: Enter the following, and then click **Save**:

   Inside IP: 0.0.0.0

   Netmask: 0.0.0.0

   Default Gateway: 0.0.0.0

   Manage IP: 0.0.0.0

   Traffic Bandwidth: 0

4. Interface >> Untrusted >> Edit: Enter the following, and then click **Save and Reset**:[4]

   IP Address: 0.0.0.0

   Netmask: 0.0.0.0

   Default Gateway: 209.122.17.253

   Manage IP: 0.0.0.0

   Traffic Bandwidth: 0

5. (NetScreen-10 and -100) Interface >> DMZ >> Edit: Enter the following, and then click **Save**:

   Inside IP: 0.0.0.0

   Netmask: 0.0.0.0

   Default Gateway: 0.0.0.0

   Manage IP: 0.0.0.0

   Traffic Bandwidth: 0

6. (NetScreen-1000) Interface >> MGT >> Edit: Enter the following, and then click **OK**:

   MGT IP (NetScreen-1000): 0.0.0.0

   Netmask: 0.0.0.0

   Traffic Bandwidth: 0

7. Address >> Trusted >> New Address: Enter the following and then click **OK**:

   Address Name: Mail Server

   IP Address/Domain Name: 209.122.17.249

   Netmask: 255.255.255.255

4. Through the CLI, you can schedule the NetScreen-5, -10, and -100 to reset at a time that is convenient for maintaining uninterrupted network operation: **set timer <mm|dd|yyyy> <hh:mm> action reset**.

8. Address >> Trusted >> New Address: Enter the following and then click **OK**:
   Address Name: FTP Server
   IP Address/Domain Name: 209.122.17.250
   Netmask: 255.255.255.255
9. Policy >> Outgoing >> New Policy: Enter the following and then click **OK**:
   Source Address: Inside Any
   Destination Address: Outside Any
   Service: Any
   Action: Permit
10. Policy >> Incoming >> New Policy: Enter the following and then click **OK**:
    Source Address: Outside Any
    Destination Address: Mail Server
    Service: Mail
    Action: Permit
11. Policy >> Incoming >> New Policy: Enter the following and then click **OK**:
    Source Address: Outside Any
    Destination Address: FTP Server
    Service: FTP
    Action: Permit

Note: *Because PC #1 and PC #2 are not specified in an Access Policy, they do not need to be added to the Trusted Address Book. The term "Inside Any" applies to any device connected to the Trusted interface.*

CLI

1. set admin sys-ip 209.122.17.252
2. set admin port 5555[5]
3. set interface trust ip 0.0.0.0 0.0.0.0
4. set interface trust gateway 0.0.0.0
5. set interface untrust ip 0.0.0.0 0.0.0.0
6. set interface untrust gateway 209.122.17.253
7. (NetScreen-1000) set interface mgt ip 0.0.0.0 0.0.0.0
8. set address trust Mail_Server 209.122.17.249 255.255.255.255
9. set address trust FTP_Server 209.122.17.250 255.255.255.255
10. set policy outgoing “inside any” “outside any” any permit
11. set policy incoming “outside any” 209.122.17.250 255.255.255.255 mail permit
12. set policy incoming “outside any” 209.122.17.249 255.255.255.255 ftp permit
13. save

5. When logging in to manage the device later, enter the following in the URL field of your Web browser: http://172.16.10.40:5555.

Network Address Translation Mode

When in Network Address Translation (NAT) mode, the NetScreen device, acting like a layer-3 switch (or router), translates two components in the header of an outgoing IP packet traversing the firewall from the Trusted side: its source IP address and source port number. The NetScreen device replaces the source IP address of the host that sent the packet with the IP address of the Untrusted port[6] of the NetScreen device. Also, it replaces the source port number with another random port number generated by the NetScreen device.

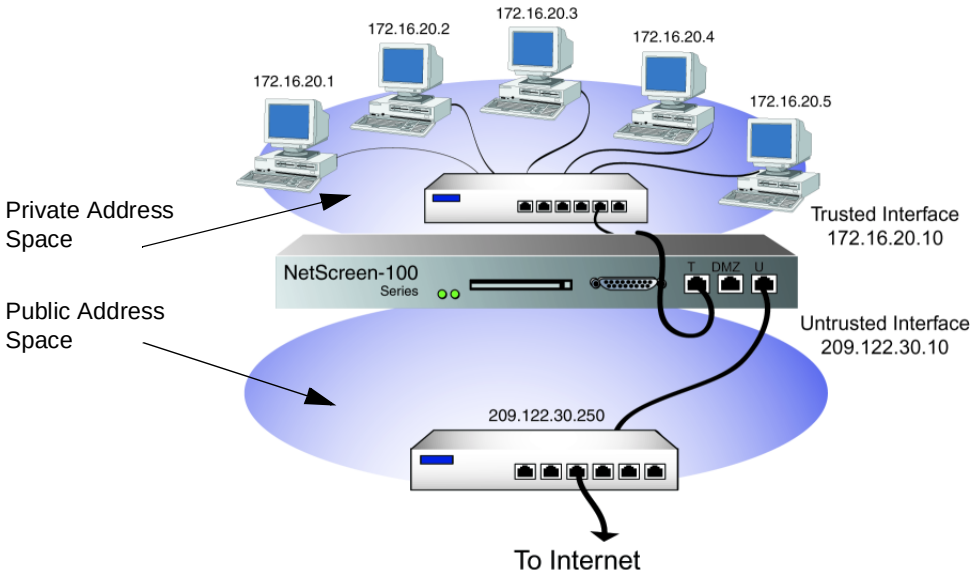

When the reply packet arrives at the NetScreen device, the device translates two components in the IP header of the incoming packet: the destination address and port number, which are translated back to the original numbers. The packet is then forwarded to its destination.

NAT adds a level of security not provided in Transparent mode: The addresses of hosts connected to the Trusted port are never exposed to the Untrusted or DMZ network.

6. If the outbound traffic is destined for the DMZ (on the NetScreen-10 and -100), then the source IP address is translated to that of the DMZ port.

Also, NAT preserves the use of Internet-routable IP addresses. With only one public, Internet-routable IP address—that of the Untrusted interface—the Trusted LAN can have a vast number of hosts with private IP addresses. The following IP address ranges are reserved for private IP networks and must not get routed on the Internet:

10.0.0.0 - 10.255.255.255
172.16.0.0 - 172.31.255.255
192.168.0.0 - 192.168.255.255

A host on the Trusted LAN can initiate traffic to the Internet, if an Access Policy allows it Internet access, but it cannot receive traffic initiated outside the firewall unless a Mapped IP or Virtual IP is set up for it. For more about Mapped and Virtual IPs, see Chapter 4, “Building Blocks for Access Policies and VPNs” .

Packet Flow Sequence: Trusted >> Untrusted

The packet flow initiating a session from a host on the Trusted side of a NetScreen device in NAT mode to a host on the Untrusted side progresses as follows:

1. Host A, on the Trusted side of the NetScreen device, sends an IP packet to Host B, which is located on the Untrusted side.
2. The NetScreen device receives the IP packet and checks if there is an Access Policy allowing outbound TCP/IP traffic from Host A to Host B with the specified service.
3. If there is such an Access Policy, the NetScreen device creates a new session in its session table and changes the source IP address and source port number on the outbound IP packet.

4. The NetScreen device forwards the modified IP packet, its true source IP address and port number unexposed.

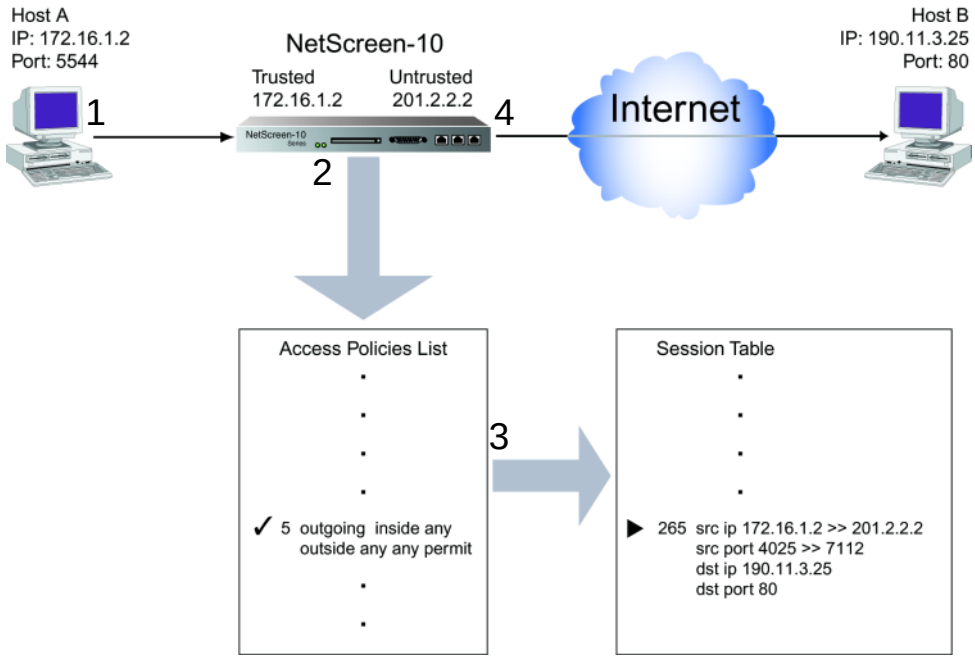

5. When the NetScreen device receives a responding IP packet from Host B, it inspects the address information in the packet header. If it matches the information stored in the session table, it forwards the packet to Host A, converting the translated IP address and port number back to the originals. (This packet filtering method is called “stateful inspection.” For more information, see “Firewall Protection” on page 2-2.)

Note: *The source port number of an outbound packet is the key to successfully directing an inbound packet to its destination. Because the source IP address for an outbound IP packet is translated to the IP address of the Untrusted interface, the distinguishing factor in the header of an inbound IP packet is its destination port number, which is unique to each inbound IP packet.*

The connection is established. Host B does not know Host A's actual IP address or port number.

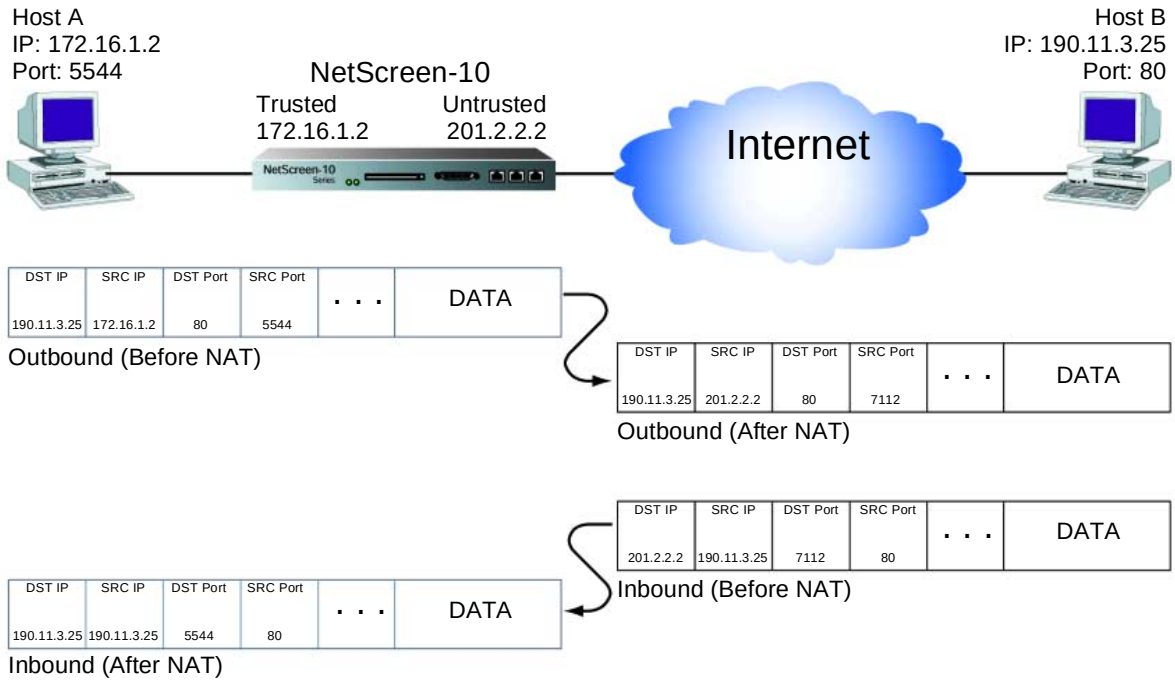

Note: *The NetScreen-10 and -100 can also perform NAT on traffic going from the Trusted interface to the DMZ.*

Packet Flow Sequence: Untrusted >> Trusted

For traffic initiated on the Untrusted side of a NetScreen device in NAT mode to reach the Trusted side, you must first create one of the following:

- A Mapped IP (MIP), mapping inbound traffic from a public IP address to a private IP address
- A Virtual IP (VIP), mapping inbound traffic from a public IP address via the port number of the incoming service to one of several possible private IP addresses
- A VPN tunnel (see Chapter 6, "Virtual Private Networks")

Mapped IP

The packet flow for a session initiating on the Untrusted side of a NetScreen device in NAT mode to a host on the Trusted side using an MIP progresses as follows:

1. Host B, on the Untrusted side of the NetScreen device, sends an IP packet to the public IP address that is mapped to the private IP address for Host A, which is located on the Trusted side.
2. The NetScreen device receives the inbound packet and checks if there is an Access Policy allowing inbound TCP/IP traffic to the MIP.
3. If there is an Access Policy, the NetScreen device creates a new session in its session table and changes the destination IP address on the inbound packet, mapping it to the private address.

4. The NetScreen device forwards the packet to Host A.

   The connection is established. Host B does not know Host A's actual IP address.

Note: *For more information on Mapped IPs, see "Mapped IP" on page 4-14.*

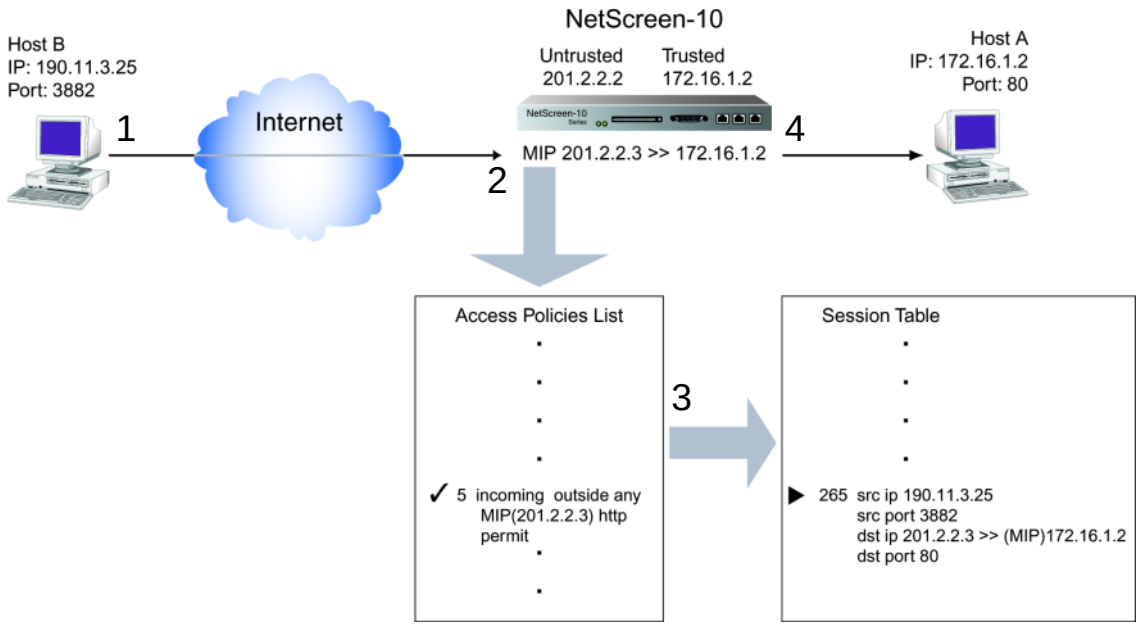

Virtual IP

The packet flow for a session initiating on the Untrusted side of a NetScreen device in NAT mode to a host on the Trusted side using an VIP progresses as follows:

1. Host B, on the Untrusted side of the NetScreen device, sends an IP packet to a public IP address in the same subnet as the Untrusted interface. (That address has been configured to route traffic to any one of several IP addresses on the Trusted side, depending on the port number carried by the incoming packet.)
2. The NetScreen device receives the inbound packet and checks if there is an Access Policy allowing inbound traffic to the VIP.
3. If there is an Access Policy, the NetScreen device creates a new session in its session table and, referring to its VIP translation table, changes the destination IP address and destination port number on the inbound packet to map it to the private IP address and port number.
4. The NetScreen device forwards the packet to Host A.

   The connection is established. Host B does not know Host A's actual IP address or port number.

Note: *For more information on Virtual IPs, see "Virtual IP" on page 4-9.*

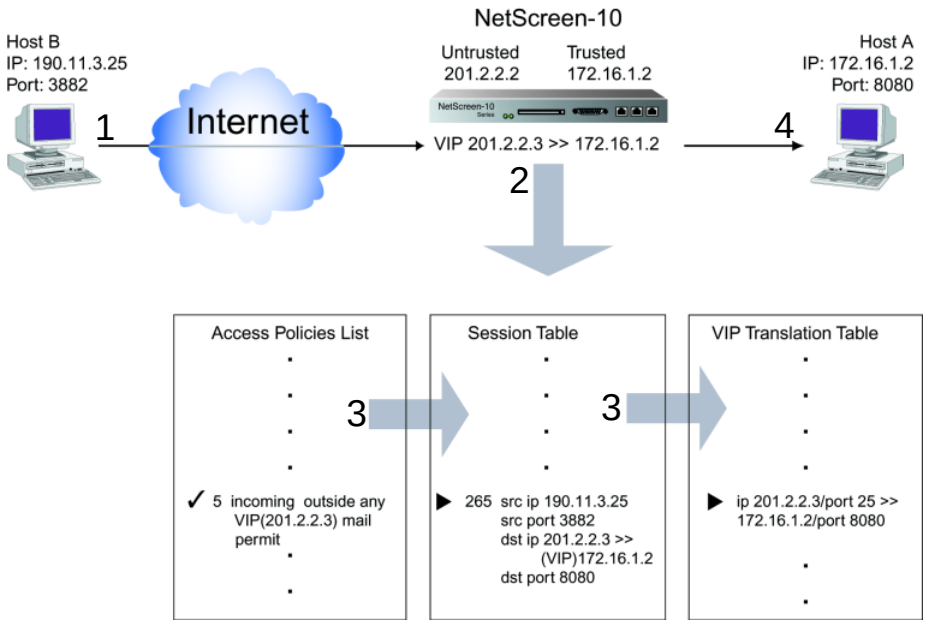

Interface Settings

For NAT mode, define the following interface settings, where <a.b.c.d>, <e.f.g.h>, and <i.j.k.l> represent numbers in an IP address, <A.B.C.D> represents the numbers in a subnet mask, and <number> represents the bandwidth size in kbps:

Trusted	IP: <a.b.c.d> Subnet Mask: <A.B.C.D> Default Gateway: <e.f.g.h> Manage IP: <i.j.k.l> Traffic Bandwidth*: <number> NAT: (select)†
Untrusted	IP: <a.b.c.d> Subnet Mask: <A.B.C.D> Default Gateway: <e.f.g.h> Manage IP: <i.j.k.l> Traffic Bandwidth*: <number>
DMZ (NetScreen-10 and -100)	IP: <a.b.c.d> Subnet Mask: <A.B.C.D> Default Gateway: <e.f.g.h> Manage IP: <i.j.k.l> Traffic Bandwidth*: <number>
Web Management	System IP: <a.b.c.d> Port: <port_number>‡
MGT (NetScreen-1000)	IP: <a.b.c.d> Subnet Mask: <A.B.C.D> Default Gateway: <e.f.g.h> Traffic Bandwidth*: <number>

* Optional setting for traffic shaping

† Selecting **NAT** for the Trusted interface defines the mode as NAT. Selecting **Route** defines the mode as Route.

‡ The default port number is 80. Changing this to any number between 1024 and 32,767 is advised for discouraging unauthorized access and modifications to the configuration.

Note: *In NAT mode, you can manage a NetScreen device from any interface—and from multiple interfaces—using the System IP address, interface IP addresses, Manage IP addresses, or the MGT IP address (NetScreen-1000).*

Example: NAT Mode

The following example illustrates a simple configuration for a LAN with a single Trusted subnet. The LAN is protected by a NetScreen-1000 in NAT mode. Access Policies permit outgoing traffic for all three Trusted hosts and incoming mail for the mail server. The incoming mail is routed to the mail server through a Virtual IP address. The device is managed through its MGT IP address.

> ***Note:*** *Compare this example with that for Route mode on page 1-27.*

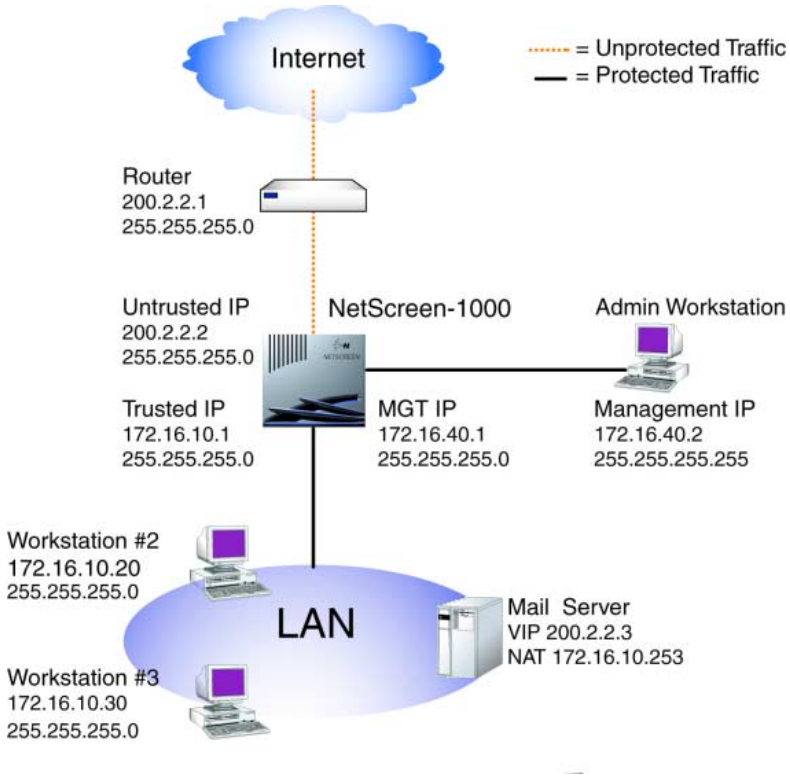

WebUI

1. Interface >> Trusted >> Edit: Enter the following, and then click **Save**:

   IP Address: 172.16.10.1

   Netmask: 255.255.255.0

   Default Gateway: 0.0.0.0

   Manage IP: 0.0.0.0

Traffic Bandwidth: 0

NAT:[7] (select)

2. Interface >> Untrusted >> Edit: Enter the following, and then click **Save**:

   IP Address:[8] 200.2.2.2

   Netmask: 255.255.255.0

   Default Gateway: 200.2.2.1

   Manage IP: 0.0.0.0

   Traffic Bandwidth: 0

3. (NetScreen-10/100) Interface >> DMZ >> Edit: Enter the following, and then click **Save and Reset**:[9]

   DMZ IP (NetScreen-10/100): 0.0.0.0

   Netmask: 0.0.0.0

   Default Gateway: 0.0.0.0

   Manage IP: 0.0.0.0

   Traffic Bandwidth: 0

4. (NetScreen-1000) Interface >> MGT >> Edit: Enter the following, and then click **OK**:

   IP Address: 172.16.40.1

   Netmask: 255.255.255.0

   Traffic Bandwidth: 0

5. Admin >> Admin >> New Management Client IP: Enter the following, and then click **OK**:

   IP Address: 172.16.40.2

   Netmask: 255.255.255.255

---

7. Selecting **NAT** determines that the NetScreen device performs NAT on traffic to and from the Trusted side.

8. If the Untrusted IP address on the NetScreen-5 and -10 is dynamically assigned by an ISP, leave the IP address and subnet mask fields empty and select DHCP. For the NetScreen-5, if the ISP is using Point-to-Point Protocol over Ethernet, select PPPoE and enter the name and password.

9. Through the CLI, you can schedule the NetScreen-5, -10, and -100 to reset at a time that is convenient for maintaining uninterrupted network operation: **set timer <mm|dd|yyyy> <hh:mm> action reset**.

6. Virtual IP >> Virtual IP 1 >> Click here to configure: Enter the following, and then click **OK**:

   Virtual IP Address: 200.2.2.3

7. Virtual IP >> New Services: Enter the following, and then click **OK**:

   Virtual Port: 25

   Service: Mail

   Map to IP: 172.16.10.253

8. Policy >> Outgoing >> New Policy: Enter the following, and then click **OK**:

   Source Address: Inside Any

   Destination Address: Outside Any

   Service: Any

   Action: Permit

9. Policy >> Incoming >> New Policy: Enter the following and then click **OK**:

   Source Address: Outside Any

   Destination Address: VIP(200.2.2.3)

   Service: Mail

   Action: Permit

CLI

1. set admin sys-ip 0.0.0.0
2. set interface trust ip 172.16.10.1 255.255.255.0
3. set interface trust NAT
4. set interface trust gateway 0.0.0.0
5. set interface untrust ip 200.2.2.2 255.255.255.0
6. set interface untrust gateway 200.2.2.1

DHCP Note: *For the NetScreen-5 and -10, if the ISP dynamically assigns the Untrusted IP address, use the following command:* ***set interface untrust dhcp***

PPPoE Note: *For the NetScreen-5, if the ISP uses PPPoE, use the* ***set pppoe*** *and* ***exec pppoe*** *commands. For more information, see the NetScreen CLI Reference Guide.*

7. (NetScreen-10/100) set interface dmz ip 0.0.0.0 0.0.0.0
8. (NetScreen-1000) set interface mgt ip 172.16.40.1 255.255.255.0
9. set admin mng-ip 172.16.40.2 255.255.255.255
10. set vip 200.2.2.3 25 mail 172.16.10.253
11. set policy outgoing “inside any” “outside any” any permit
12. set policy incoming “outside any” “vip 200.2.2.3” mail permit
13. save

Route Mode

In Route mode, the NetScreen device routes traffic between different interfaces without performing NAT; that is, the source address and port number in the IP packet header remain unchanged as it traverses the NetScreen device. Unlike NAT, the hosts on the Trusted side must have public IP addresses, and you do not need to establish Mapped and Virtual IP addresses to allow sessions initiated on the Untrusted side to reach hosts on the Trusted side. Unlike Transparent mode, the Trusted and Untrusted interfaces are on different subnets.

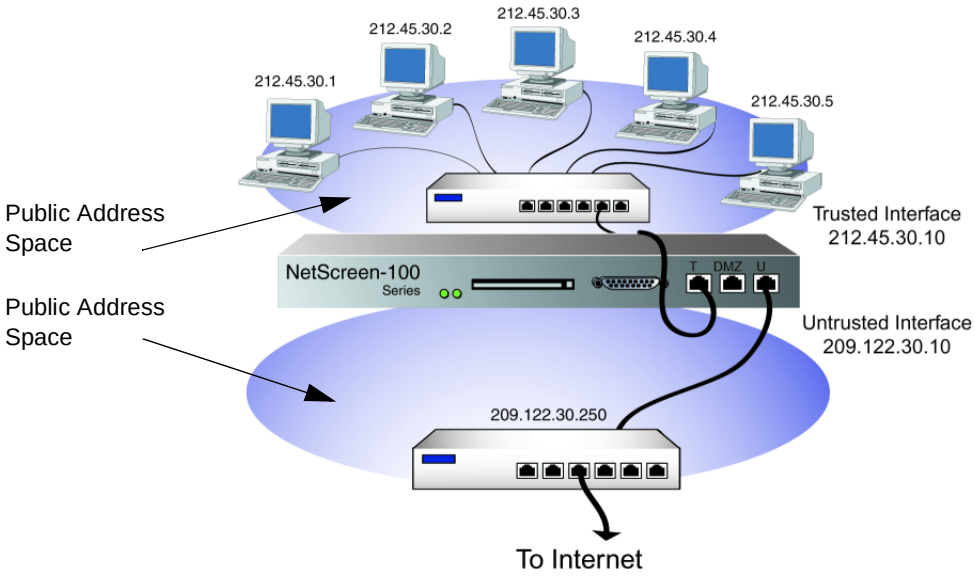

With the NetScreen-10 or -100 operating in Route mode (or Transparent mode), you do not need to set up Virtual or Mapped IPs for servers in the DMZ; the servers only require Internet-routable IP addresses. Using Route mode for the Trusted side likewise eliminates the need to create Virtual or Mapped IPs.

Interface Settings

For Route mode, define the following interface settings, where <a.b.c.d> and <e.f.g.h> represents numbers in an IP address, <A.B.C.D> represents the numbers in a subnet mask, and <number> represents the bandwidth size in kbps:

Trusted	IP: <a.b.c.d> Subnet Mask: <A.B.C.D> Default Gateway: <e.f.g.h> Manage IP: <i.j.k.l> Traffic Bandwidth*: <number> Route: (select)†
Untrusted	IP: <a.b.c.d> Subnet Mask: <A.B.C.D> Default Gateway: <e.f.g.h> Manage IP: <i.j.k.l> Traffic Bandwidth*: <number>
DMZ (NetScreen-10 and -100)	IP: <a.b.c.d> Subnet Mask: <A.B.C.D> Default Gateway: <e.f.g.h> Manage IP: <i.j.k.l> Traffic Bandwidth*: <number>
Web Management	System IP: <a.b.c.d> Port: <port_number>‡
MGT (NetScreen-1000)	IP: <a.b.c.d> Subnet Mask: <A.B.C.D> Default Gateway: <e.f.g.h> Traffic Bandwidth†: <number>

* Optional setting for traffic shaping

† Selecting **Route** for the Trusted interface defines the mode as Route. Selecting **NAT** defines the mode as NAT.

‡ The default port number is 80. Changing this to any number between 1024 and 32,767 is advised for discouraging unauthorized access and modifications to the configuration.

Note: *In Route mode, you can manage a NetScreen device from any interface—and from multiple interfaces—using the System IP address, Manage IP addresses, or interface IP addresses.*

Example: Route Mode

In the previous example for NAT mode on page 1-21, the hosts on the protected LAN have private IP addresses and a Mapped IP for the mail server. In the following example of the same network protected by a NetScreen-1000 operating in Route mode, note that the hosts have public IP addresses and that a MIP is unnecessary for the mail server. The device is managed through its MGT IP address.

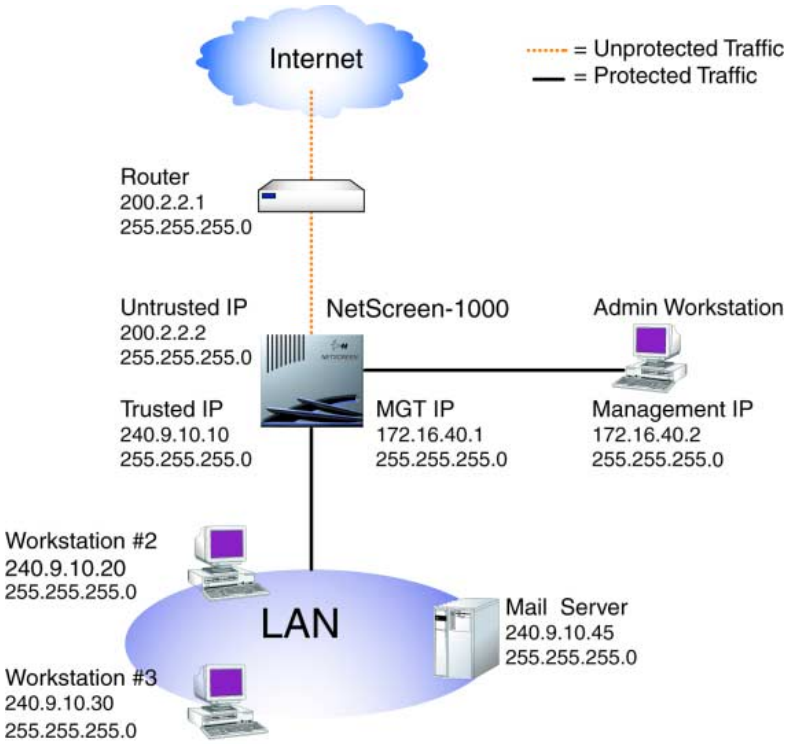

WebUI

1. Interface >> Trusted >> Edit: Enter the following, and then click **Save**:

   IP Address: 240.9.10.10

   Netmask: 255.255.255.0

   Default Gateway: 0.0.0.0

   Manage IP: 0.0.0.0

   Traffic Bandwidth: 0

   Route:[10] (select)

2. Interface >> Untrusted >> Edit: Enter the following, and then click **Save**:

   IP Address:[11] 200.2.2.2

   Netmask: 255.255.255.0

   Default Gateway: 200.2.2.1

   Manage IP: 0.0.0.0

   Traffic Bandwidth: 0

3. Interface >> DMZ (NetScreen-10/100) >> Edit: Enter the following, and then click **Save**:

   DMZ IP: 0.0.0.0

   Netmask: 0.0.0.0

   Default Gateway: 0.0.0.0

   Manage IP: 0.0.0.0

   Traffic Bandwidth: 0

4. (NetScreen-1000) Interface >> MGT >> Edit: Enter the following, and then click **Save and Reset**:

   MGT IP: 172.16.40.1

   Netmask 255.255.255.0

10. Selecting **Route** determines that the NetScreen device operates in Route mode, without performing NAT on traffic to or from the Trusted side.

11. If the Untrusted IP address on the NetScreen-5 and -10 is dynamically assigned by an ISP, leave the IP address and subnet mask fields empty and select DHCP. For the NetScreen-5, if the ISP is using Point-to-Point Protocol over Ethernet, select PPPoE and enter the name and password.

5. Admin >> Admin: Enter the following, and then click **Apply**:
   Management Client IP: 172.16.40.2
   Netmask: 255.255.255.255
6. Address >> Trusted >> New Address: Enter the following and then click **OK**:
   Address Name: Mail Server
   IP Address/Domain Name: 240.9.10.45
   Netmask: 255.255.255.255
7. Policy >> Outgoing >> New Policy: Enter the following and then click **OK**:
   Source Address: Inside Any
   Destination Address: Outside Any
   Service: Any
   Action: Permit
8. Policy >> Incoming >> New Policy: Enter the following and then click **OK**:
   Source Address: Outside Any
   Destination Address: Mail Server
   Service: Mail
   Action: Permit

CLI

1. set admin sys-ip 0.0.0.0
2. set interface trust ip 240.9.10.10 255.255.255.0
3. unset interface trust NAT[12]
4. set interface trust gateway 0.0.0.0
5. set interface untrust ip 200.2.2.2 255.255.255.0
6. set interface untrust gateway 200.2.2.1
7. (NetScreen-10/100) set interface dmz ip 0.0.0.0 0.0.0.0
8. (NetScreen-1000) set interface mgt ip 172.16.40.1 255.255.255.0
9. set admin mng-ip 172.16.40.2 255.255.255.255
10. set address trust mail_server 240.9.10.45 255.255.255.0
11. set policy outgoing “inside any” “outside any” any permit
12. set policy incoming “outside any” mail_server mail permit

12. The **unset interface trust NAT** command determines that the NetScreen device operates in Route mode.

System Parameters 2

This chapter focusses on the concepts involved in establishing system parameters affecting the following areas of a NetScreen security appliance:

FIREWALL PROTECTION

NetScreen firewalls secure a network by inspecting, and then allowing or denying, all connection attempts that require crossing the Untrusted, Trusted, and DMZ (NetScreen-10 and -100) interfaces.

By default, a NetScreen firewall denies all traffic in all directions.[1] Through the creation of Access Policies, you can then control the traffic flow across an interface by defining the kinds of traffic permitted to pass from specified sources to specified destinations at scheduled times. At the broadest level, you can allow all kinds of traffic from any Trusted source to any Untrusted destination without any scheduling restrictions. At the narrowest level, you can create an Access Policy that allows only one kind of traffic between a specified server on the Trusted side and a specified client on the Untrusted side during a scheduled period of time. In the first case, the firewall keeps all Internet traffic out of the protected network while providing all Trusted hosts access to the Internet. In the second case, you completely separate the two sides of the firewall except for a single hole connecting a point on the Trusted side to another point on the Untrusted side.

Broadly defined Internet Access: Any service from any point on the Trusted side to any point on the Untrusted side at any time

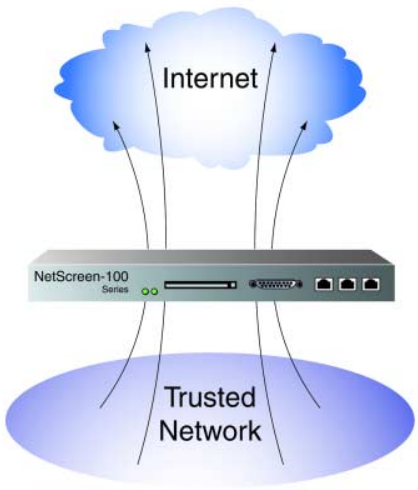

Narrowly defined Internet Access: SMTP service from a mail server on the Trusted side to a mail server on the Untrusted side from 5:00 AM to 7:00 PM

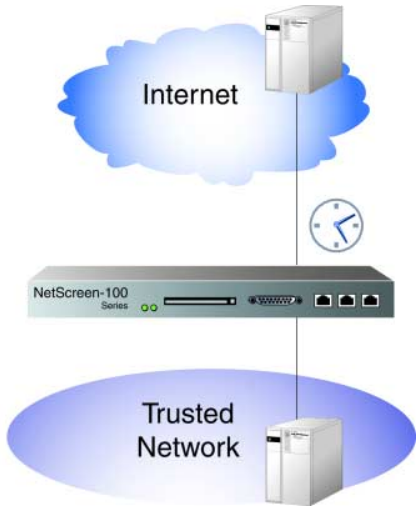

1. The NetScreen-5 default Access Policy denies all inbound traffic but allows all outbound traffic.

> ***Note:*** *For more information on the creation and application of Access Policies, see Chapter 5, "Access Policies" .*

To secure all connection attempts originating from Trusted hosts, NetScreen devices use a dynamic packet filtering method known as stateful inspection. Using this method, the NetScreen device notes various components in an outgoing TCP packet header— source and destination IP addresses, source and destination port numbers, and packet sequence numbers—and maintains the state of each TCP session traversing the firewall. (The NetScreen device also modifies session states based on changing elements such as dynamic port changes or session termination.) When a responding TCP packet arrives, the NetScreen device compares the information reported in its header with the state of its associated session stored in the inspection table. If they match, the incoming packet is allowed to pass the firewall. If the two do not match, the packet is dropped.

To protect against attacks from the Untrusted interface, you can enable defense mechanisms that can detect and deflect over a dozen common network attacks:

- **SYN Flood:** A SYN flood attack occurs when a network becomes so overwhelmed by SYN packets initiating uncompletable connection requests that it can no longer process legitimate connection requests, resulting in a denial of service (DoS).

- **ICMP Flood:** An ICMP flood occurs when ICMP pings overload a system with so many echo requests that the system expends all its resources responding until it can no longer process valid network traffic. After enabling the ICMP flood protection feature, you can set a threshold that once exceeded invokes the ICMP flood attack protection feature. (The default threshold value is 1000 packets per second.) If the threshold is exceeded, the NetScreen device ignores further ICMP echo requests for the remainder of that second.

- **UDP Flood:** Similar to the ICMP flood, UDP flooding occurs when UDP packets are sent with the purpose of slowing down the system to the point that it can no longer handle valid connections. After enabling the UDP flood protection feature, you can set a threshold that once exceeded invokes the UDP flood attack protection feature. (The default threshold value is 1000 packets per second.) If the threshold is exceeded, the NetScreen device ignores further UDP packets for the remainder of that second.

- **Ping of Death:** The TCP/IP specification requires a specific packet size for datagram transmission. Many ping implementations allow the user to specify a larger packet size if desired. A grossly oversized ICMP packet can trigger a range of adverse system reactions such as denial of service (DoS), crashing, freezing, and rebooting. If you enable the NetScreen device to do so, it can detect and reject such oversized and irregular packet sizes.

- **IP Spoofing:** Spoofing attacks occur when an attacker attempts to bypass the firewall security by imitating a valid client IP address. When IP Spoofing defense is enabled, the NetScreen device guards against this attack by analyzing the IP addresses with its own route table. If the IP address is not in the route table, traffic from that source is not allowed to communicate through the NetScreen device and any packets from that source are dropped.

- **Port Scan Attack:** Port scan attacks occur when packets are sent with different port numbers with the purpose of scanning the available services in hopes that one port will respond. The NetScreen device internally logs the number of different ports scanned from one remote source. If a remote host scans 10 ports in 0.3 seconds, NetScreen flags this as a port scan attack, and drops the connection.

- **Land Attack:** Combining a SYN attack with IP spoofing, a Land attack occurs when an attacker sends spoofed SYN packets containing the IP address of the victim as both the destination and source IP address. The receiving system responds by sending the SYN-ACK packet to itself, creating an empty connection that lasts until the idle timeout value is reached. Flooding a system with such empty connections can overwhelm the system, causing a DoS. By combining elements of the SYN flood defense and IP Spoofing protection, the NetScreen device blocks any attempts of this nature.

- **Tear Drop Attack:** Tear Drop attacks exploit the reassembly of fragmented IP packets. In the IP header, one of the options is offset. When the sum of the offset and size of one fragmented packet differ from that of the next fragmented packet, the packets overlap, and the server attempting to reassemble the packet can crash. If the NetScreen sees this discrepancy in a fragmented packet, it drops it.

- **Filter IP Source Route Option:** IP header information has an option to contain routing information that may specify a different source than the header source. Enable this option to block all IP traffic that employs the Source Route Option. Source Route Option can allow an attacker to enter a network with a false IP address and have data sent back to his real address.

- **Address Sweep Attack:** Similar to a port scan attack, an address sweep attack occurs when an attacker sends ICMP echo requests (or pings) to different destination addresses hoping that one will reply, thus uncovering an address to target. The NetScreen device internally logs the number of different addresses being pinged from one remote source. If a remote host pings 10 addresses in 0.3 seconds, NetScreen flags this as an address sweep attack, and drops the connection.

- **Block *Java/ActiveX/ZIP/EXE* Component:** Malicious Java or ActiveX components can be hidden in Web pages. When downloaded, these applets install a Trojan horse[2] on your computer. Similarly, Trojan horses can be hidden in compressed files such as .zip, .gzip, and .tar, and executable (.exe) files. Enabling this feature blocks all embedded Java and ActiveX applets from Web pages and strips attached .zip, .gzip, .tar and .exe files from e-mail.

- **Winnuke Attack:** WinNuke is a pervasive application, whose sole intent is to cause any computer on the Internet running Windows to crash. WinNuke sends out-of-band (OOB) data—usually to NetBIOS port 139—to a host with an established connection, and introduces a NetBIOS fragment overlap, which causes many machines to crash. After rebooting, the following message appears, indicating that an attack has occurred:

  An exception OE has occurred at 0028:[address] in VxD MSTCP(01) + 000041AE. This was called from 0028:[address] in VxD NDIS(01) + 00008660. It may be possible to continue normally.

  – Press any key to attempt to continue.
  – Press CTRL+ALT+DEL to restart your computer. You will lose any unsaved information in all applications.
  – Press any key to continue.

  If you enable the WinNuke attack defense mechanism on a NetScreen device, it scans any incoming Microsoft NetBIOS Session Service (port139) packets. If the NetScreen device observes that TCP URG code bit is set on one of those packets, it inspects the offset, removes the fragmented overlap, and corrects the offset as necessary to prevent an OOB error. The modified packet is then passed, and a WinNuke attack log entry is created in the Alarm Event log.

2. A Trojan horse is a program that when surreptitiously installed on a computer provides direct control of the computer to an outside party.

To enable the firewall features designed to counter the network attacks listed above, do either of the following:

WebUI

Configure >> General: Select the features you want enabled (and set threshold values for SYN Attack, ICMP Flood, and UDP Flood), and then click **Apply**:

Detect SYN Attack (and threshold value)
Detect ICMP Flood (and threshold value)
Detect UDP Flood (and threshold value)
Detect Ping of Death Attack
Detect IP Spoofing Attack
Detect Port Scan Attack
Detect Land Attack
Default Packet Deny
Detect Tear Drop Attack
Filter IP Source Route Option
Detect Address Sweep Attack
Block Java/ActiveX/ZIP/EXE Component

CLI

set firewall {applet | bypass-others-ipsec | default-deny | icmp-flood [threshold <number>] | ip-spoofing | ip-sweep [threshold <microseconds>] | land | log-self | ping-of-death | port-scan [threshold <number>] | src-route | syn-flood [alarm-threshold <number> | queue-size <number> | timeout <number>] | tear-drop | udp-flood [threshold <number>] | winnuke}

Note: *See the NetScreen CLI Reference Guide for an explanation of the* ***set firewall*** *arguments, plus examples and related commands.*

Example: SYN Flood Attack

A TCP connection is established with a triple exchange of packets known as a three-way handshake: A sends a SYN packet to B; B responds with a SYN/ACK packet; and A responds with an ACK packet. A SYN Flood attack inundates a site with SYN packets containing forged ("spoofed") IP source addresses with nonexistent or unreachable addresses. The firewall responds with SYN/ACK packets to these addresses and then waits for responding ACK packets. Because the SYN/ACK packets are sent to nonexistent or unreachable IP addresses, they never elicit responses and eventually time out.

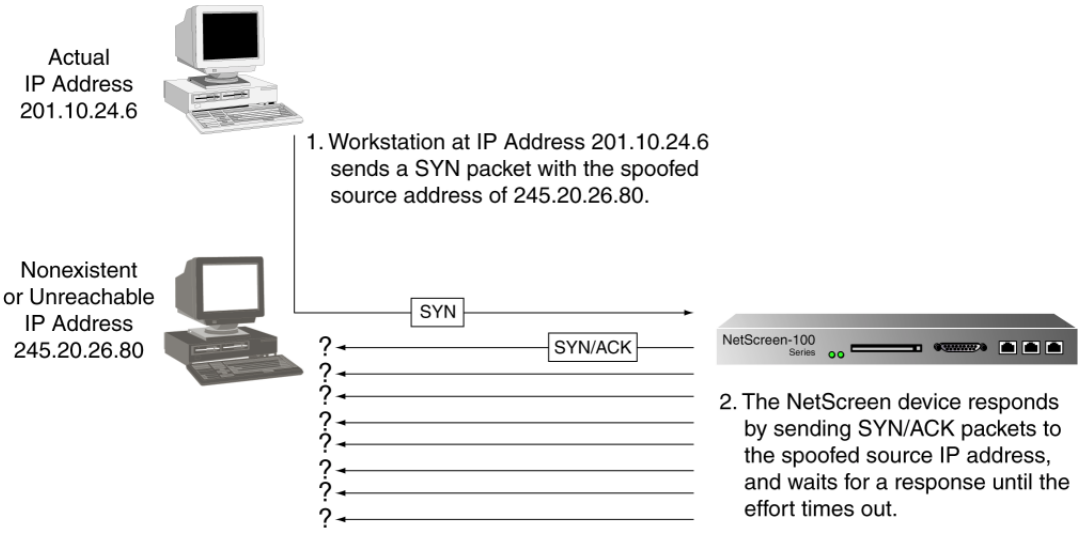

By flooding a server or host with uncompletable connections, the attacker eventually fills the host's memory buffer. Once this buffer is full, no further connections can be made and the host's operating system might be damaged. Either way, the attack disables the host and its normal operations. A SYN Flood attack is classified as a denial-of-service (DoS) attack.

SYN Flood Attack Protection

NetScreen devices can impose a limit on the number of SYN packets per second permitted to pass through the firewall. When that threshold is reached, the NetScreen device starts proxying incoming SYN packets, sending out SYN/ACK responses for the host and storing the incomplete connections in a connection queue.[3] The incomplete connections remain in the queue until the connection is completed or the request times out.

In the following illustration, the SYN threshold has been passed and the NetScreen device has begun proxying SYN packets.

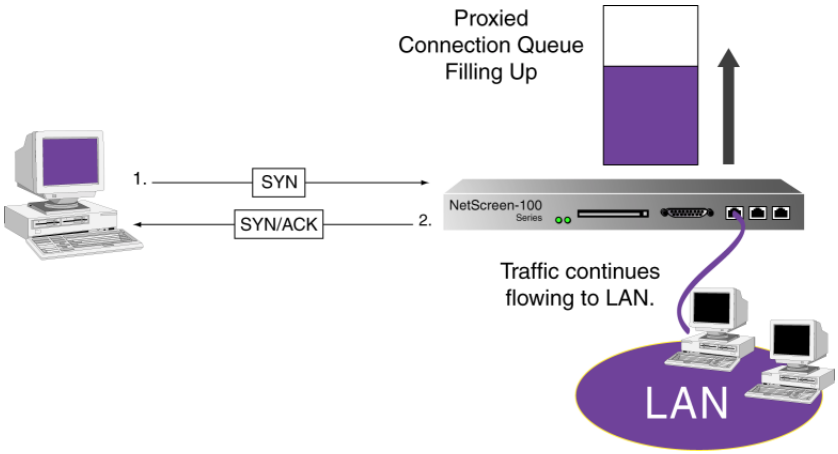

In the next illustration, the proxied connection queue has completely filled up, and new incoming SYN packets are being rejected.

3. Because the NetScreen-1000 proxies all incoming SYN packets, setting a threshold is unnecessary.

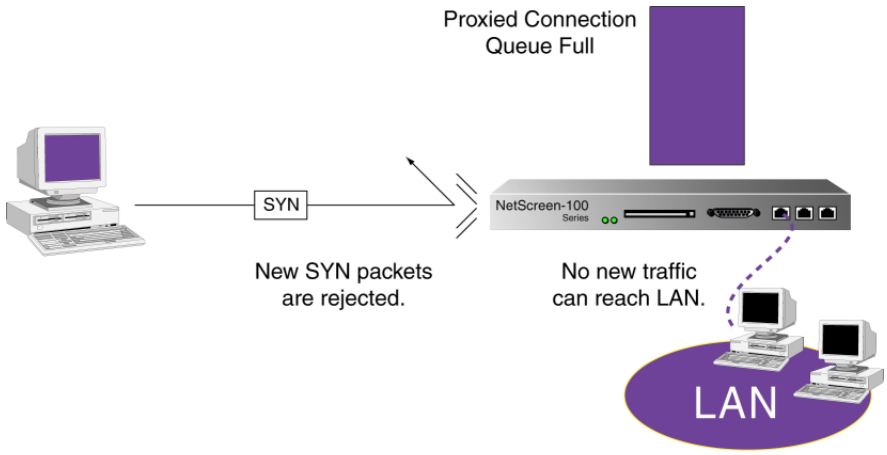

This action attempts to shield hosts on the protected network from the bombardment of incomplete three-way handshakes.

> ***Note:*** *The procedure of proxying incomplete SYN connections above a set threshold pertains only to traffic permitted by existing Access Policies. Any traffic for which an Access Policy does not exist is automatically dropped.*

WebUI: Enabling SYN Flood Attack Protection

1. Configure >> General: Enter the following settings, and then click **Apply**:

   Detect SYN Attack check box: Select

   SYN Attack Threshold (NetScreen-5/10/100): 20,000/Sec.

> ***Note:*** *The NetScreen-1000 proxies all sessions; therefore, there is no threshold to set. On the NetScreen-5/10/100, proxying is enabled when SYN Attack detection is enabled.*
>
> *Through the WebUI, you can set the threshold at which the NetScreen-5, -10, and -100 begin proxying sessions. Through the CLI, you can also set the queue length, timeout value, and alarm threshold.*

CLI: Enabling SYN Flood Attack Protection and Defining Parameters

1. Enable SYN Flood attack protection.

   ```
   set firewall syn-attack
   ```

   You can set the following four parameters for proxying uncompleted SYN connections:

2. **Threshold:** The number of SYN packets per second required to activate the SYN proxying mechanism. Although you can set the threshold at any number, you need to know the normal traffic patterns at your site to set an appropriate threshold for it. For example, if it is an e-business site that normally gets 20,000 SYN packets per second, you might want to set the threshold at 30,000/second. If a smaller site normally gets 20 SYN packets/second, you might consider setting the threshold at 40.

   ```
   set syn-threshold <number>
   ```

3. **Queue size:** The number of proxied connection requests held in the proxied connection queue before the system starts rejecting new connection requests. The longer the queue size, the longer the NetScreen device needs to scan the queue to match a valid ACK response to a proxied connection request. This can slightly slow the initial connection establishment; however, because the time to begin data transfer is normally far greater than any minor delays in initial connection setup, users would not see a noticeable difference. The queue size can be from 0–2000 for the NetScreen-10, and 0–20,000 for both the NetScreen-100 and -100p.

   ```
   set syn-qsize <number>
   ```

4. **Timeout:** The maximum length of time before a half-completed connection is dropped from the queue. The default is 20 seconds, and you can set the timeout from 0–50 seconds on the NetScreen-10, -100, and -100p. You might try decreasing the timeout value to a shorter length until you begin to see any dropped connections during normal traffic conditions. 20 seconds is a very conservative timeout for a threeway-handshake ACK response.

   ```
   set syn-timeout <number>
   ```

5. **Alarm:** The number of proxied, half-complete connections per second at which an alarm is entered in the Event Alarm log. The value you set for an alarm threshold triggers an alarm when the number of proxied, half-completed connections per second exceeds that value. For example, if the SYN threshold is set at 2000 SYN packets per second and the alarm at 1000, then a total of 3001 SYN packets per second are required to trigger an alarm entry in the log. More precisely:

   1. The firewall passes the first 2000 SYN packets per second that meet Access Policy requirements.
   2. The firewall proxies the next 1000 SYN packets in the same second.
   3. The 1001st proxied connection (or 3001st connection request in that second) triggers the alarm.

   If an attack persists, the Event Alarm log enters an alarm for each second of the attack until the attack stops and the queue empties.

   ```
   set syn-alarm <number>
   ```

ROUTE TABLE CONFIGURATION

The route table provides information that helps the NetScreen device direct traffic to different interfaces[4] and subnets. You need to define static routes for conditions such as the following:

- If the Trusted interface is on a subnet with more than one router leading to other subnets, you must define static routes that specify which router to use when forwarding traffic destined for those subnets.
- If the Untrusted interface is on a subnet with more than one router leading to multiple Internet connections, you must define static routes that specify which router to use for forwarding traffic to specific ISPs.
- You *must* define static routes that direct management traffic originating from the device itself (as opposed to user traffic traversing the firewall). For example, you need to define static routes directing syslog, SNMP, OneSecure, and WebTrends messages to the administrator's address, authentication requests to the RADIUS, SecurID, and LDAP servers, and URL checks to the Websense server.

Note: *When the NetScreen device is in Transparent mode, you must define a static route for management traffic from the device even if the destination is on the same subnet as the device. This route is necessary to define the interface through which to send traffic.*

4. When you set the interface IP addresses for a NetScreen device in NAT mode, the route table automatically creates static routes for traffic traversing the interfaces.

Example: Setting Up a Route Table

In the following example, a NetScreen-100 operating in NAT mode protects a multilevel network. There is both local and remote management (via NetScreen-Global Manager and Global Pro). SNMP traps and syslog reports are sent to the local administrator, located on the Trusted network, while NetScreen-Global PRO reports are sent to the remote administrator, located on the Untrusted network. A SecurID server on the DMZ is used to authenticate users, and a Websense server on the Trusted side performs URL blocking.

There must be statements in the NetScreen-100 route table specifying the destination network address and subnet mask, and the gateway IP address and interface[5] through which the NetScreen-100 directs traffic to the following destinations:

1. Default gateway to the Internet
2. Remote administrator in the 3.3.3.0/24 subnet
3. The Trusted 10.10.0.0/16 subnet
4. The Trusted 10.20.0.0/16 subnet
5. The Trusted 10.30.1.0/24 subnet
6. The DMZ 2.2.40.0/24 subnet
7. The DMZ 2.20.0.0/16 subnet

Note: *The following example assumes that you have already configured the Untrusted, DMZ, and Trusted interfaces as 2.2.2.1/24, 2.2.10.1/24, and 10.1.1.1/24 respectively.*

5. For each route table entry, there is also metric statement of either 0 or 1. This parameter specifies the priority of the route; that is, when there are multiple route entries for the same subnet in the route table, the NetScreen device uses the one with the lowest metric value. When using the WebUI, all route table entries that are automatically created when you define the Trusted, Untrusted, or DMZ interface have a value of 0, and any user-defined routes have a metric value of 1. Although you cannot redefine this value through the WebUI, the CLI does allow you to set it.

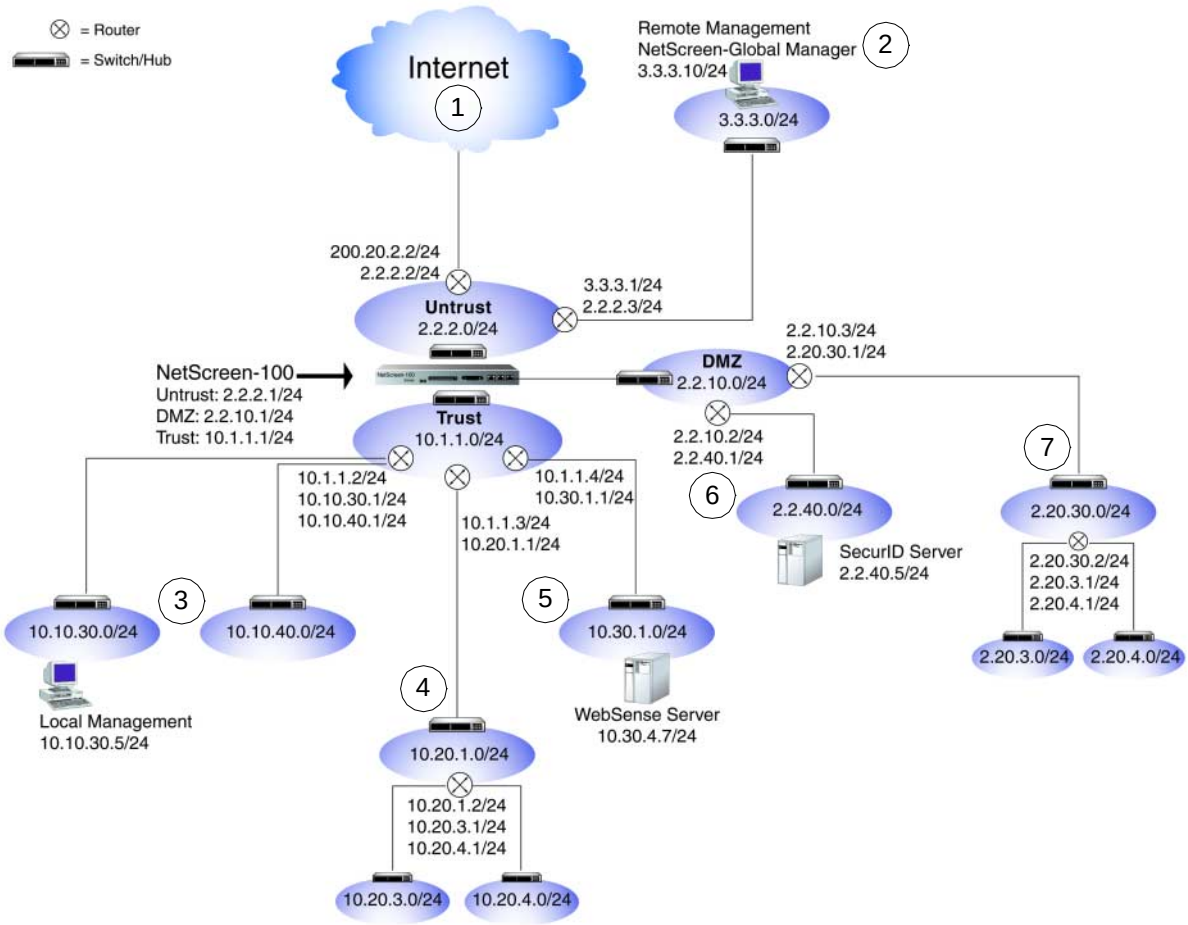

WebUI

1. Interface >> Untrusted >> Edit: Enter the following to create the Untrusted default gateway, and then click **Save and Reset**:

   Default Gateway: 2.2.2.2

2. Configure >> Route Table >> New Entry: Enter the following to direct system reports generated by the NetScreen-100 to remote management, and then click **Apply**:

   Network Address: 3.3.3.0

   Netmask: 255.255.255.0

   Gateway IP Address: 2.2.2.3

   Interface: Untrusted

3. Configure >> Route Table >> New Entry: Enter the following, and then click **Apply**:

   Network Address: 10.10.0.0
   Netmask: 255.255.0.0
   Gateway IP Address: 10.1.1.2
   Interface: Trusted

4. Configure >> Route Table >> New Entry: Enter the following, and then click **Apply**:

   Network Address: 10.20.0.0
   Netmask: 255.255.0.0
   Gateway IP Address: 10.1.1.3
   Interface: Trusted

5. Configure >> Route Table >> New Entry: Enter the following, and then click **Apply**:

   Network Address: 10.30.1.0
   Netmask: 255.255.255.0
   Gateway IP Address: 10.1.1.4
   Interface: Trusted

6. Configure >> Route Table >> New Entry: Enter the following, and then click **Apply**:

   Network Address: 2.2.40.0
   Netmask: 255.255.255.0
   Gateway IP Address: 2.2.10.2
   Interface: DMZ

7. Configure >> Route Table >> New Entry: Enter the following, and then click **Apply**:

   Network Address: 2.20.0.0
   Netmask: 255.255.0.0
   Gateway IP Address: 2.2.10.3
   Interface: DMZ

Note: *To modify a route table entry, click* ***Edit*** *under the Configure section for the entry you want to modify. The Route Table Configuration dialog box for that entry opens. Make your changes and click* ***Apply****.*

To remove an entry, click ***Remove****. A System Message appears prompting you to confirm the removal. Click* ***Yes*** *to proceed, or* ***No*** *to cancel the action.*

CLI

1. set interface untrust gateway 2.2.2.2
2. set route 3.3.3.0 255.255.255.0 interface untrust gateway 2.2.2.3
3. set route 10.10.0.0 255.255.0.0 interface trust gateway 10.1.1.2
4. set route 10.20.0.0 255.255.0.0 interface trust gateway 10.1.1.3
5. set route 10.30.1.0 255.255.255.0 interface trust gateway 10.1.1.4
6. set route 2.2.40.0 255.255.255.0 interface dmz gateway 2.2.10.2
7. set route 2.20.0.0 255.255.0.0 interface dmz gateway 2.2.10.3
8. save

DOMAIN NAME SYSTEM SUPPORT

The NetScreen device incorporates Domain Name System (DNS) support allowing you to use domain names as well as IP addresses for identifying locations. A DNS server keeps a table of the IP addresses associated with domain names. Using DNS makes it possible to reference locations by domain name (such as www.netscreen.com) in addition to using the routable IP address, which for www.netscreen.com is 209.125.148.135. DNS translation is supported in all the following programs:

- Address Book
- Syslog
- E-mail
- WebTrends®
- Websense®
- LDAP
- SecurID®
- RADIUS
- NetScreen Global-Manager

Before you can use DNS for domain name/address resolution, you must enter the addresses for DNS servers (the primary and secondary DNS servers) in the NetScreen device.

> ***Note:*** *When enabling the NetScreen-5 or -10 as a Dynamic Host Configuration Protocol server (see "DHCP" on page 2-20), you must also enter the IP addresses for DNS servers in the DHCP page on the WebUI or through the* ***set dhcp*** *command in the CLI.*

DNS Lookup

When the NetScreen device connects to the DNS server to resolve a domain name/IP address mapping, it stores that entry in its DNS status table. Some details involved in a DNS lookup follow:

- In the WebUI, the DNS lookup is performed as soon as you press **Apply** or **OK** on a page that supports DNS. In the CLI, the DNS lookup occurs when you enter a command that supports DNS.
- When a DNS lookup returns multiple entries, the address book accepts all entries. The other programs listed above accept only the first one.

- The NetScreen device reinstalls all Access Policies if it finds that anything in the domain name table has changed when you refresh a lookup using the **Refresh Now** button in the WebUI or enter the **exec dns refresh** CLI command.
- If a DNS server fails, the NetScreen device looks up everything again.
- If a lookup fails, the NetScreen device removes it from the cache table.
- If the domain name lookup fails when adding addresses to the address book, the NetScreen device displays an error message and prompts you to choose to continue adding the entry to the address book or not.

The NetScreen device must do a new lookup once a day, which you can schedule the NetScreen device to do at a specified time:

WebUI

Configure >> DNS: Enter the following, and then click **Apply**:

Lookup DNS every day at: Select check box and enter time <hh:mm>

CLI

1. set dns host schedule <hh:mm>
2. save

The DNS Status Table

The DNS status table reports all the domain names looked up, their corresponding IP addresses, whether the lookup was successful, and when the domain name/IP address was last resolved. The report format looks like the example below:

Domain Name	Corresponding IPs	Status	Last Resolved
www.yahoo.com	204.71.200.74 204.71.200.75 204.71.200.67 204.71.200.68	Success	8/13/2000 16:45:33
www.hotbot.com	209.185.151.28 209.185.151.210 216.32.228.18	Success	8/13/2000 16:45:38

Example: Defining DNS Server Addresses and Scheduling Lookups

To implement DNS functionality, the IP addresses for the DNS servers at 24.0.0.3 and 24.1.64.38 are entered in the NetScreen-5, protecting a single host in a home office. The NetScreen-5 is scheduled to refresh the DNS settings stored in its DNS status table everyday at 11:00 P.M.

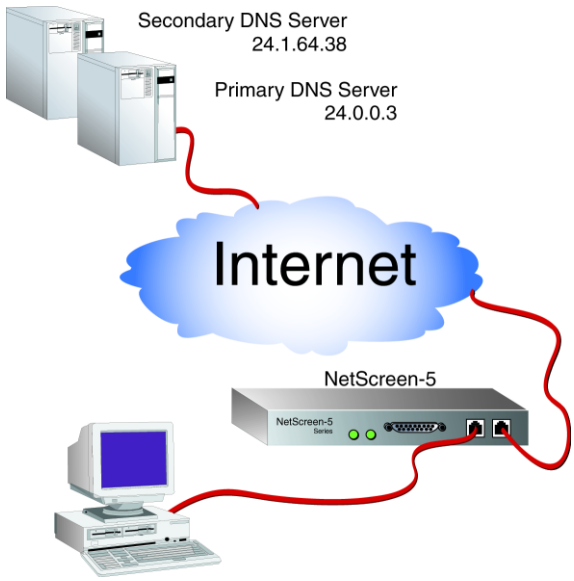

WebUI

Configure >> DNS: Enter the following, and then click **Apply**:

Primary DNS Server: 24.0.0.3
Secondary DNS Server: 24.1.64.38
Lookup DNS every day at: 23:00

CLI

1. set dns host dns1 24.0.0.3
2. set dns host dns2 24.1.64.38
3. set dns host schedule 23:00
4. save

DHCP

Dynamic Host Configuration Protocol (DHCP) was designed to reduce the demands on network administrators by automatically assigning the TCP/IP settings for the hosts on a network. Instead of requiring administrators to assign, configure, track, and change (when necessary) all the TCP/IP settings for every machine on a network, DHCP does it all automatically. Furthermore, DHCP ensures that duplicate addresses are not used, reassigns unused addresses, and automatically assigns IP addresses appropriate for the subnet on which a host is connected.

Both the NetScreen-5 and -10 can act as a DHCP client, receiving a dynamically assigned IP address for the Untrusted interface from an ISP. The NetScreen-5 and -10 can also act as a DHCP server, allocating dynamic IP addresses to hosts, acting as DHCP clients, on the Trusted network.

> ***Note:*** *While using DHCP to assign addresses to hosts on the Trusted network such as workstations and printers, you can still use fixed IP addresses for other machines such as mail servers and WINS servers.*

DHCP consists of two components: a protocol for delivering host-specific TCP/IP configuration settings and a mechanism for allocating IP addresses. The NetScreen device provides the following TCP/IP settings to each host when that host boots up:

- Default gateway IP address of the router—if there is one—that connects to the Trusted interface.
- The IP addresses of the following servers:
  - WINS servers (2):[6] A Windows Internet Naming Service (WINS) server maps a NetBIOS name used in a Windows NT network environment to an IP address used on an IP-based network.
  - DNS servers (3): A Domain Name System (DNS) server maps a uniform resource locator (URL) to an IP address.
  - SMTP server (1): A Simple Mail Transfer Protocol (SMTP) server delivers SMTP messages to a mail server, such as a POP3 server, which stores the incoming mail.

6. The number in parentheses indicates the number of servers supported.

- POP3 server (1): A Post Office Protocol version 3 (POP3) server stores incoming mail. A POP3 server must work conjointly with an SMTP server.
- News server (1): A news server receives and stores postings for news groups.

Note: *If a DHCP client to which the NetScreen device is passing the above parameters has a specified IP address, that address overrides all the dynamic information received from the DHCP server.*

When using DHCP, a NetScreen device allocates IP addresses and subnet masks in two modes:

- In Dynamic mode, the NetScreen device, acting as a DHCP server, assigns (or "leases") an IP address from an address pool[7] to a host, acting as a DHCP client. The IP address is leased for a determined period of time or until the client relinquishes the address. (To define an unlimited lease period, enter 0.)
- In Reserved mode, the NetScreen device assigns a designated IP address from an address pool exclusively to a specific client every time that client goes online.

Note: *The NetScreen device saves every IP address assigned through DHCP in flash memory. Consequently, rebooting the NetScreen device does not affect address assignments.*

7. An address pool is a defined range of IP addresses within the same subnet from which the NetScreen device can draw DHCP address assignments. You can group up to 255 IP addresses in up to 64 address pools.

Example: NetScreen-10 as DHCP Server

Using DHCP, the Trusted network behind a NetScreen-10 is sectioned into three IP address pools. All IP addresses are assigned dynamically, except for two workstations that have reserved IP addresses, and four servers that have static IP addresses. The NetScreen-10 is operating in NAT mode. The domain name is dynamic.com.

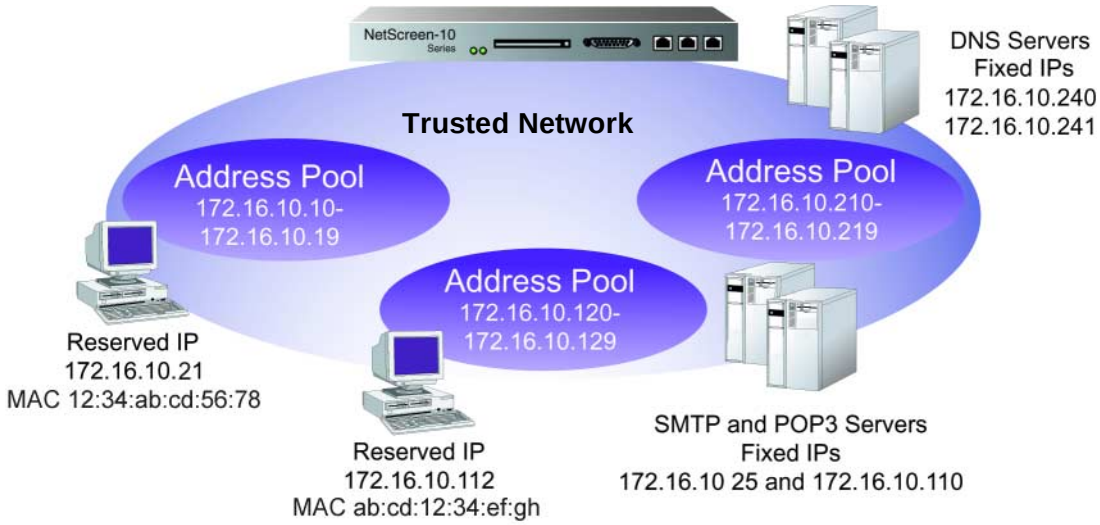

WebUI

1. Address >> Trust >> New Address: Enter the following, and then click **OK**:
   Address Name: DNS#1
   IP Address/Domain Name: 172.16.10.240
   Netmask: 255.255.255.255
   Comment: Primary DNS Server
   Trust: (select)
2. Address >> Trust >> New Address: Enter the following, and then click **OK**:
   Address Name: DNS#2
   IP Address/Domain Name: 172.16.10.241
   Netmask: 255.255.255.255
   Comment: Secondary DNS Server
   Trust: (select)

3. Address >> Trust >> New Address: Enter the following, and then click **OK**:
   Address Name: SMTP
   IP Address/Domain Name: 172.16.10.25
   Netmask: 255.255.255.255
   Comment: SMTP Server
   Trust: (select)
4. Address >> Trust >> New Address: Enter the following, and then click **OK**:
   Address Name: POP3
   IP Address/Domain Name: 172.16.10.110
   Netmask: 255.255.255.255
   Comment: POP3 Server
   Trust: (select)
5. Admin >> DHCP >> Enter the following information and click **Apply**:
   Enable DHCP Server: (select)
   Lease: Unlimited (select)
   Gateway: 0.0.0.0
   Netmask: 255.255.255.0
   Domain Name: dynamic.com
   WINS#1: 0.0.0.0
   WINS#2: 0.0.0.0
   DNS#1: 172.16.10.240
   DNS#2: 172.16.10.241
   DNS#3: 0.0.0.0
   SMTP: 172.16.10.25
   POP3: 172.16.10.110
   NEWS: 0.0.0.0
6. Admin >> DHCP >> New Address: Enter the following, and then click **OK**:
   Dynamic: (select)
   IP Address Start: 172.16.10.10
   IP Address End: 172.16.10.19
7. Admin >> DHCP >> New Address: Enter the following, and then click **OK**:
   Dynamic: (select)
   IP Address Start: 172.16.10.120
   IP Address End: 172.16.10.129

8. Admin >> DHCP >> New Address: Enter the following, and then click **OK**:
   Dynamic: (select)
   IP Address Start: 172.16.10.210
   IP Address End: 172.16.10.219
9. Admin >> DHCP >> New Address: Enter the following, and then click **OK**:
   Reserved: (select)
   IP Address: 172.16.10.21
   Ethernet Address: 1234 abcd 5678
10. Admin >> DHCP >> New Address: Enter the following, and then click **OK**:
    Reserved: (select)
    IP Address: 172.16.10.112
    Ethernet Address: abcd 1234 efgh

CLI

1. set address trust dns#1 172.16.10.240 255.255.255.255 “primary dns server”
2. set address trust dns#2 172.16.10.241 255.255.255.255 “secondary dns server”
3. set address trust snmp 172.16.10.25 255.255.255.255 “snmp server”
4. set address trust pop3 172.16.10.110 255.255.255.255 “pop3 server”
5. set dhcp server service
6. set dhcp server option domainname dynamic.com
7. set dhcp server option lease 0
8. set dhcp server option netmask 255.255.255.0
9. set dhcp server option dns1 172.16.10.240
10. set dhcp server option dns2 172.16.10.241
11. set dhcp server option smtp 172.16.10.25
12. set dhcp server option pop3 172.16.10.110
13. set dhcp server ip 172.16.10.10 to 172.16.10.19
14. set dhcp server ip 172.16.10.120 to 172.16.10.129
15. set dhcp server ip 172.16.10.210 to 172.16.10.219
16. set dhcp server ip 172.16.10.11 mac 1234abcd5678
17. set dhcp server ip 172.16.10.112 mac abcd1234efgh
18. save

Example: NetScreen-5 as DHCP Client

The Untrusted interface of the NetScreen-5 has a dynamically assigned IP address. When the NetScreen-5 requests its IP address from its ISP, it receives its IP address, subnet mask, gateway IP address, and the length of its lease on the address. The IP address of the DHCP server is 222.33.44.55.

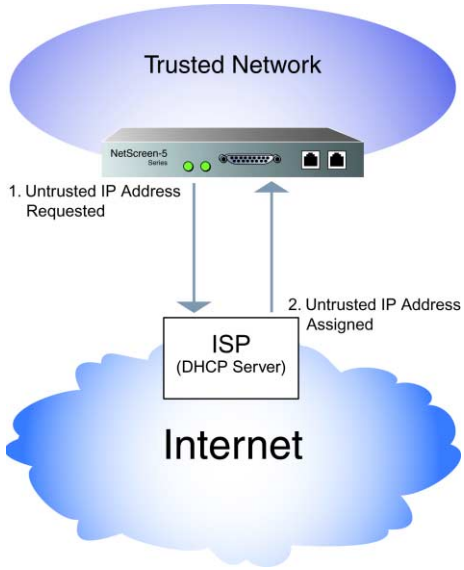

Note: *Before setting up a site for DHCP service, you must have the following:*

- *Digital subscriber line (DSL) modem and line*
- *Account with ISP*

WebUI

Interface >> Untrust >> Edit: Select **Obtain IP using DHCP**, and then click **Save and Reset**.[8]

8. You cannot specify the IP address of the DHCP server through the WebUI; however, you can do so through the CLI. Also, through the CLI, you can schedule the NetScreen-5, -10, and -100 to reset at a time that is convenient for maintaining uninterrupted network operation: **set timer <mm|dd|yyyy> <hh:mm> action reset**.

CLI

1. set interface untrust dhcp
2. set dhcp client server 222.33.44.55
3. save

PPPoE

Point-to-Point Protocol over Ethernet (PPPoE) is a protocol that allows the members of an Ethernet LAN to make individual PPP connections with their ISP by encapsulating the IP packet within the PPP payload, which is encapsulated inside the PPPoE payload.

The NetScreen-5 supports PPPoE, allowing it to operate compatibly on DSL, Ethernet Direct, and cable networks run by ISPs using PPPoE for their clients' Internet access.

Example: Setting Up PPPoE

The following example illustrates how to define the Untrusted interface of the NetScreen-5 for PPPoE connections, and how to initiate PPPoE service.

In this example, the NetScreen-5 receives a dynamically assigned IP address for its Untrusted interface from the ISP, and the NetScreen-5 also dynamically assigns IP addresses for the three hosts on its Trusted side. In this case, the NetScreen-5 acts both as a PPPoE client and DHCP server. The NetScreen-5 must be in either NAT mode or Route mode.

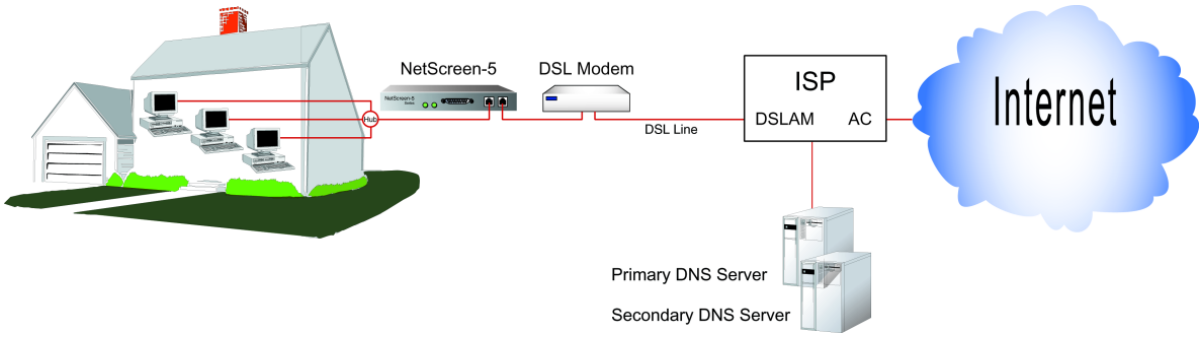

Before setting up the site in this example for PPPoE service, you must have the following:

- Digital subscriber line (DSL) modem and line
- Account with ISP
- User name and password (obtained from the ISP)

WebUI

1. Interface >> Trusted >> Edit: Enter the following, and then click **Save**:
   Address: 172.16.30.10
   Subnet Mask: 255.255.255.0
2. Interface >> Untrusted >> Edit: Enter the following:
   Obtain IP using PPPoE (select)
   User Name: <name>
   Password: <password>
3. Interface >> Untrusted >> Edit: To test your PPPoE connection, click **Connect**.

Note: *When you initiate a PPPoE connection, your ISP automatically provides the IP addresses for the Untrusted interface and the IP addresses for the Domain Name Service (DNS) servers.*

If you use a static IP address for the Untrusted interface, you must obtain the DNS servers' IP addresses and then manually enter them on the NetScreen-5 and on the Trusted hosts.

4. Admin >> DHCP: Enter the following, and then click **Apply**:

   Enable DHCP Server (select)
   Lease: 1 hour
   Gateway: 0.0.0.0
   Netmask: 0.0.0.0
   Domain Name: (leave blank)
   DNS#1: 0.0.0.0
   DNS#2: 0.0.0.0

5. Admin >> DHCP >> New Address: Enter the following, and then click **OK**:

   Dynamic: (select)
   IP Address Start: 172.16.30.2
   IP Address End: 172.16.30.5

6. Turn off the power to the DSL modem, the NetScreen-5, and the three workstations.
7. Turn on the DSL modem.
8. Turn on the NetScreen-5.

   The NetScreen-5 makes a PPPoE connection to the ISP and, through the ISP, gets the IP addresses for the DNS servers.

9. Turn on the workstations.

   The workstations automatically receive the IP addresses for the DNS servers. They get an IP address for themselves when they attempt a TCP/IP connection.

Note: *When you use DHCP to assign IP addresses to hosts on the Trusted side, the NetScreen-5 automatically forwards the IP addresses of the DNS servers that it receives from the ISP to the Trusted hosts.*

If the IP addresses for the hosts are not dynamically assigned through DHCP, you must manually enter the IP addresses for the DNS servers on each host.

Every TCP/IP connection that a Trusted host makes to the Untrusted side, automatically goes through the PPPoE encapsulation process.

CLI

1. set interface trust ip 172.16.30.10 255.255.255.0
2. set pppoe interface untrust
3. set pppoe username <name> password <password>
4. To test your PPPoE connection:

   exec pppoe connect

   get pppoe
5. set dhcp server service
6. set dhcp server ip 172.16.30.2 to 172.16.30.5
7. set dhcp server option lease 60
8. save
9. Turn off the power to the DSL modem, the NetScreen-5, and the three workstations.
10. Turn on the DSL modem.
11. Turn on the NetScreen-5.
12. Turn on the workstations.

    The workstations automatically receive the IP addresses for the DNS servers. They get an IP address for themselves when they attempt a TCP/IP connection.

    Every TCP/IP connection that a Trusted host makes to the Untrusted side, automatically goes through the PPPoE encapsulation process.

URL FILTERING CONFIGURATION

NetScreen URL filtering features the Websense Enterprise Engine, which enables you to block or permit access to different sites based on their URLs, domain names, and IP addresses. With the Websense API built directly into the NetScreen firewall, the NetScreen device creates a direct link to a Websense URL-blocking server, running on either Microsoft Windows NT 4.0 or Solaris 2.5 or 2.6.

Using Websense manager, the NetScreen administrator can do the following:

- Alter the URL-blocking database to block or allow access to any sites they choose
- Schedule different URL filtering profiles for different times of the day
- Download Websense Reporter logs of blocked or viewed URLs

Note: *For additional information about Websense, visit www.websense.com.*

To specify URL filtering options:

WebUI

Configure >> URL Filtering: Enter the following information, and then click **Apply**:

Enable URL Filtering via Websense Server: (select)

Websense Server Name: The IP address of the computer running the Websense server.

Websense Server Port: The default port for Websense is 15868. If you have changed the default port on the Websense server you must also change it on the NetScreen device. Please see your Websense documentation for full details.

Communication Timeout: The time interval, in seconds, that the NetScreen device waits for a response from the Websense filter. If Websense does not respond within the time interval, the NetScreen device will ultimately block the request.

Current Server Status: The NetScreen device reports the status of the Websense server.

URL Block Return Message: This is the message the NetScreen device returns to the user after blocking the site. You can use the message sent from the Websense server, or create a message (up to 220 characters) to be sent from the NetScreen device.

CLI

1. set url server {<domain_name> | <a.b.c.d>} <port_number> <timeout_value>
2. set url msg-type {0 | 1}
3. set url message <string>
4. set url config {enable | disable}

Note: *See the NetScreen CLI Reference Guide for an explanation of the* ***set url*** *arguments, plus examples and related commands.*

DOWNLOADING/UPLOADING SETTINGS AND SOFTWARE

You can upload and download configuration settings and software to and from a NetScreen device. The kinds of location that you upload from and download to depend on whether you use the WebUI or the CLI to perform the operation. Using the WebUI and Web browser support, you can upload and download configuration settings and upload ScreenOS software from any local directory. Through the CLI, you can upload and download settings and software from and to a TFTP server or PCMCIA card.

Saving and Importing Settings

It is good practice to backup your settings after every significant change you make. Through the WebUI, you can download the configuration to any local directory as a backup precaution. Through the CLI, you can download the configuration to a TFTP server or PCMCIA card (NetScreen-10, -100, and -1000). Should you need the saved backup configuration, you can then simply upload it to the NetScreen device.

The ability to download and upload a configuration also provides the means for mass distribution of configuration templates.

To download a configuration:

WebUI

1. Admin >> Settings: Click **Download Configuration**.
2. Browse to the location where you want to keep the configuration file, and then click **Save**.

CLI

save config to {tftp <a.b.c.d>| slot} <filename>

> ***Note:*** *On the NetScreen-1000, you must specify slot 1 or slot 2.*

To upload a configuration:

WebUI

Admin >> Settings: Specify the file name and location, and then click **Apply**:

Configure Script Upload: Type the configuration file location.

Or

Click **Browse** and navigate to the file location, select the file, and then click **Open**.

CLI

save config from {tftp <a.b.c.d> | slot} <filename>

Note: *On the NetScreen-1000, you must specify slot 1 or slot 2.*

Uploading and Downloading Software

When the NetScreen ScreenOS operating system is updated, a customer can purchase and upload it to their NetScreen device. Through the WebUI, you can upload software from a local directory. Through the CLI, you can upload the software from a TFTP server or PCMCIA card (NetScreen-10, -100, and -1000), and you can download software to a TFTP server.

> ***Note:*** *After the software is upgraded, the NetScreen device reboots. This process takes a few minutes.*

WebUI

Configure >> General: Specify the file name and location, and then click **Apply**:

Software Update: Type the software file location.

Or

Click **Browse** and navigate to the file location, select the file, and then click **Open**.

CLI

save software from {tftp <a.b.c.d>| slot} <filename>

> ***Note:*** *On the NetScreen-1000, you must specify slot 1 or slot 2.*

Through the CLI, you can also download software to a TFTP server, using the **save** command:

save software from flash to tftp <a.b.c.d> <filename>

Software Keys

The software key feature allows you to expand the capabilities of your NetScreen device without having to upgrade to a different device or system image. You can purchase a key that unlocks specified features already loaded in the software, such as the following:

- VPN tunnels
- User capacity
- Virtual Systems

Each NetScreen device ships with a standard set of features enabled and might support the activation of optional features or the increased capacity of existing features. For information regarding which features are currently available for upgrading, refer to the latest marketing literature from NetScreen.

Example: Expanding User Capacity

A small company using a single NetScreen-5 with a license for 10 users has grown to the point where it now needs an unrestricted user license. The NetScreen administrator expands the capabilities of the NetScreen-5 by obtaining a software key for an unrestricted number of users.

1. Contact the value-added reseller (VAR) who sold you the NetScreen device or contact NetScreen Technologies directly.
2. Provide the serial number of your device and state the feature option you want—an unrestricted user license, in this example.

   A combination of the serial number, the feature keyword (vpn, capacity, vsys), and the feature option keyword (<number> or “unlimited” tunnels, users, users) is used to generate the software key (for example, 7e58e876ca050192). The key is then sent to you via e-mail.

3. Enter the key through either the WebUI or CLI:

WebUI

Admin >> Software Key: Specify the path and file name, and then click **Apply and Reset**:[9]

Software Update: Type the software key file location.

Or

Click **Browse** and navigate to the software key file location, select the file, and then click **Open**.

CLI

set software-key {vpn | capacity | vsys} <key_value>

reset

9. Through the CLI, you can schedule the NetScreen-5, -10, and -100 to reset at a time that is convenient for maintaining uninterrupted network operation: **set timer <mm|dd|yyyy> <hh:mm> action reset**.

Administration 3

This chapter describes various management methods and tools, ways to secure administrative traffic, and the administrative privilege levels that you can assign to admin users:

MANAGEMENT METHODS AND TOOLS

The management methods and the tools with which to apply each method are presented in the following sections:

Web User Interface

For administrative ease and convenience, you can use the Web user interface (WebUI). NetScreen devices use Web technology that provides a Web-server interface to configure and manage the software.

To use the WebUI, you must have the following:

- Netscape® Communicator® (version 4.5 or later) or Microsoft® Internet Explorer (version 5 or later)
- TCP/IP network connection to the NetScreen device

> ***Note:*** *For a complete description of WebUI, refer to the* NetScreen WebUI Reference Guide.

HTTP

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

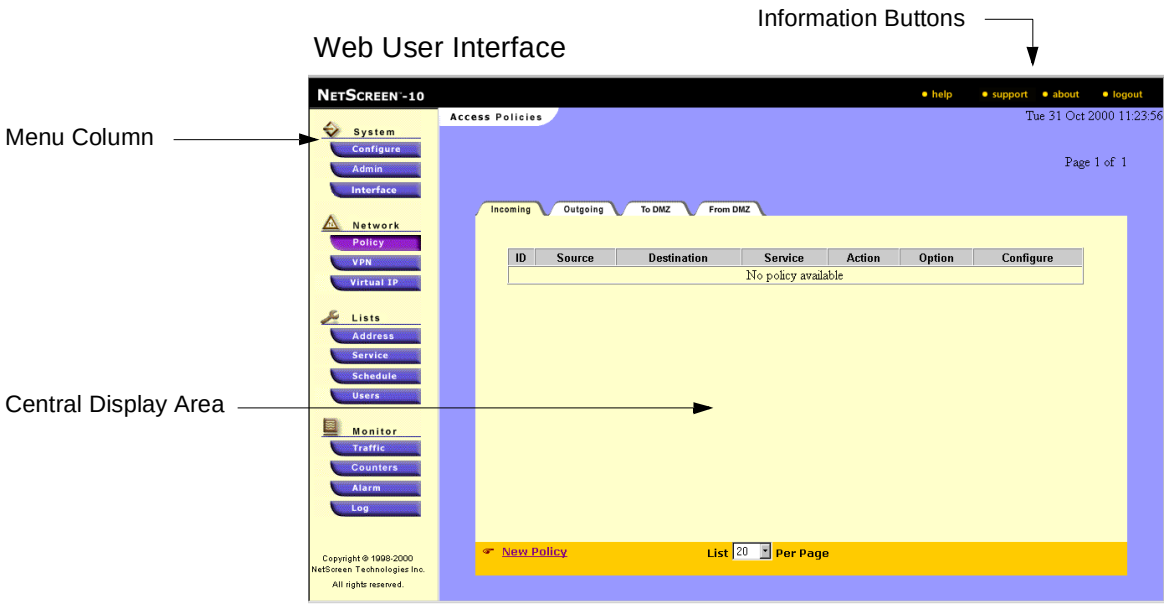

You can secure HTTP traffic by either encapsulating it in a virtual private network (VPN) tunnel or through the Secure Sockets Layer (SSL) protocol. You can also secure it by completely separating management traffic from network user traffic. You can run all administrative traffic through the MGT interface (NetScreen-1000) or devote an interface such as the DMZ (NetScreen-10 and -100) entirely to administrative traffic.

> ***Note:*** *For more information, see "Virtual Private Networks" on page 3-24, "Secure Sockets Layer" (below), and "Manage IP" on page 3-20.*

Secure Sockets Layer

Secure Sockets Layer (SSL) is a set of protocols that can provide a secure connection between a Web client and server communicating over a TCP/IP network. NetScreen ScreenOS provides:

- Web SSL support
- SSL version 3 compatibility
- Netscape Navigator 4.7x and Internet Explorer 5.x compatibility[1]
- Public Key Infrastructure (PKI) key management integration (see "PKI" on page 6-15.)

SSL is not a single protocol, but consists of the SSL Handshake Protocol (SSLHP), which allows the server and client 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)

1. Check your Web browser to see how strong the ciphers can be and which ones your browser supports. (Both the NetScreen device and your Web 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 "Cipher Strength." To obtain the advanced security package, click the **Update Information** link. In Netscape Navigator, click **Help**, **About Communicator**, and read the section about RSA®. To change the SSL configuration settings, click **Security**, **Navigator**, **Configure SSL v3**.

Independent of application protocol, SSL uses TCP to provide secure service. SSL uses certificates to authenticate first the server or both the client and the server, and then encrypt the traffic sent during the session. Before using SSL, you must first create a public/private key pair and then load a certificate. Because SSL is integrated with PKI key/certificate management, you can select the SSL certificate from one of the certificates in the certificate list. You can also use the same certificate for an IPSec VPN.

> ***Note:*** *For information on requesting and loading certificates, see "Certificates and CRLs" on page 6-17.*

NetScreen supports the following encryption algorithms for SSL:

- RC4 with 40-bit and 56-bit keys
- DES: Data Encryption Standard
- 3DES: Triple DES

NetScreen supports the same authentication algorithms for SSL as for VPNs—Message Digest version 5 (MD5) and Secure Hash Algorithm version 1 (SHA-1). The RC4 algorithms are always paired with MD5; DES and 3DES with SHA-1.

The basic steps for setting up SSL are as follows:

1. Obtain a certificate[2].

   For details on requesting and loading a certificate, see “Certificates and CRLs” on page 6-17.

2. Enable SSL management:

   Admin >> Web: Enter the following, and then click **OK**:

   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.

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

   Interface >> Trusted | Untrusted | DMZ >> Edit: Select the **SSL** management service option, and then click **Save and Reset**.

4. Connect to the NetScreen device via the SSL port:

   When you type the IP address for managing the NetScreen device in your browser's URL field, change “http” to “https”, and follow the IP address with a colon and the HTTPS (SSL) port number (for example, https://123.45.67.89:1443).

2. Be sure to specify a bit length that your Web browser also supports.

Command Line Interface

Advanced administrators can attain finer control by using the command line interface (CLI). To configure a NetScreen device with the CLI, you can use any software that emulates a VT100 terminal. With a terminal emulator, you can configure the NetScreen device using a console from any Windows®, UNIX™, or Macintosh® operating system. For remote administration through the CLI, you can use Telnet or Secure Command Shell (SCS). With a direct connection through the console port, you can use Hyperterminal®.

> ***Note:*** *For a complete listing of the CLI commands for the NetScreen devices, refer to the* NetScreen CLI Reference Guide.

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 NetScreen device. After logging in, the administrator can issue CLI commands, which are sent to the Telnet program on the NetScreen device, effectively configuring the device as if operating through a direct connection. Using Telnet to manage NetScreen devices requires the following:

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

You can secure Telnet traffic by encapsulating it in a virtual private network (VPN) tunnel or by completely separating it from network user traffic. You can run all administrative traffic through the MGT interface (NetScreen-1000) or devote an interface such as the DMZ (NetScreen-10 and -100) entirely to administrative traffic.

> ***Note:*** *For more information, see "Virtual Private Networks" on page 3-24, "Secure Sockets Layer" on page 3-3, and "Manage IP" on page 3-20.*

Secure Shell

You can use secure shell (SSH™) for secure CLI access over unsecure channels. SSH allows you to open a remote command shell[3] securely, execute commands, and copy files to or from the remote device. Secure Command Shell (SCS) is a SSH-compatible utility that allows you to remotely manage your NetScreen device without establishing a VPN.

Using SCS, you can administer NetScreen devices from an Ethernet connection or a dial-in modem. The built-in SCS server on the NetScreen device allows the SSH client, installed on the administrator's workstation, to open an instance of the NetScreen device console, which makes secure configuration and management possible.

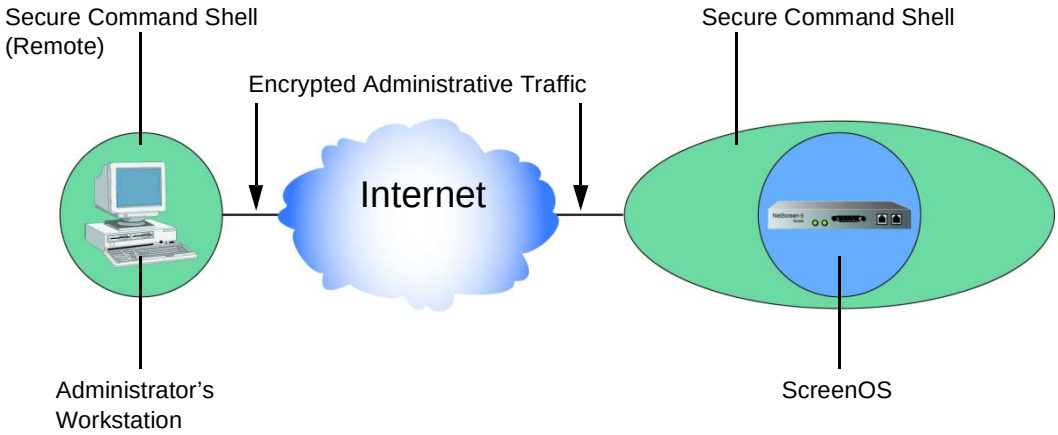

3. A command shell is an operating system's outer layer, providing an environment in which you can launch and operate programs running within the operating system's inner layer, or kernel.

Serial Console

You can manage a NetScreen device through a direct serial connection from the administrator's workstation to the NetScreen device via the Console port (Diagnostics port on the NetScreen-5). Although a direct connection is not always possible, this is surely the most secure method for managing the device.

You need the following items to create a serial connection:

- A DB-9 female to DB-25 male serial cable (NetScreen-10 and -100)
- A DB-9 female to DB-9 male serial cable (NetScreen-5)
- A MiniDIN-8 to DB-9 female serial cable (NetScreen-1000)
- 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, see the "Getting Started" chapter in the* NetScreen CLI Reference Guide.

Central Administration

If you manage large or dispersed systems, you can use either NetScreen-Global Manager independently or in conjunction with NetScreen-Global PRO to manage and configure all of your NetScreen devices from a central location.

NetScreen-Global Manager

NetScreen-Global Manager allows you to deploy and control up to 1000 NetScreen devices over multiple local-area networks (LANs) or a wide-area network (WAN) from a central location. NetScreen-Global Manager runs on Windows NT and requires network access to each device.

> ***Note:*** *For more information, refer to the* NetScreen-Global Manager User's Guide.

NetScreen-Global PRO

The NetScreen-Global PRO system allows you to control up to 10,000 NetScreen devices from a central location. NetScreen-Global PRO contains the following components:

- The database, which collects reports and statistics
- The master controller, which communicates with the database to retrieve management information and update tables
- The data collector, which collects performance- and fault-related data from the NetScreen devices

These additional components work with the NetScreen-Global PRO system:

- NetScreen devices, which provide data to the data collector
- The administration tool, which allows you to administer the system
- NetScreen-Global Manager™ Report Viewer, which displays the Global PRO reports

NetScreen-Global PRO runs on a UNIX® (Solaris™) platform.

> ***Note:*** *For more information, refer to the* NetScreen-Global PRO User's Guide.

Administrative Interface Options

You can configure the NetScreen-5, -10, -100, and -1000 to allow administration of the device through one or more interfaces. For example, you might have local management access the device through the Trusted interface and remote management through the Untrusted interface. With a NetScreen-10 or -100, you might use the DMZ interface exclusively for administration, separating management traffic completely from network user traffic for the Trusted and Untrusted interfaces.

To enable an interface to allow various methods of administration to traverse it through the WebUI and the CLI, do the following:

WebUI

Interface >> Trusted | Untrusted | DMZ: Select the following management service options, and then click **Save and Reset**[4]:

WebUI: Selecting this option allows the interface to receive HTTP traffic to manage the NetScreen device via the Web user interface (WebUI).

SSL: Selecting this option allows the interface to receive HTTPS traffic for secure management of the NetScreen device via the Web user interface (WebUI).

NS-Global: NetScreen offers two applications for central management of multisite networks—NetScreen-Global Manager and NetScreen-Global PRO. Selecting this option allows the interface to receive management traffic from NetScreen-Global Manager.

4. Through the CLI, you can schedule the NetScreen-5, -10, and -100 to reset at a time that is convenient for maintaining uninterrupted network operation: **set timer <mm|dd|yyyy> <hh:mm> action reset**.

NS-GlobalPRO: NetScreen offers two applications for central management of multisite networks—NetScreen-Global Manager and NetScreen-Global PRO. Selecting this option allows the interface to receive management traffic from NetScreen-Global PRO[5].

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.

SCS: You can administer the NetScreen device from an Ethernet connection or a dial-in modem using Secure Command Shell (SCS), which is SSH-compatible. You must have a SSH client that is compatible with Version 1.5 of the SSH protocol. These clients are available for Windows 95, Windows 98, Windows NT, Linux, and UNIX. The NetScreen device communicates with the SSH client through its built-in SCS server, which provides device configuration and management services. Selecting this option enables SCS manageability.

SNMP: The NetScreen device supports the Simple Network Management Protocol version 1.5 (SNMPv1), described in RFC-1157, and all relevant Management Information Base II (MIB II) groups, as defined in RFC-1213. Selecting this option enables SNMP manageability.

CLI

set interface {trust | untrust | dmz | mgt} manage {global | global-pro | ping | scs | snmp | ssl | telnet | web}

5. NetScreen-Global PRO requires the use of NetScreen-Global Manager, so if you want to enable this option, you also need to select the **NetScreen-Global Manager** option.

LEVELS OF ADMINISTRATION

NetScreen devices support multiple administrative users. The privileges on the NetScreen-1000 differ somewhat from those on the other NetScreen devices because of the administration of virtual systems. Therefore, the administration privileges are treated separately in the following sections.

NetScreen-5, -10, and -100 Administrators

The NetScreen-5 and -10 support up to four admin users, and the NetScreen-100 supports up to ten. (For the NetScreen-1000, see "NetScreen-1000 Administrators" on page 3-13.)

On the NetScreen-5, -10, and -100, there are three administrative levels with the following privileges:

- Level 1: Root Administrator

  The Root Administrator has complete administrative privileges.

- Level 2: Super Administrator

  The Super Administrator has the same privileges as the Root Administrator, but cannot create, modify, or remove other admin users.

- Level 3: Sub Administrator

  The Sub Administrator has viewing privileges only for the WebUI, and can only issue the get and ping CLI commands.

For any configuration changes that an administrator makes, the following information is logged:

- Name of the administrator making the change
- IP address from which the change was made
- Time of the change

NetScreen-1000 Administrators

There are four levels of administrative privilege possible for the NetScreen-1000:

- Level 1: Root Administrator
- Level 2: Super Administrator
- Level 3: Sub Administrator
- Level 4: Virtual System Administrator

Root Administrator

The Root Administrator has the following privileges:

- Manages the root system of the NetScreen device
- Adds and manages all other administrators
- Establishes and manages Virtual Systems

Super Administrator

The Super Administrator, who has root level access (similar to “root privilege” in UNIX), has the following privileges:

- Manages the root system of the NetScreen device
- 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)

Sub Administrator

The Sub 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

You can configure the NetScreen-1000 with up to 100 subsystems called virtual systems. Virtual systems are unique security domains that can be managed by their own administrators (Virtual System Administrators).

Virtual System Administrators independently manage their own virtual systems, either through CLI or the WebUI. On each virtual system, the Virtual System Administrator has the following privileges:

- Creates and edits users
- Creates and edits services
- Creates and edits Access Policies
- Creates and edits addresses
- Creates and edits VPNs
- Creates his or her login password

If necessary, a Virtual System Administrator can set up a VPN tunnel for managing a virtual system securely from a remote location, and for remote users to secure their connections to the virtual system.

Adding 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 Sub Administrator

The Root Administrator is adding a new Sub Administrator named Roger with the password 2bd21wG7 to the NetScreen-100.

WebUI

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

Name: Roger

New Password: 2bd21wG7[6]

Confirm Password: 2bd21wG7

CLI

set admin user Roger password 2bd21wG7 privilege read-only

6. The password can be up to 31 characters long. It must be alphanumeric, without any spaces or special characters.

SECURING ADMINISTRATIVE TRAFFIC

To secure the NetScreen device during setup, perform these four steps:

1. On the Web interface, change the administrative port.

   See “Changing the System IP Port Number” on page 3-16.

2. Turn off any unnecessary interface management service options.

   See “Administrative Interface Options” on page 3-10.

3. 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

   Interface >> Trusted | Untrusted | DMZ: Clear the following service options, and then click **OK**:

   Ping: A utility that enables you to determine whether a specific IP address is accessible. Selecting this option allows people to ping the IP address of the NetScreen device through the Trusted, Untrusted, or DMZ interface.

   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 NetScreen 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.

   CLI

   unset interface {trust | untrust | dmz | mgt} manage ping

   unset interface {trust | untrust | dmz | mgt} ident-reset

4. Change the user name and password for administration access.

   See “Changing the Admin Login Name and Password” on page 3-17.

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

   See “Restricting Administrative Access” on page 3-19.

Changing the System IP Port Number

Changing the port number to which the NetScreen 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 type the new port number in the URL field in your Web browser when you next attempt to contact the NetScreen device. (In the following example, the administrator needs to enter http://188.30.12.2:15522.)

Example: Changing the Port Number

In this example, the System IP is 188.30.12.2 with the standard port number 80. You change the port number from 80 to 15522.

WebUI

1. Admin >> Web >> Port: 15522
2. Click **Apply and Reset**[7].

CLI

1. set admin port 15522
2. save

7. Through the CLI, you can schedule the NetScreen-5, -10, and -100 to reset at a time that is convenient for maintaining uninterrupted network operation: **set timer <mm|dd|yyyy> <hh:mm> action reset**.

Changing the Admin Login Name and Password

By default, the initial login name for NetScreen devices is *netscreen*. The initial password is also *netscreen*. Because these have been widely published, you should change the login name and password immediately. The login name and password are both case-sensitive. Each must be one word, alphanumeric, with no symbols. 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 cannot reset or gain access to the device. It must then be returned to the factory for resetting.*

Administrative users for the NetScreen device can be authenticated using the internal database and an external RADIUS server[8]. When the admin user logs in to the NetScreen device, it first checks the local internal database for authentication. If there is no entry present, it then uses RADIUS to authenticate.The purpose of this feature is to extend the authentication schemes to the management of administrative users.

Note: *For more information about admin user levels, see "Click the Save button." on page 3-29. For more about using a RADIUS server for authentication, see "RADIUS" on page 4-19.*

8. Remote Authentication Dial-In User Service (RADIUS) is a protocol for authenticating and authorizing dial-up users. The NetScreen device can act as a client of a RADIUS server.

Example: Changing an Admin User Login Name and Password

The Root Administrator has decided to change a Super Administrator's login name from John to Smith and his password from xL7s62a1 to 3MAb99j2.

> ***Note:*** *For information on the different levels of administrators, see "Click the Save button." on page 3-29.*

WebUI

Admin >> Admin >> (John) Edit: Enter the following, and then click **OK**:

Name: Smith

New Password: 3MAb99j2

Confirm Password: 3MAb99j2

CLI

1. unset admin user John
2. set admin user Smith password 3MAb99j2 privilege all
3. save

Example: An Admin User Changing Her Own Password

Non-root users can change their own administrator password, but not their login name. In this example, a Super Administrator with the login name "starling" is changing her password from 3MAb99j2 to ru494Vq5.

WebUI

Admin >> Admin >> (starling) Edit: Enter the following, and then click **OK**:

Name: starling

Old Password: 3MAb99j2

New Password: ru494Vq5

Confirm Password: ru494Vq5

CLI

1. set admin password ru494Vq5
2. save

Restricting Administrative Access

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

Caution *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 NetScreen device will immediately terminate your current session and you will no longer be 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 NetScreen-10.

WebUI

Admin >> Admin >> New Management Client IP: Enter the following, and then click **OK**:

IP Address: 172.16.40.42

Netmask: 255.255.255.255

CLI

1. set admin manage-ip 172.16.40.42 255.255.255.255
2. 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 the NetScreen-10.

WebUI

Admin >> Admin >> New Management Client IP: Enter the following, and then click **OK**:

IP Address: 172.16.40.0

Netmask: 255.255.255.0

CLI

1. set admin manage-ip 172.16.40.0 255.255.255.0
2. save

Manage IP

The Trusted, Untrusted, and DMZ (NetScreen-10 and -100) interfaces can have two IP addresses: an interface IP address that corresponds with the physical port through which that interface connects to a network, and a Manage IP address that can be used to receive administrative traffic.

You can specify a Manage IP address for managing a NetScreen device through every available interface. Also, when a NetScreen-100 or -1000 is a slave unit in a redundant group for High Availability, you can access and configure the unit through its Manage IP address (or addresses).

Example: Setting Manage IPs for Multiple Interfaces

In this example, a small group of local administrators in the DMZ use the DMZ interface exclusively for HTTP, SNMP, and Telnet traffic. The Untrusted interface must also be able to support administrative traffic from a remote administrator using NetScreen-Global Manager. Manage IP addresses are set for both the DMZ and Untrusted interfaces to allow administrative access from both of those directions.

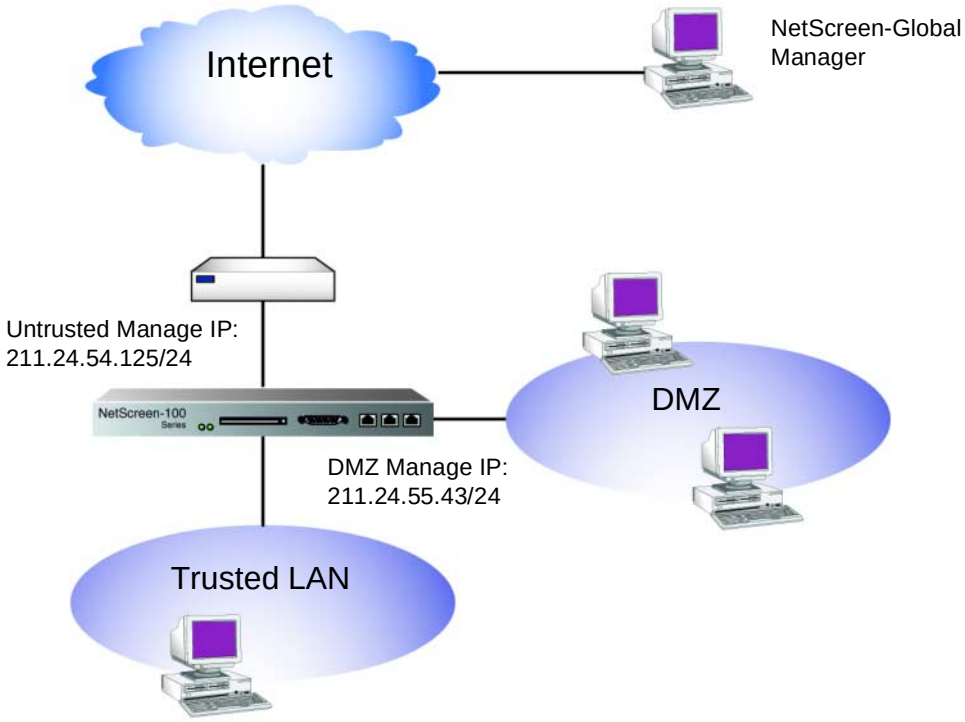

WebUI

1. Interface >> DMZ >> Edit: Enter the following, and then click **Save**:

   IP Address: 211.24.55.144
   Netmask: 255.255.255.0
   Default Gateway: 0.0.0.0
   Manage IP: 211.24.55.43
   Traffic Bandwidth: 0
   WebUI: (select)
   Telnet: (select)
   SNMP: (select)

2. Interface >> Untrusted >> Edit: Enter the following, and then click **Save and Reset**:

   IP Address: 211.24.54.10
   Netmask: 255.255.255.0
   Default Gateway: 0.0.0.0
   Manage IP: 211.24.54.125
   Traffic Bandwidth: 0
   NS-Global: (select)
   NS-GlobalPRO: (select)

CLI

1. set interface dmz ip 211.24.55.144 255.255.255.0
2. set interface dmz manage-ip 211.24.55.43
3. set interface dmz manage web
4. set interface dmz manage telnet
5. set interface dmz manage web
6. set interface untrust ip 211.24.54.10 255.255.255.0
7. set interface untrust manage-ip 211.24.54.125
8. set interface untrust manage global
9. set interface untrust manage global-pro
10. save

Management Interface

The Management (MGT) interface allows you to manage the NetScreen-1000 through a separate interface, moving administrative traffic outside the regular network user traffic. Separating administrative traffic from network user traffic greatly increases administrative security and assures constant management bandwidth.

On the NetScreen-1000, the Management (MGT) interface provides a dedicated connection for management traffic. Connect one end of a Cat-5 serial cable to the MGT interface and the other end to your management network or workstation.

With this arrangement, you can use a Web browser to manage through the WebUI (see "Web User Interface" on page 3-2) or use Telnet (see "Telnet" on page 3-6) to manage through the CLI. You can also manage through the MGT interface by connecting a workstation directly to the console port or modem port and accessing the device through its MGT IP address.

Example: Administration Through the MGT Interface

You can configure the NetScreen-1000 to allow administration through one or more of the Trusted, Untrusted, or Management (MGT) interfaces. To maintain the highest level of security, NetScreen recommends that you limit administrative traffic exclusively to the MGT interface and user traffic to the Trusted and Untrusted interfaces. This prohibits administrative access from Trusted and Untrusted workstations that are connected to your network and assures bandwidth availability for administrative traffic.

In this example, the IP address of the MGT interface is 192.168.20.2/24, and the MGT interface is enabled to receive Telnet and Web administrative traffic.

WebUI

Interface >> Management >> Edit: Enter the following, and then click **Save and Reset**:

IP Address: 192.168.20.2

Netmask: 255.255.255.0

Default Gateway: 0.0.0.0.

Traffic Bandwidth: 0

Enable Manageability: WebUI (select), Telnet (select)

CLI

1. set interface mgt ip 192.168.20.2 255.255.255.0
2. set interface mgt manage web
3. set interface mgt manage telnet
4. save

Virtual Private Networks

You can use a Virtual Private Network (VPN) to secure remote management and monitoring of a NetScreen device from either a dynamically assigned or fixed Untrusted IP address. Using a VPN, you can protect any kind of traffic, such as HTTP, Telnet, or SNMP.

NetScreen supports three methods for creating a 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 Preshared Key:* One or two preshared 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 "Virtual Private Networks" on page 6-1. For more information on NetScreen-Remote, refer to the* NetScreen-Remote User's Guide.

By default, NetScreen VPN tunnels use the Untrusted interface IP address (in NAT mode) or the System IP address (in Transparent mode) as the tunnel endpoint. Optionally, you can designate the Trusted interface as the endpoint when directing management traffic through a VPN tunnel to an address on the

Untrusted side. This allows you to create an Access Policy encrypting management traffic, such as SNMP or syslog, originating within the NetScreen device (with the source address being the Trusted interface) and destined for a remote server on the Untrusted side. To enable this, do the following:

WebUI

Select one or more of the following check boxes, and then click **OK**:

Admin >> Syslog: Enable Syslog VPN encryption: (select)

Admin >> Syslog: Enable WebTrends VPN encryption: (select)

Admin >> SNMP: Enable SNMP VPN encryption: (select)

Admin >> NS Global: Enable Global Manager/PRO VPN encryption: (select)

CLI

set {global | snmp | syslog | webtrends} vpn

Note: *You also need to define the VPN tunnel and create an Access Policy.*

Example: Administration through a VPN Tunnel on the Trusted Side

In this example, the network security administrator uses a VPN to keep security separate from general network administration. She creates a Manual Key VPN tunnel from her workstation at 10.10.11.56/24 to 10.10.10.1/24, the IP address of the Trusted interface. She has NetScreen-Remote 5.0 installed on her workstation.

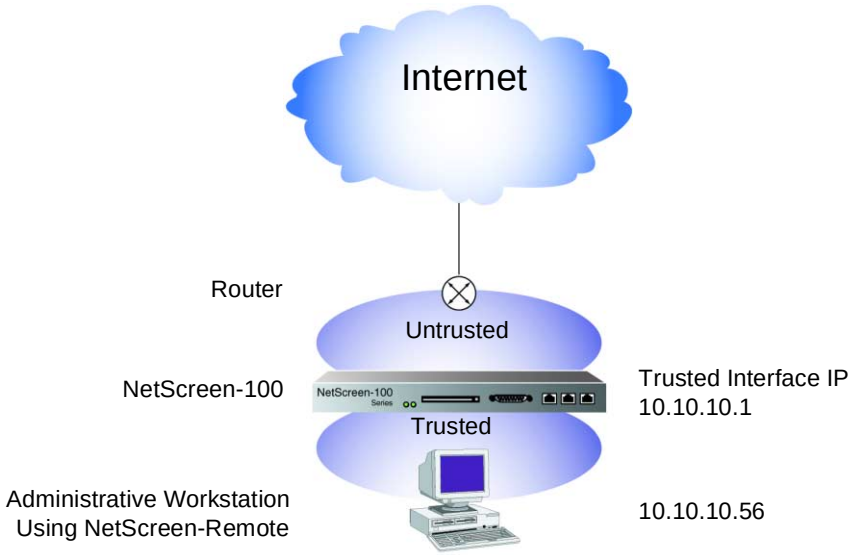

WebUI

NetScreen-100

1. Address >> Trusted: New Address: Enter the following, and then click **OK**:

   Name: Trusted Interface

   Address: 10.10.10.1

   Netmask: 255.255.255.255

2. Address >> Untrusted: New Address: Enter the following, and then click **OK**:

   Name: Admin 1

   Address: 10.10.10.56

   Netmask: 255.255.255.255

   Comment: For VPN Admin

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

   Name: Admin Tunnel
   Gateway IP: 10.10.10.56
   Security Index: 4567 (Local) 5555 (Remote)
   ESP-CBC: (select)
   Encryption Algorithm: DES-CBC
   Generate Key by Password[9]: netscreen1
   Authentication Algorithm: MD5
   Generate Key by Password: netscreen2
   Tunnel to Trusted Interface: (select)

Note: *By default, a VPN tunnel to a NetScreen device terminates at the Untrusted interface. After you select the* ***Tunnel to Trusted Interface*** *option, you cannot clear it. To modify the tunnel to terminate at the Untrusted interface, you must first remove the existing tunnel, and then create a new one.*

If the NetScreen device is in Transparent mode, then the tunnel from the Trusted side terminates at the system IP address.

4. Policy >> Outgoing >> New Policy: Enter the following, and then click **OK**:

   Name: Admin VPN Policy
   Source Address: Trusted Interface
   Destination Address: Admin 1
   Service: Any
   Action: Tunnel
   VPN Tunnel: Admin VPN

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

CLI

1. set address trust “Trusted Interface” 10.10.10.1 255.255.255.255
2. set address untrust “Admin 1” 10.10.11.56 255.255.255.255
3. set vpn trust manual 4567 5555 “Admin tunnel” gateway 10.10.10.56 esp des password netscreen1 auth md5 password netscreen2
4. set policy outgoing “Trusted Interface” “Admin 1” any tunnel vpn “Admin tunnel”

NetScreen-Remote Security Policy Editor

1. Click **Options >> Secure >> Specified Connections**.
2. Click the **Add a new connection** button, and type **ns100** next to the new connection icon that appears.
3. Configure the connection options:

   Connection Security: Secure

   Remote Party ID Type: IP Address

   IP Address: 10.10.10.1
4. Click the **PLUS** symbol, located to the left of the new connection icon, to expand the connection policy.
5. ns100 >> Security Policy: Use Manual Keys: (select)
6. 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.
7. Key Exchange (Phase 2) >> Proposal 1: Select the following IPSec Protocols:

   Encapsulation Protocol (ESP): (select)

   Encrypt Alg: DES

   Hash Alg: MD5

   Encapsulation: Tunnel
8. Proposal 1 >> Inbound Keys: In the Security Parameters Index field, type 5555, and then click **Enter Key**.

9. Inbound Keys >> Enter Key: Enter the following[10], and then click **OK**:

   Choose key format: Binary

   ESP Encryption Key: dccbee96c7e546bc

   ESP Authentication Key: dccbe9e6c7e546bcb0b667794ab7290c

10. Proposal 1 >> Outbound Keys: In the Security Parameters Index field, type 4567, and then click **Enter Key**.
11. Outbound Keys >> Enter Key: Enter the following, and then click **OK**:

    Choose key format: Binary

    ESP Encryption Key: dccbee96c7e546bc

12. ESP Authentication Key: dccbe9e6c7e546bcb0b667794ab7290c
13. Click the **Save** button.

10. These are the two generated keys that you copied after configuring the NetScreen device.

Building Blocks for Access Policies and VPNs

4

This chapter discusses the concepts common to Access Policies and Virtual Private Networks (VPNs). The specific topics discussed are:

ADDRESSES

The NetScreen ScreenOS classifies the addresses of all other devices by location and netmask. Trusted addresses are located behind the Trusted interface and appear as green in the WebUI. Untrusted addresses are located behind the Untrusted interface and appear as red. DMZ addresses are located behind the DMZ interface (NetScreen-10 and -100) and appear as brown.

Individual hosts have only a single IP address defined and are represented by a single computer icon in the WebUI. Individual hosts must have a netmask setting of 255.255.255.255 (which masks out all but this host).

Subnets have an IP address and a netmask (for example, 255.255.255.0 or 255.255.0.0) and are represented by multiple computer icons in the WebUI.

Before you can configure Access Policies to permit, deny, or tunnel traffic to and from individual hosts and subnets, you must make entries for them in the Trusted, Untrusted, and DMZ (NetScreen-10 and -100) sections of the NetScreen address book.

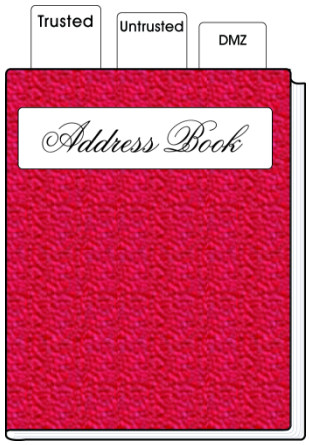

Note: *You do not have to make address book entries for "Inside Any," "Outside Any," or "DMZ Any." These terms automatically apply to all devices physically located beyond these respective interfaces.*

Address Book Entries

Before you can set up many of the NetScreen firewall, VPN, and traffic shaping features, you need to define addresses in the address book. The address book contains the IP addresses or domain names[1] of hosts or subnets whose traffic is either allowed, blocked, encrypted, or user-authenticated.

1. Before you can use domain names for address book entries, you must configure the NetScreen device for Domain Name System (DNS) services. For information on DNS configuration, see "Domain Name System Support" on page 2-17.

Example: Adding Addresses

In this example, you add the subnet “Santa Clara Eng” with the IP address 192.10.10.0/24 as a Trusted address, and the address www.firenet.com as an Untrusted address.

WebUI

1. Address >> Trusted >> New Address: Enter the following information, and then click **OK**:

   Address Name: Santa Clara Eng
   IP Address/Domain Name: 192.10.10.0
   Netmask: 255.255.255.0
   Trust: (select)

2. Address >> Untrusted >> New Address: Enter the following information, and then click **OK**:

   Address Name: FireNet
   IP Address/Domain Name: www.firenet.com
   Netmask: 255.255.255.255
   Untrust: (select)

CLI

1. set address trust “Santa Clara Eng” 192.10.10.0 255.255.255.0
2. set address untrust www.firenet.com 255.255.255.255
3. save

Example: Modifying Addresses

In this example, you change the address entry for the host “Santa Clara Eng” to reflect that this host has moved to Dallas and reassigned an IP address of 192.10.40.0/24.

WebUI

Address >> Trusted >> Edit (for Santa Clara Eng): Change the name and IP address to the following, and then click **OK**:

Address Name: Dallas Eng

IP Address/Domain Name: 192.10.40.0

CLI

1. unset address trust “Santa Clara Eng”
2. set address trust “Dallas Eng” 192.10.40.0 255.255.255.0
3. save

Note: *After you define an address—or an address group—and associate it with an Access Policy, you cannot change the address location to another interface (such as from Trusted to Untrusted). To change its location, you must first disassociate it from the underlying Access Policy.*

Example: Deleting Addresses

In this example, you remove the address entry for the subnet “Dallas Eng.”

WebUI

Address >> Trusted: Click **Remove** in the Configure column for Dallas Eng.

CLI

1. unset address trust “Dallas Eng”
2. save

Address Groups

The previous section explained how you create, modify, and delete address book entries for individual hosts and subnets. As you add addresses to the address book, it becomes difficult to manage how Access Policies affect each address entry. NetScreen allows you to create groups of addresses. Rather than manage a large number of address book entries, you can manage a small number of groups. Changes you make to the group are applied to each address entry in the group.

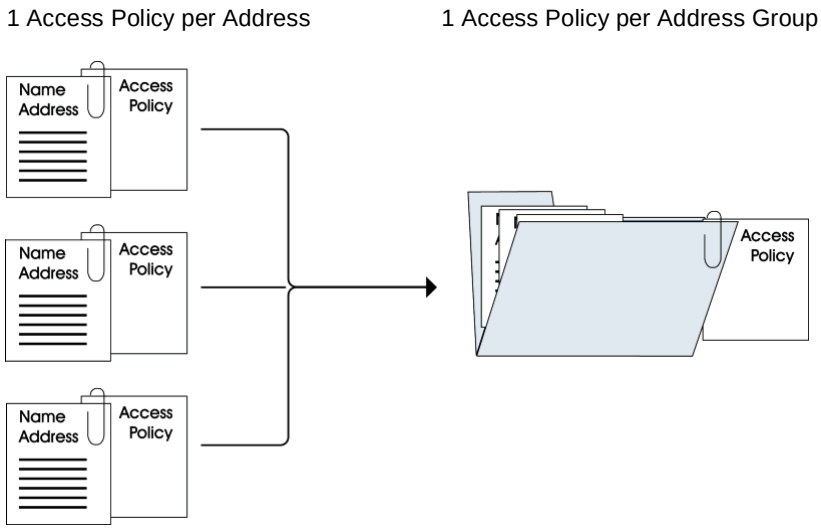

The Address Group option has the following features:

- You can create address groups on the Trusted, Untrusted, or DMZ sides.
- You can create address groups with existing users, or you can create empty address groups and then fill them with users.
- Address group entries can be used like individual address book entries.
- NetScreen applies Access Policies to each member of the group by creating individual Access Policies for each group member. While you only have to create one Access Policy for a group, NetScreen actually creates an Access Policy for each member in the group (as well as for each service configured for each user).[2]
- When an individual address book entry is deleted from the address book, it is also removed from all groups in which it was referenced.

The following constraints apply:

- Address groups can only contain addresses for one type of interface (Trusted, Untrusted, or DMZ).
- Address names cannot be the same as group names. If the name “Paris” is used for an individual address entry, it cannot be used for a group name.
- If an address group is referenced in an Access Policy, the group cannot be removed. It can, however, be edited.
- When a single Access Policy is assigned to an address group, it is applied to each group member individually, and the NetScreen device makes an entry for each member in the access control list (ACL). If you are not vigilant, it is possible to exceed the number of available Access Policy resources, especially if both the source and destination are address groups
- You cannot add the predefined addresses: “Outside Any,” “Inside Any,” “DMZ Any,” “All Virtual IPs,” and “Dial-Up VPN” to groups.
- The following table lists the group size limits for each platform.

Hardware Platform	Number of Groups	Members per Group
NetScreen-5	16	16
NetScreen-10	32	32
NetScreen-100	64	64
NetScreen-1000	256 (Root); 8 (Virtual System)	256 (Root); 16 (Virtual System)

2. The automatic nature by which NetScreen applies Access Policies to address group members, saves you from having to create them one by one for each address. Furthermore, NetScreen writes these Access Polices to ASIC which makes lookups run very fast.

Example: Creating an Address Group

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

WebUI

Address >> Trusted >> New Group: Enter the following group name, move the following addresses, and then click **OK**:

Group Name: HQ 2nd Floor

Group Members << Available Members:

Santa Clara Eng

Tech Pubs

CLI

1. set group address trust “HQ 2nd Floor” add “Santa Clara Eng”
2. set group address trust “HQ 2nd Floor” add “Tech Pubs”
3. save

Example: Editing a Group Address Entry

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

WebUI

Address >> Trusted >> Edit (for HQ 2nd Floor): Move the following address, and then click **OK**:

Group Members << Available Members:

Support

CLI

1. set group address trust “HQ 2nd Floor” add Support
2. save

Example: Removing an Address Group Member and a Group

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

WebUI

1. Address >> Trusted >> HQ 2nd Floor >> Edit: Move the following address, and then click **OK**:

   Group Members >> Available Members:
   Sales

2. Address >> Trusted: Click **Remove** in the Configure column for Sales.

CLI

1. unset group address trust “HQ 2nd Floor” remove Support
2. unset group address trust Sales
3. save

Note: *The NetScreen device does not automatically delete a group from which you have removed all names.*

VIRTUAL IP

The Virtual IP (VIP) feature provides network flexibility and security. In a Network Address Translation (NAT) environment, host computers use non-routable IP addresses inside the firewall while maintaining full Internet connection and functionality. This feature gives network administrators flexibility to expand their networks without being constrained by the scarcity of legal IP addresses. In addition, NAT also provides better network security by hiding internal network topology and host information from the outside world.

However, to maintain some Internet services (for example, e-mail, POP3, FTP), a server with a legal IP address must be present to service the requests. VIP allows you to map routable IP addresses to internal servers, thereby providing transparent connections for a NAT network to the Internet. Other benefits of using VIP include:

Scalability: As Internet service demand increases, companies need to improve servers' performance in order to maintain the quality of their services. While upgrading the server to a larger, faster machine generally relieves the short-term pressures, the disruption to services and the prohibitive cost of upgrading quickly make this solution undesirable.

Redundancy: With Virtual IP, servers can be assigned to the same IP address and mirrored to provide High Availability (HA) for network services. Individual servers can also be taken off-line for maintenance without disruption to network traffic.

Reduction in capital cost: Multiple domains and Web servers can be mapped to the same physical server, thus reducing the cost of computer equipment as well as the associated administration tasks.

Flexibility in assigning ports: By setting up Virtual IP (VIP) addresses, you can configure your NetScreen device to route traffic destined for many different IP addresses on the subnet of the Untrusted interface to specific addresses on the Trusted network.

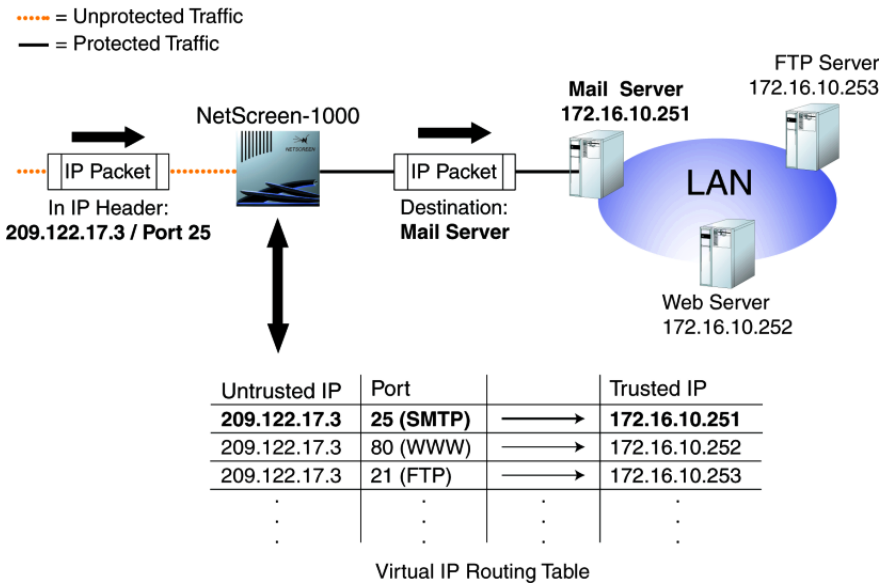

The maximum number of VIPs, and the maximum number of services per VIP that are supported by each NetScreen device are as follows:

	VIPs	Services/VIP
NetScreen-5	1	64
NetScreen-10	2	64
NetScreen-100	4	8
NetScreen-1000	6	8

Required Information

You need the following information to define a Virtual IP:

- The IP address for the VIP, which must be in the same subnet as the Untrusted interface and can even be the same address as the Untrusted interface[3]
- The port number for communication between the Trusted server and the Untrusted interface on the NetScreen device
- The IP address for the server on the Trusted interface that processes the requests

Example: Configuring Virtual IP Servers

In this example, you configure a VIP at 2.2.2.20 to route inbound HTTP traffic to a pool of two Web servers at 172.16.12.10 and 172.16.12.11. (The Untrusted IP address of the NetScreen-100 is 2.2.2.10/24.) The port number for HTTP is translated from 80 (the standard protocol ID number) to 1142.

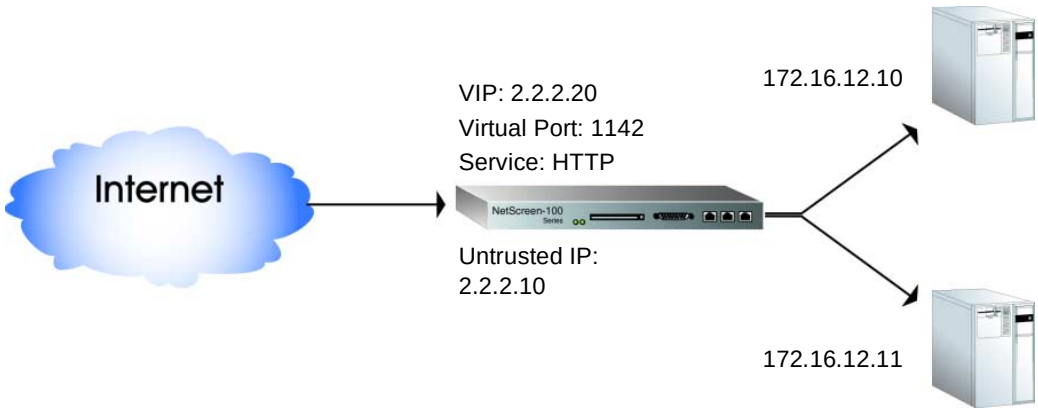

3. On the NetScreen-5, the Untrusted interface can receive its Untrusted IP address dynamically via DHCP or PPPoE. If you want to use a VIP in such a situation, do either of the following: In the WebUI (Virtual IP >> Virtual Server IP >> Click here to configure), select the **Same as the Untrusted IP address** option when setting up the VIP; in the CLI, use the **set vip untrust-ip** command.

WebUI

1. Virtual IP >> Virtual IP1 >> Virtual Server IP: Enter the following address, and then click **OK**:

   Virtual IP Address: 2.2.2.20

2. New Service: Enter the following, and then click **OK**:

   Virtual Port: 1142[4]

   Service: HTTP[5]

   (NetScreen-100) Load Balance: None

   1 Server IP: 172.16.12.10

   (NetScreen-100) Server Weight: 1

   2 Server IP: 172.16.12.11

   (NetScreen-100) Server Weight: 1

CLI

1. set vip 2.2.2.20 1142 http none 172.16.12.10/1
2. set vip 2.2.2.20 1142 http none 172.16.12.11/1
3. save

4. Using non-standard port numbers adds another layer of security, thwarting attacks that check for services at standard port numbers.

5. When initially configuring a VIP, you can only map one service at a time. For example, if you are mapping six services to a Virtual IP, you must enter each one individually.

Example: Editing a VIP Configuration

In this example, you modify the Virtual IP server configuration you just created. In this case, you add an additional server 172.16.12.12.

WebUI

Virtual IP >> Virtual IP1 >> Edit (in the HTTP row): Enter the following, and then click **OK**:

3 Server IP: 172.16.12.12
(NetScreen-100) Server Weight: 1

CLI

1. set vip 2.2.2.20 1142 http none 172.16.12.12/1
2. save

Example: Removing a VIP Configuration

In this example, you delete the VIP configuration that you just created and modified.

WebUI

Virtual IP >> Virtual IP1 >> Virtual Server IP 2.2.2.20: Click **Clear**.

CLI

1. unset vip 2.2.2.20
2. save

Note: *You cannot edit or remove a Virtual IP entry when existing Access Policies are still associated with it.*

MAPPED IP

Mapped IP (MIP) is a direct one-to-one mapping of traffic destined for one IP address to another IP address, and is based solely on IP addresses. When the NetScreen device is operating in NAT mode, an MIP provides a means for incoming traffic to reach a private address on the Trusted network. You can configure an MIP address to route traffic destined for an address on the Untrusted subnet to a different address on the Trusted subnet, regardless of the service and corresponding port number involved.

By setting up MIP addresses, you can configure the NetScreen device to route traffic destined for many different IP addresses on the subnet of the Untrusted interface to specific addresses on the Trusted network.

Mapped IP: Inbound traffic is mapped from 200.2.2.10 on the Untrusted subnet to 172.16.40 on the Trusted subnet.

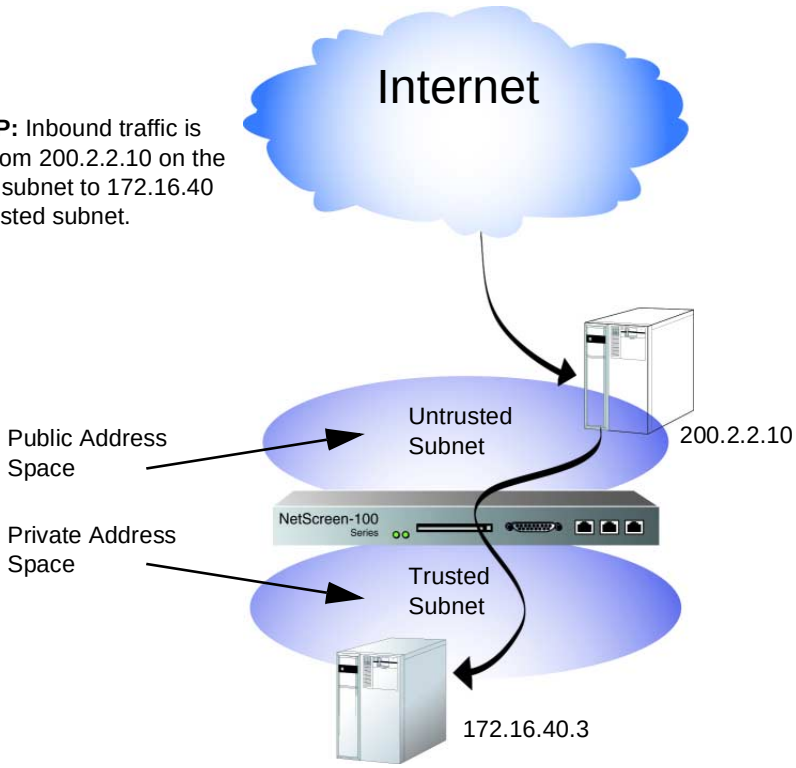

Example: Creating a Mapped IP Address

This example explains how to map incoming traffic destined to the Untrusted IP address 209.122.17.6 to the Trusted IP address 172.16.17.6.

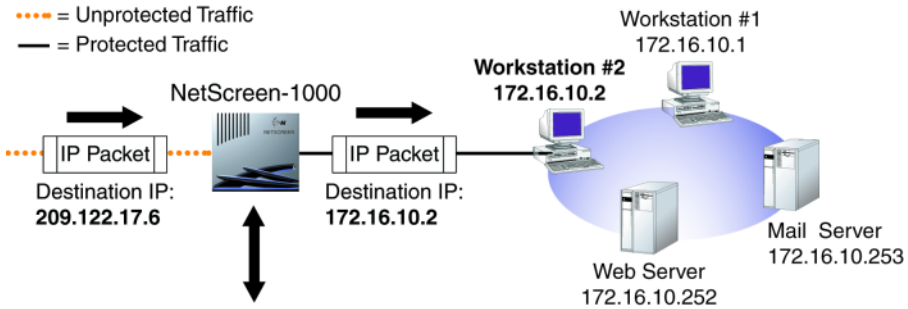

Untrusted IP		Trusted IP
209.122.17.5	→	172.16.10.1
209.122.17.6	→	**172.16.10.2**
.	.	.
.	.	.
.	.	.
209.122.17.249	→	172.16.10.252
209.122.17.250	→	172.16.10.253

Mapped IP Routing Table

WebUI

Virtual IP >> Mapped IP >> New Entry: Enter the following and then click **OK**:

Untrusted IP Address: 209.122.17.6

Netmask: 255.255.255.255

Map to IP Address: 172.16.10.2

CLI

1. set mip 209.122.17.6 host 172.16.10.2 255.255.255.255
2. save

> ***Note:*** *You must define an Access Policy allowing the mapped IP address to be accessed. No address book entry is required for a Mapped IP.*
>
> *You can map an address-to-address or subnet-to-subnet relationship. When a subnet-to-subnet mapped IP configuration is defined, the netmask is applied to both the mapped IP subnet and the original IP subnet.*

USERS

NetScreen supports three kinds of users:

- Authentication User – A network user who must provide a user name and password for authentication when initiating a connection across the firewall.
- IKE Dynamic Peer – A VPN user with a dynamically assigned IP address. The user provides his or her identity using an e-mail address, an IP address, or a domain name. The VPN can use either AutoKey IKE with a preshared key or AutoKey IKE with a certificate.
- VPN Dialup User – A VPN user with a dynamically assigned IP address. The VPN uses the manual key method for encryption and/or authentication.

Before traffic from an authentication user can traverse the firewall, and before a VPN user can participate in a VPN, you must create a configuration profile for each one.

> ***Note:*** *For more information about creating VPNs, see Chapter 4, "Virtual Private Networks".*

Example: Creating Three New Users

In this example, you create the following users:

- An authentication user named Alice with the password “Nd4syt4.”
- An IKE dynamic peer named Carol with the ID carol@netscreen.com
- A VPN dialup user named Phil, who is assigned to the dialup user group “Western” and uses 3DES encryption with SHA-1 authentication.

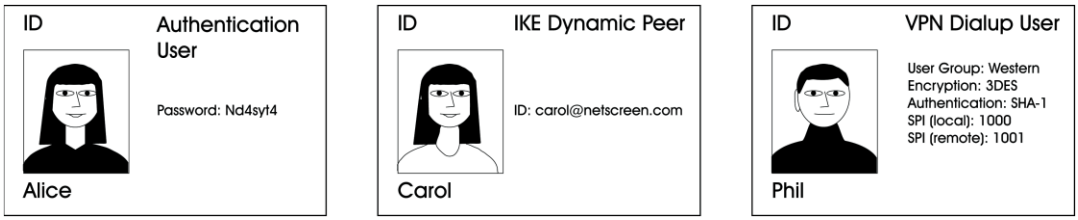

WebUI

1. Users >> Users >> New User: Enter the following, and then click **OK**:
   User Name: Alice
   Authentication User: (select)
   Authentication Password: Nd4syt4
   Confirm Password: Nd4syt4
   Status: Enable (select)
2. Users >> Users >> New User: Enter the following, and then click **OK**:
   User Name: Carol
   IKE Dynamic Peer: (select)
   User Group: None
   Identity: carol@netscreen.com
3. Users >> Users >> New User: Enter the following, and then click **OK**:
   User Name: Phil
   VPN Dialup User - Manual Key Only: (select)
   User Group: Western
   Security Index: 1000 (Local); 1001 (Remote)
   ESP: (select)
   ESP-Encryption-Algorithm: 3DES CBC

Generate Key by Password: 12345678
Authentication Algorithm: SHA-1
Generate Key by Password: 99999999

CLI

1. set user Alice password Nd4syt4
2. set user Carol ike-id carol@netscreen.com
3. set user Carol enable
4. set user Phil dialup 1000 1001 esp 3des pass 12345678 auth sha-1 pass 99999999
5. save

User Authentication

There are a number of different protocols that your NetScreen device can use to verify that a user is who they say they are. These different techniques are discussed in this section.

Internal Database

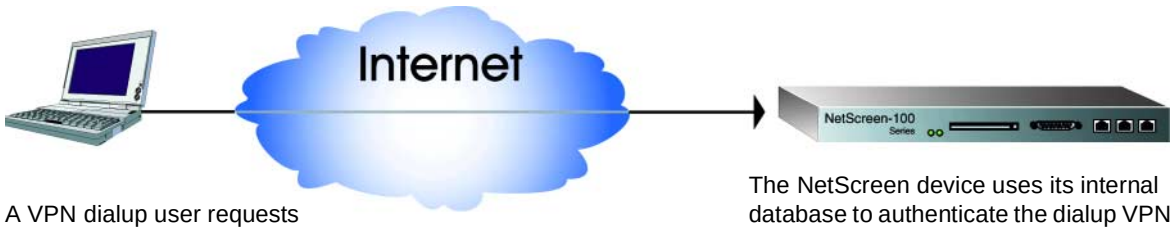

A VPN dialup user requests a connection to the protected network.

The NetScreen device uses its internal database to authenticate the dialup VPN user before approving the connection request.

All NetScreen devices support a built-in user database for authentication. The maximum number of entries supported by each device are:

Platform	Total Number of Entries
NetScreen-5	100
NetScreen-10	500
NetScreen-100	1500
NetScreen-1000	2000

After entering the user name and password in the database, you must create an Access Policy that requires a user to authenticate him or herself when initiating a specified connection (for example, outbound or inbound HTTP, or Telnet traffic). When the user attempts to initiate traffic for which the Access Policy applies, he or she is prompted to enter his or her name and password. Before granting permission, the NetScreen device validates the user name and password by checking them against those stored in the database.

RADIUS

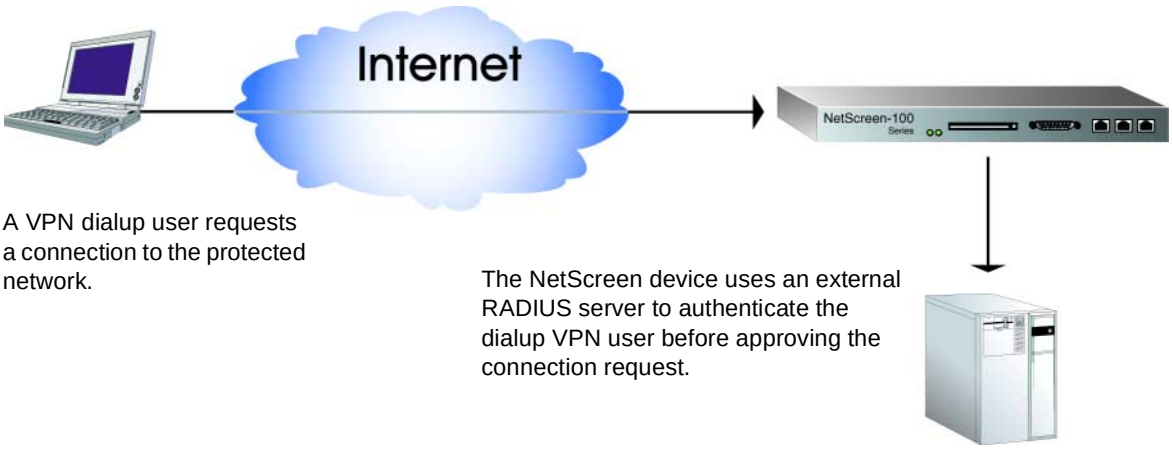

The Remote Authentication Dial-In User Service (RADIUS) is a protocol for an authentication server which can be modified to run on different kinds of networks, and makes it easy and efficient to manage large modem pools. The focus for RADIUS is the remote user who needs to dial into the network.

RADIUS uses an authentication server to solve the security problems associated with remote computing. Distributed security separates user authentication and authorization from the communications process and creates a single, central location for user authentication data.

One RADIUS server can support up to tens of thousands of users, making it a very practical service for rapidly growing networks.

The RADIUS client (that is, the NetScreen device) authenticates users through a series of communications between the client and the server. Basically, RADIUS asks the person logging on to enter his or her user name and password. It then compares these values to those in its database, and once a user is authenticated, the client provides the user with access to the appropriate network services.

SecurID

The relationship of NetScreen device and a Security Dynamics Technologies® SecurID® ACE™ server is similar to that of a NetScreen device and a RADIUS server; that is, the NetScreen device acts as a client, forwarding authentication requests to the external server for approval. SecurID differs from RADIUS in that the user password involves a continually changing string of numbers.

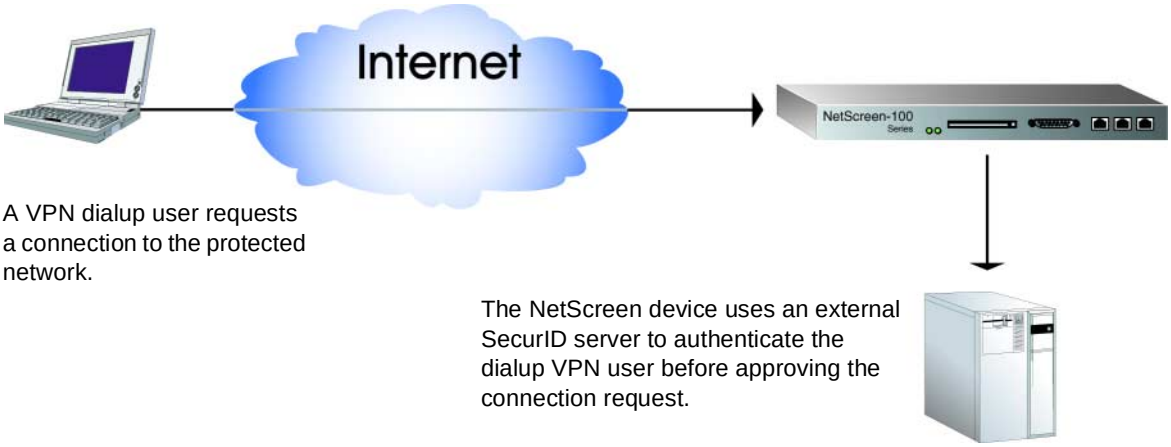

SecurID issues a credit card sized device with an LCD window that displays a randomly generated string of numbers that changes every minute. There is no other information on the card besides the number in the LCD display.

Security Dynamics issues a card and a personal ID number (PIN) to a registered user and maintains the user profile in their database. When the user is prompted to authenticate himself, he enters his name and password, which is his PIN followed by the string of numbers currently displayed on his card. The numbers displayed on the card change every minute. The values that display are generated by an algorithm known only by Security Dynamics. This value is saved to the

Security Dynamics database entry for this PIN. When the user to be authenticated enters his PIN and the number on his card, Security Dynamics compares these values to those in the database. If they match, the user is authenticated.

Lightweight Directory Access Protocol

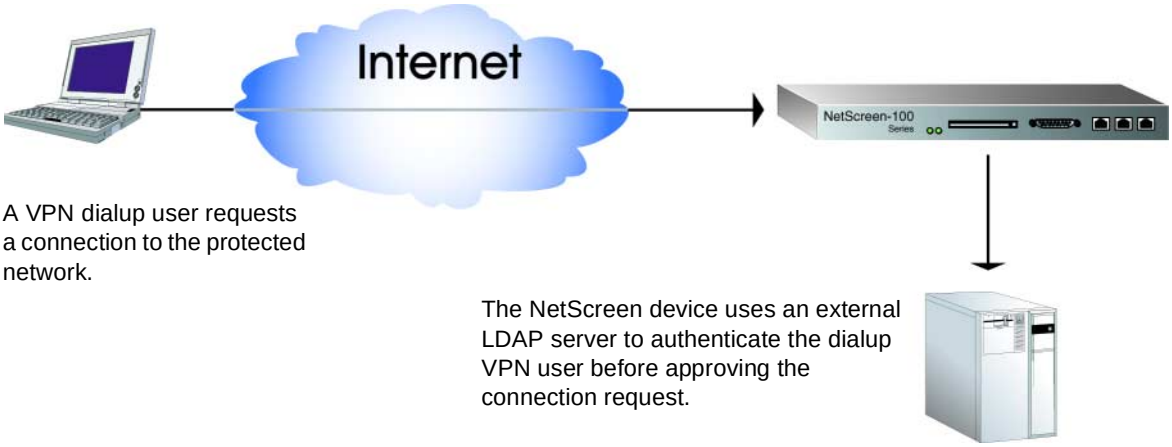

Lightweight Directory Access Protocol (LDAP) is a directory server standard developed by Netscape® to help authenticate users attempting to connect to networks controlled by directory servers.

LDAP is a client-server protocol for accessing a directory service. You can use it as a front-end to X.500, as a stand-alone protocol, or as a directory server.

LDAP does not require the upper layers of the OSI stack, it is a simpler protocol to implement (especially in clients), and LDAP is under IETF change control and so can more easily evolve to meet Internet requirements.

The LDAP information model is based on the entry, which contains information about some object (for example, a person). Entries are composed of attributes, which have a type and one or move values. Each attribute has a syntax that determines what kind of values are allowed in the attribute and how those values behave during directory operations.

Dialup User Groups

One of the main reasons organizations use VPNs is to allow remote dialup users to be able to traverse the firewall from anywhere in the world and access their data in a secure environment. The VPN tunnel connection from them to the corporate site assures security as well as access.

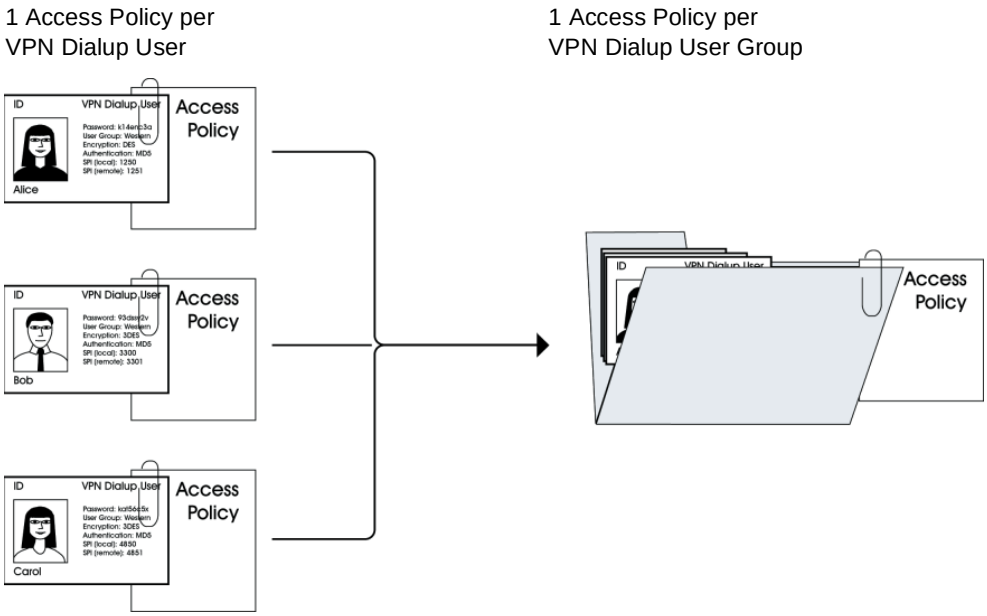

To manage a number of remote dialup users, NetScreen enables you to create dialup user groups. Rather than manage each user individually, you can aggregate users into a group. Changes you make to the group are then propagated to each group member. The examples that follow illustrate how to create new dialup user groups and then add users to it. Other examples show how to remove members from a group and move members from one group to another.

Example: Defining a New Dialup User Group

In this example, you define a new dialup user group named Tahoe.

WebUI

Users >> Dialup Group >> New Group: Enter the following, and then click **OK**:

Dialup Group Name: Tahoe

CLI

1. set dialup-group tahoe
2. save

Example: Adding a Member to a Dialup User Group

In this example, you add a user named Fred to the dialup user group Tahoe.

WebUI

Users >> Users >> Edit (for the user named Fred): Select the following, and then click **OK**:

VPN Dialup User: (select)

User Group: Tahoe

CLI

1. set dialup-group tahoe + Fred
2. save

Example: Removing an Existing Group Member

In this example, you delete Phil from the Tahoe dialup user group.

WebUI

Users >> Users: Click **Remove** (for Phil).

CLI

1. set dialup-group Tahoe - Phil
2. save

Example: Moving a Group Member to Another Group

In this example, you move Phil from the dialup user group Tahoe to the group Santa Cruz.

WebUI

Users >> Users >> Edit (for Phil): Enter the following, and then click **OK**:

VPN Dialup User: Select

User Group: Santa Cruz

CLI

1. set dialup-group Tahoe - Phil
2. set dialup-group “Santa Cruz” + Phil
3. save

SERVICES

Services are types of IP traffic for which protocol standards exist. Each service has a port number associated with it, such as 21 for FTP and 23 for Telnet.

The illustration below shows the services supported in ScreenOS 2.5. For information on each service, hold your cursor over the service icon. In this illustration, the mouseover information block is displayed for X-Windows.

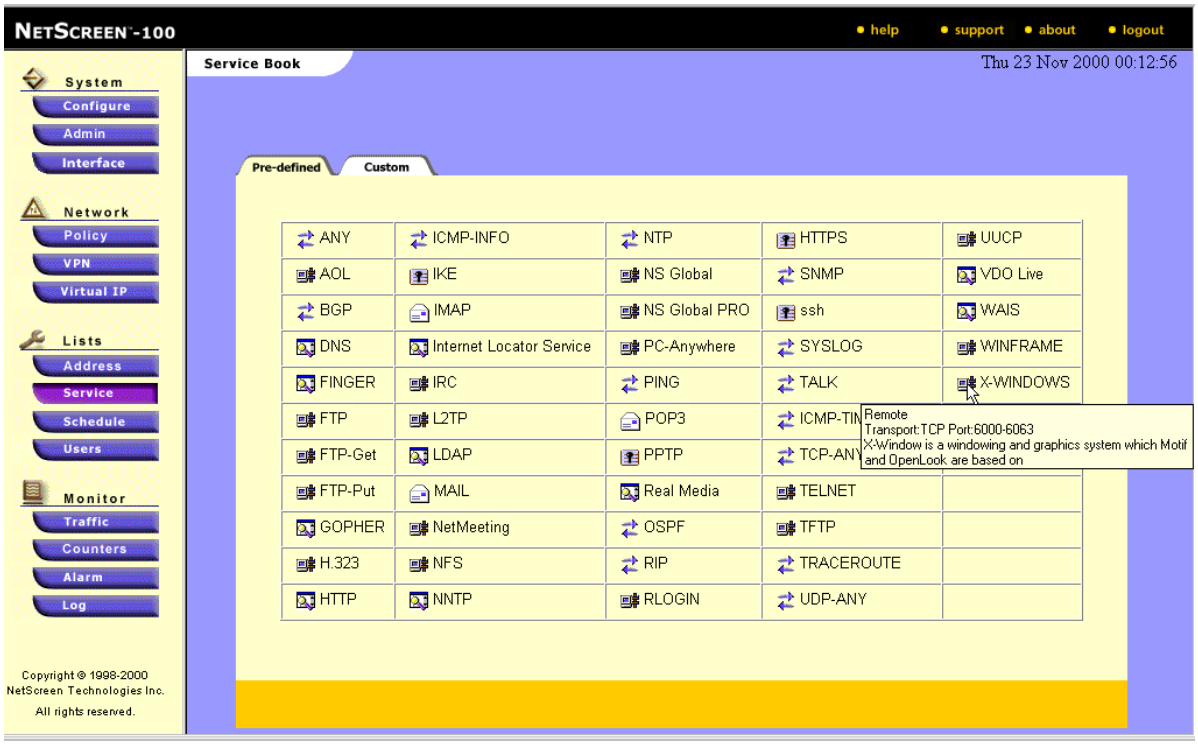

When you create an Access Policy, you must specify a service for it. You can select one of the pre-configured services from the service book, or a custom service or service group that you created. You can see which service you can use in an Access Policy by examining the Service drop-down List in the Policy Configuration dialog box (WebUI), or by using the **get service** command (CLI).

The following section provides examples for viewing the service book and for creating, modifying, and deleting custom services.

Example: Viewing the Service Book

In this example, you view the predefined and custom services in the service book.

WebUI

1. Service >> Pre-defined
2. Service >> Custom

CLI

get service

The output from the CLI is similar to that shown below.

```
transport   src port(low/high)   dst port(low/high)   ack/reverse
Name:ANY(Group:other/5, Id#0, Flag:128)
 0           0/65535      0/65535               0/0
Name:AOL(Group:remote/1, Id#1, Flag:128)
 tcp         1/65535      5190/5194             0/0
Name:DNS(Group:info seeking/3, Id#2, Flag:128)
 udp         0/65535      53/53                 0/0
 tcp         1024/65535   53/53                 0/0
```

Example: Adding a Custom Service

To add a custom service to the service book, you need the following information:

- A name for the service, in this example “Corporate”
- A range of source port numbers valid for the service. For example, 1500-10000.
- A range of destination port numbers to receive the service request; for example, 15000-25000.
- Whether the service uses TCP or UDP protocol, or some other protocol as defined by the Internet specifications. In this example, the protocol is TCP

WebUI

Service >> Custom >> New Service: Enter the following, and then click **OK**:

Service Name: Corporate
Source Port Low: 1500
Source Port High: 10000
Destination Port Low: 15000
Destination Port High: 25000
Transport: TCP

CLI

1. set service Corporate + protocol tcp src-port 1500-10000 dst-port 15000-25000
2. set service Corporate + timeout 30
3. save

Example: Modifying a Custom Service

In this example, you change a custom service. In this case, the Transport is UDP, and the Source Port range changes to 1 through 1000.

Use the **set service <name> clear** command to remove the definition of a custom service without removing the service from the service book:

WebUI

Service >> Custom: Enter the following and then click **OK**:

Service Name: Corporate
Source Port Low: 1
Source Port High: 1000
Destination Port Low: 15000
Destination Port High: 25000
Transport: UDP

CLI

1. set service Corporate clear
2. set service Corporate + protocol udp src-port 1-1000 dst-port 15000-25000
3. save

Example: Removing a Custom Entry

In this example, you remove the custom service “Corporate.”

WebUI

Service >> Custom: Click **Remove** in the Configure column for “Corporate.”

CLI

1. unset service Corporate
2. save

Service Groups

A service group is a set of services that you have gathered together under one name. After you create a group containing several services, you can then apply services at the group level to Access Policies, thus simplifying administration.

The NetScreen service group option has the following features:

- Each service book entry can be referenced by one or more service groups.
- Each service group can contain pre-defined and user-defined service book entries.
- Each service group can be referenced by other service groups, providing that a group referencing other groups does not include itself in the reference list.

Service groups are subject to the following limitations:

- Service groups cannot have the same names as services; therefore, if you have a service named “FTP,” you cannot have a service group named “FTP.”
- If a service group is referenced in an Access Policy, you can edit the group but you cannot remove it until you have first removed the reference to it in the Access Policy.
- If a custom service book entry is deleted from the service book, the entry is also removed from all the groups in which it was referenced.
- The all-inclusive service term “ANY” cannot be added to groups.
- The following table lists the number of service groups supported by platform.

Hardware Platform	Number of Groups	Number of Members
NetScreen-5	16	16
NetScreen-10	32	32
NetScreen-100	64	64
NetScreen-1000	256 (Root); 8 (Virtual System)	64 (Root); 16 (Virtual System)

Example: Creating a Service Group

This example illustrates how you create a custom service named Wiget that supports IKE, FTP, and LDAP services.

WebUI

Service >> Custom >> New Group: Enter the following, move the following services, and then click **OK**:

Group Name: Wiget

Group Members << Available Members:

IKE

FTP

LDAP

CLI

1. set group service Wiget
2. set group service Wiget add ike
3. set group service Wiget add ftp
4. set group service Wiget add ldap
5. save

Note: *If you attempt to add a service to a service group that does not exist, the NetScreen device creates the group. Also, ensure that groups referencing other groups do not include themselves in the reference list.*

Example: Modifying a Service Group

Although you cannot modify any of the pre-defined NetScreen services, you can modify existing user-defined custom services and service groups.

In this example, you change the existing user-defined services from IKE, FTP, and LDAP to HTTP, FINGER, IMAP, and H.323 protocols.

WebUI

Service >> Custom >> Edit (for Wiget): Move the following services, and then click **OK**:

Group Members >> Available Members:

IKE

FTP

LDAP

Group Members << Available Members:

HTTP

FINGER

IMAP

H.323

CLI

1. clear group service Wiget
2. set group service Wiget add http
3. set group service Wiget add finger
4. set group service Wiget add imap
5. set group service Wiget add h.323
6. save

Example: Deleting a Service

Although you cannot remove any of the pre-defined NetScreen services, you can remove existing user-defined custom services and service groups.

In this example, you delete HTTP from the service group Wiget.

WebUI

Service >> Custom >> Edit (for Wiget): Move the following service, and then click **OK**:

Group Members >> Available Members:
HTTP

CLI

1. unset group service Wiget remove http
2. get service Wiget
3. save

Example: Deleting a Service Group

In this example, you delete the service group Wiget.

WebUI

Service >> Custom: Click **Remove** (for Wiget).

CLI

1. unset group service Wiget
2. save

Note: *The NetScreen device does not automatically delete a group from which you have removed all members.*

SCHEDULES

A schedule is a configurable object that you can associate with one or more Access Policies to define when they are in effect. Through the application of schedules, you can control network traffic flow and enforce network security.

When you define a schedule, enter values for the following parameters:

Schedule Name: The name that appears in the Schedule drop-down list in the Policy Configuration dialog box. Choose a descriptive name to help you identify the schedule. The name must be unique and is limited to 19 characters.

Comment: Any additional information that you want to add.

Recurring: Enable this when you want the schedule to repeat on a weekly basis.

Start and End Times: You must configure both a start time and an end time. You can specify up to two time periods within the same day.

Once: Enable this when you want the schedule to start and end only once.

mm/dd/yyyy hh:mm: You must enter both start and stop dates and times.

Example: Recurring Schedule

In this example, there is a short-term employee named Tom who is using the company's Internet access for personal pursuits after work. You create a schedule for non-business hours that you can then associate with an Access Policy to deny outbound TCP/IP traffic from that worker's computer (10.10.4.5/24) outside of regular business hours.

WebUI

1. Schedule >> New Schedule: Enter the following, and then click **OK**:

   Schedule Name: After Hours

   Comment: For non-business hours

   Recurring: (select)

   Period 1:

Week Day	Start Time	End Time
Sunday	00:00	23:59
Monday	00:00	06:00
Tuesday	00:00	06:00
Wednesday	00:00	06:00
Thursday	00:00	06:00
Friday	00:00	06:00
Saturday	00:00	23:59

   Period 2:

Week Day	Start Time	End Time
Sunday	17:00	23:59
Monday	17:00	23:59
Tuesday	17:00	23:59
Wednesday	17:00	23:59
Thursday	17:00	23:59
Friday	17:00	23:59
Saturday	17:00	23:59

2. Address >> Trusted >> New Address: Enter the following, and then click **OK**:

   Address Name: Tom

   IP Address/Domain Name: 10.10.4.5

   Netmask: 255.255.255.255

   Comment: Temp

   Location: Trust

3. Policy >> Outgoing: New Policy: Enter the following, and then click **OK**:

   Name: No Net

   Source Address: Tom

   Destination Address: Outside Any

   Service: HTTP

   Action: Deny

   Schedule: After Hours

CLI

1. set schedule “after hours” recurrent sunday start 00:00 stop 23:59
2. set schedule “after hours” recurrent monday start 00:00 stop 06:00 start 17:00 stop 23:59
3. set schedule “after hours” recurrent tuesday start 00:00 stop 06:00 start 17:00 stop 23:59
4. set schedule “after hours” recurrent wednesday start 00:00 stop 06:00 start 17:00 stop 23:59
5. set schedule “after hours” recurrent thursday start 00:00 stop 06:00 start 17:00 stop 23:59
6. set schedule “after hours” recurrent friday start 00:00 stop 06:00 start 17:00 stop 23:59
7. set schedule “after hours” recurrent saturday start 00:00 stop 23:59 comment “for non-business hours”
8. set address trust tom 10.10.4.5 255.255.255.0 “temp”
9. set policy outgoing tom outside-any http deny schedule “after hours”
10. save

Access Policies 5

This chapter describes what Access Policies do and how the various elements that comprise an Access Policy are related. It is divided into the following two main sections:

ACCESS POLICIES DEFINED

A firewall provides a network boundary with a single point of entry and exit—a choke point. Because all incoming and outgoing traffic must pass through the choke point, you can screen and direct all that traffic through the implementation of a set of Access Policies—the Access Control List (ACL).

Access Policies allow you to permit, deny, encrypt, authenticate, prioritize, schedule, and monitor the traffic attempting to cross your firewall, whether incoming, outgoing, to the DMZ (NetScreen-10 and -100), or from the DMZ. You decide which users and what information can enter and leave, and when and where they can go..

> ***Note:*** *Access Policies set in the root system of the NetScreen-1000 do not affect Access Policies set in Virtual Systems.*

Anatomy of a Policy

An Access Policy must contain the following elements:

- Addresses (source and destination)
- Service
- Action (permit, deny, tunnel)

An Access Policy can also contain the following elements:

- VPN tunneling
- Authentication
- Logging
- Counting
- Traffic alarm settings
- Scheduling
- Traffic shaping

The remainder of this section examines each of the above elements in turn.

Policy Configuration Dialog Box

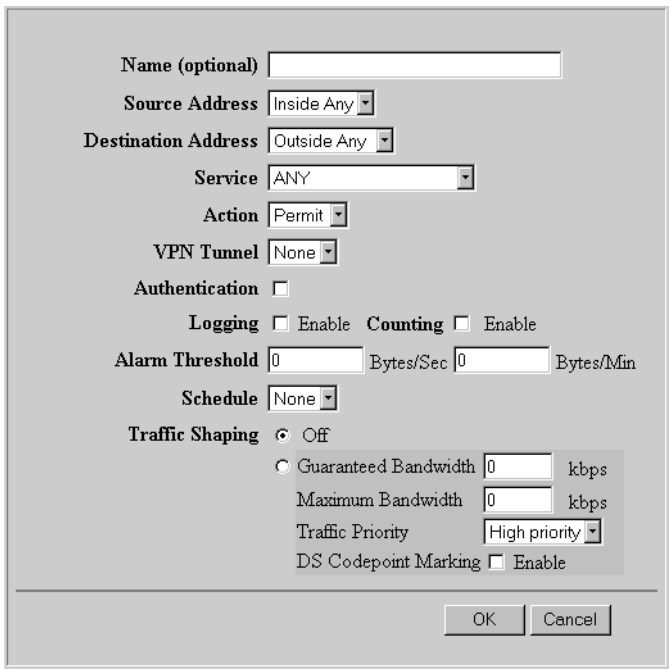

Addresses

Addresses are objects that identify network devices such as hosts and networks by their location in relation to the firewall—on the Trusted side, the Untrusted side, or in the DMZ (NetScreen-10 and -100). Individual hosts are specified using the mask 255.255.255.255, indicating that all 4 bytes of the address are significant. Networks are specified using their subnet mask to indicate which bytes are significant. To create an Access Policy for specific addresses, you must first create entries for the relevant hosts and networks in the address book.

You can also create address groups and apply Access Policies to them as you would to other address book entries.

> ***Note:*** *For more information on creating address groups, see "Address Groups" on page 4-5.*

When using address groups as elements of Access Policies, be aware that because the NetScreen device applies the Access Policy to each address in the group, the number of available Access Policies can become depleted more quickly than expected. This is a danger especially when you use address groups for both the source and destination.

Services

Services are objects that identify application protocols using layer 4 information such as standard and accepted TCP and UDP port numbers for application services like Telnet, FTP, SMTP, and HTTP. NetScreen includes predefined core Internet services. Additionally, the administrator can define custom services.

> ***Note:*** *For more information on both predefined and custom services, see "Services" on page 4-25.*

Services are defined in "Services" on page 4-25.

You can define Access Policies that specify which services are permitted, denied, encrypted, authenticated, logged, or counted, and which trigger an alarm.

Actions

Actions are objects that describe what the firewall does to the traffic it receives.

- **Permit** allows the packet to pass the firewall.
- **Deny** blocks the packet from traversing the firewall.
- **Tunnel** encrypts and authenticates data using IPSec (see "IPSec Concepts" on page 6-3). After selecting Tunnel, specify which VPN tunnel to use.

The NetScreen device applies the specified action on traffic that matches the first two criteria: addresses (source and destination) and service.

VPN Tunnel

You can apply a single Access Policy or multiple Access Policies to any VPN tunnel that you have configured. In the WebUI, the VPN Tunnel option provides a drop-down list of all such tunnels. In the CLI, you can see all available tunnels with the **get vpn** command.

Authentication

Selecting this option requires the user at the source address to authenticate his/her identity by supplying a user name and password before traffic is allowed to traverse the firewall or enter the VPN tunnel. The NetScreen device can use the internal user database or an external RADIUS, SecurID, or LDAP server to perform the authentication check.

Schedules

By associating a schedule to an Access Policy, you can determine when the Access Policy is in effect. You can configure schedules on a recurring basis and as a one-time event. Schedules provide a powerful tool in controlling the flow of network traffic and in enforcing network security. For an example of the latter, if you were concerned about employees transmitting important data outside the company, you might set an Access Policy that blocked outbound FTP-Put and MAIL traffic after normal business hours.

In the WebUI, define schedules in the Schedule section. In the CLI, use the set schedule command. For more information on setting schedules, see.

> ***Note:*** *In the WebUI, scheduled Access Policies appear in green to indicate that the current time is not within the defined schedule. When a scheduled Access Policy becomes active, it appears in red.*

Logging

When you enable logging in an Access Policy, the NetScreen device logs all connections to which that particular Access Policy applies. You can view the logs through either the WebUI or CLI, and the graphs in the Monitor section of the WebUI.

> ***Note:*** *For more information about viewing logs and graphs, see Chapter 9, "Monitoring NetScreen Devices" .*

Counting

When you enable counting in an Access Policy, the NetScreen device counts the total number of bytes of traffic to which this Access Policy applies and records the information in historical graphs.

Alarm Threshold

You can set a threshold that triggers an alarm when the traffic permitted by the Access Policy exceeds a specified number of bytes per second, bytes per minute, or both. Because the traffic alarm requires the NetScreen device to monitor the total number of bytes, you must also enable the counting feature.

Note: *For more information about traffic alarms, see "Alarms" on page 9-15.*

Traffic Shaping

You can set parameters for the control and shaping of traffic for each Access Policy. The traffic shaping parameters include:

Guaranteed Bandwidth: Guaranteed throughput in kilobits per second (kbps). Traffic below this threshold passes with the highest priority without being subject to any traffic management or shaping mechanism.

Maximum Bandwidth: Secured bandwidth available to the type of connection being specified in kilobits per second (kbps). Traffic beyond this threshold will be throttled and dropped.

Note: *It is advised that you do not use rates less than 10 kbps. Rates below this threshold lead to dropped packets and excessive retries that defeat the purpose of traffic management.*

Traffic Priority: When traffic bandwidth falls between the guaranteed and maximum settings, the NetScreen device passes higher priority traffic first, and lower priority traffic only if there is no other higher priority traffic. There are eight priority levels.

DS Codepoint Marking: Differentiated Services (DiffServ) is a system for tagging (or "marking") traffic at a position within a hierarchy of priority. The eight NetScreen priority levels can be mapped to the DiffServ system. By default, the highest priority (priority 0) maps to the first three bits (111) in the DS field (see RFC 2472) or the IP precedence field in the TOS byte (see RFC 1349) in the IP packet header. The lowest priority (priority 8) in the NetScreen system maps to 000 in the DiffServ system.

Note: *For a more detailed discussion of traffic management and shaping, see Chapter 7, "Traffic Shaping" .*

To change the mapping between the NetScreen priority levels and the DS system, use the following CLI command:

set traffic-shaping ip_precedence <number for priority 0 (highest priority)> <number for priority 1> <number for priority 2> <number for priority 3> <number for priority 4> <number for priority 5> <number for priority 6> <number for priority 7>

ACCESS POLICIES APPLIED

This section describes the management of Access Policies: viewing, creating, ordering and reordering, modifying, and removing Access Policies.

Viewing Access Policies

To view Access Policies through the WebUI, click **Policy >> Incoming | Outgoing | To DMZ | From DMZ**. In the CLI, use the **get policy** command.

Access Policy Icons

When viewing a list of Access Policies, the WebUI uses icons to provide you a graphical summary of policy components. The table below defines the different icons used in the Access Policies page.

Icon	Function	Description
	Permit	All traffic meeting the criteria is passed.
	Deny	All traffic meeting the criteria is denied.
	Encrypt enabled	All traffic meeting the criteria is encrypted.
	Encrypt disabled	There is a VPN configuration error (Action: Tunnel; VPN Tunnel: None), so no encryption is applied.
	Authenticate	The user must authenticate himself/herself when initiating a connection.
	Log	All traffic is logged and made available for syslog and e-mail, if enabled.
	Count	The amount of traffic is counted in bytes per second.
	Alarm	Indicates that you have set alarm thresholds.
	Traffic Shaping	Bandwidth shaping is active.

Icon	Function	Description
	Schedule	An Access Policy is only active during the time defined by the chosen schedule.

Creating Access Policies

Access Policies define the security of your network. You can set Access Policies to accept, deny, encrypt, and authenticate the network traffic travelling through the Netscreen device.

> ***Note:*** *The default policy for the NetScreen-10, -100, and 1000 is to deny all access. The NetScreen-5 default Access Policy denies all inbound traffic but allows all outbound traffic.*

Access Policy Location

You assign an Access Policy for one of four directions, based on the intended source and destination addresses: Incoming, Outgoing, To DMZ, or From DMZ.

The differences are categorized as follows:

TRAFFIC	Outgoing	Incoming	To DMZ	From DMZ
Source	Trusted	Untrusted	Trusted Untrusted	DMZ
Destination	Untrusted	Trusted MIP VIP	DMZ	Trusted Untrusted

Example: Typical ACL for a Small-to-Medium Enterprise

A small software firm, ABC Design, has divided its Trusted network into two subnets:

- Engineering (with the defined address “Engineering”)
- The rest of the company (with the defined address “Office”).

It also has a DMZ for its Web and mail servers.

The following example presents a typical set of Access Policies for the following users:

- Engineering is permitted to use all the services for outbound traffic except FTP-Put, IMAP, MAIL, and POP3.
- Office is permitted to use e-mail and access the Internet, provided they authenticate themselves.
- The entire company can access the company Web and mail servers on the DMZ.
- There is also a group of system administrators (with the defined address "Sys-admins"), who have complete user and administrative access to the servers on the DMZ.

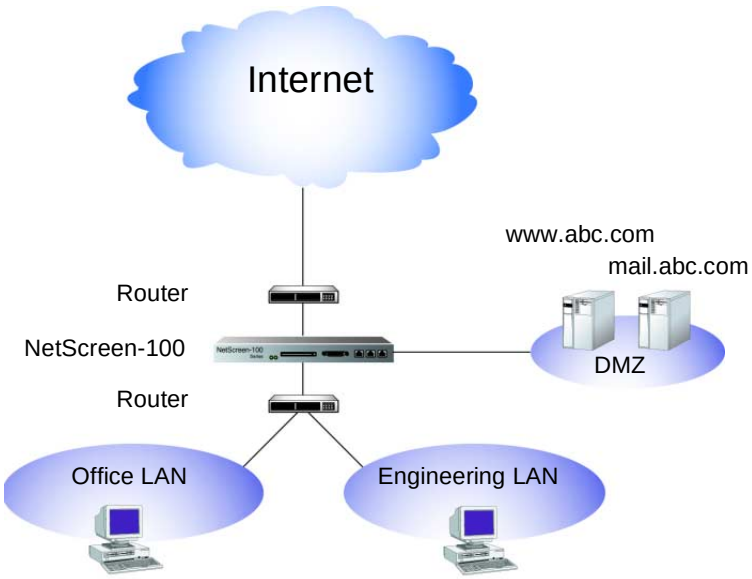

Outgoing

Source	Destination	Service	Action
Inside Any	Outside Any	Com (service group: FTP-Put, IMAP, MAIL, POP3)	Deny
Engineering	Outside Any	Any	Permit
Office	Outside Any	Internet (service group: FTP-Get, HTTP, HTTPS)	Permit (+ Authentication)

Incoming (Default Access Policy)

Source	Destination	Service	Action
Outside Any	Inside Any	Any	Deny

To DMZ

Source	Destination	Service	Action
Outside Any	mail.abc.com	MAIL	Permit
Outside Any	www.abc.com	Web (service group: HTTP, HTTPS)	Permit
Inside Any	mail.abc.com	e-mail (service group: IMAP, MAIL, POP3)	Permit
Inside Any	www.abc.com	Internet	Permit

From DMZ

Source	Destination	Service	Action
mail.abc.com	Outside Any	MAIL	Permit

WebUI

1. Policy >> Outgoing >> New Policy: Enter the following, and then click **OK**:
   Source Address: Inside Any
   Destination Address: Outside Any
   Service: Com[1]
   Action: Deny
2. Policy >> Outgoing >> New Policy: Enter the following, and then click **OK**:
   Source Address: Engineering
   Destination Address: Outside Any
   Service: ANY
   Action: Permit
3. Policy >> Outgoing >> New Policy: Enter the following, and then click **OK**:
   Source Address: Office
   Destination Address: Outside Any
   Service: Internet[2]
   Action: Permit
   Authentication: (select)

Note: *For incoming traffic, use the default Access Policy to deny everything.*

4. Policy >> To DMZ >> New Policy: Enter the following, and then click **OK**:
   Source Address: Outside Any
   Destination Address: mail.abc.com
   Service: MAIL
   Action: Permit

1. "Com" is a service group with the following members: FTP-Put, MAIL, IMAP, and POP3.
2. "Internet" is a service group with the following members: FTP-Get, HTTP, and HTTPS.

5. Policy >> To DMZ >> New Policy: Enter the following, and then click **OK**:
   Source Address: Outside Any
   Destination Address: www.abc.com
   Service: Web[3]
   Action: Permit
6. Policy >> To DMZ >> New Policy: Enter the following, and then click **OK**:
   Source Address: Inside Any
   Destination Address: mail.abc.com
   Service: e-mail[4]
   Action: Permit
7. Policy >> To DMZ >> New Policy: Enter the following, and then click **OK**:
   Source Address: Inside Any
   Destination Address: www.abc.com
   Service: Internet
   Action: Permit
8. Policy >> To DMZ >> New Policy: Enter the following, and then click **OK**:
   Source Address: Sys-admins
   Destination Address: DMZ Any
   Service: Any
   Action: Permit
9. Policy >> From DMZ >> New Policy: Enter the following, and then click **OK**:
   Source Address: mail.abc.com
   Destination Address: Outside Any
   Service: MAIL
   Action: Permit

3. "Web" is a service group with the following members: HTTP and HTTPS.
4. "e-mail" is a service group with the following members: MAIL, IMAP, and POP3.

CLI

1. set policy outgoing “inside any” “outside any” com deny
2. set policy outgoing engineering “outside any” any permit
3. set policy outgoing office “outside any” internet permit auth
4. set policy todmz “outside any” mail.abc.com mail permit
5. set policy todmz “outside any” www.abc.com web permit
6. set policy todmz “inside any” mail.abc.com e-mail permit
7. set policy todmz “inside any” www.abc.com internet permit
8. set policy todmz sys-admins “dmz any” any permit
9. set policy fromdmz mail.abc.com “outside any” mail permit
10. save

Modifying Access Policies

After you create an Access Policy, you can always return to it to make modifications. In the WebUI, you click the **Edit** link in the Configure column for the Access Policy that you want to change. In the Policy Configuration dialog box that appears for that Access Policy, make your changes and then click **OK**. In the CLI, you use the **set policy** command.

Example: Disabling an Access Policy through the Schedule Feature

NetScreen does not provide a specific method for enabling and disabling Access Policies. After you create an Access Policy, it is automatically enabled. However, you can use the schedule feature to effectively accomplish the same enabling and disabling function.

You must first, create a schedule for a one-time event that started and stopped in the past and name it “disable.” Then you apply that schedule to whatever Access Policy you want to disable. When you want to enable it again, change the schedule back to None (or to another schedule).

WebUI

Policy >> Incoming | Outgoing | To DMZ | From DMZ >> Edit: In the Schedule drop-down list, select **disable**, and then click **OK**.

CLI

1. set policy {incoming | outgoing | todmz | fromdmz} <source address> <destination address> <service> <action> schedule disable
2. save

Reordering Access Policies

The NetScreen device checks all attempts to traverse the firewall against Access Policies, beginning with the first one listed in the ACL for the appropriate direction (outgoing, incoming, to DMZ, from DMZ) and moving through the list. Because action applies to the first matching Access Policy, you must arrange them from the most specific to the most general. (Whereas a specific Access Policy does not preclude the application of a more general Access Policy located down the list, a general Access Policy appearing before a specific one does.)

To move an Access Policy to a different position in the ACL, do the following:

WebUI

1. Policy >> Incoming | Outgoing | To DMZ | From DMZ: Click the circular arrows in the Configure column to display the Move Policy Micro dialog box:

2. Change the order of the Access Policy to fit your needs, and then click the **OK** button.

   The Access Policies page reappears with the Access Policy you moved in its new position.

CLI

1. set policy move <id number> {before | after} <number>
2. save

Example: Reordering Home-to-Office Access Policies

By setting priority levels and guaranteed bandwidth levels for outbound traffic, you can ensure that important traffic always has enough bandwidth. At home, you might want to set up the following three Access Policies on your NetScreen-5 to ensure that you can still reach your office through your home-to-office VPN even when your children are playing games on the Internet. (These Access Policies also ensure that you have enough bandwidth to play games on the Internet when your children are doing the same thing.)

Outgoing Access Policies

ID	Source	Destination	Service	Action	Guaranteed and Maximum Bandwidth*	Priority
0	Inside Any	Outside Any	Any	Permit	0 Kbps	Low priority
1	Mom/Dad	corp-net	Any	Tunnel (VPN Tunnel: home-corp)	3500 Kbps	High priority
2	Mom/Dad	Outside Any	Any	Permit	1500 Kbps	2nd priority

* The bandwidth for the Trusted and Untrusted interface is set at 5 Mbps per interface.

Note that if the three Access Policies are ordered as shown above, the NetScreen device only applies the first Access Policy to outgoing traffic. You must move the Access Policy #0 to the bottom of the list.

WebUI

1. Policy >> Outgoing: Click the circular arrows in the Configure column for Access Policy ID #0.
2. In the Move Policy Micro dialog box that appears, enter the following, and then click **OK**:

   After: (select)

   ID: 2

CLI

1. set policy move 0 after 2
2. save

Removing an Access Policy

In addition to modifying an Access Policy, you can also delete it from the ACL. In the WebUI, you click **Remove** in the Configure column for the Access Policy that you want to remove. When the system message prompts for confirmation to proceed with the removal, click **Yes**. In the CLI, use the **unset policy <number>** command.

Virtual Private Networks 6

The first part of this chapter discusses the following aspects of virtual private network (VPN) technology:

The second part of this chapter presents examples of the following kinds of VPN tunnels:

INTRODUCTION TO VPNS

A virtual private network (VPN) provides a means for securely communicating between remote computers across a public wide area network (WAN), such as the Internet.

A VPN connection can link two local area networks (LANs) or a remote dialup user and a LAN. The traffic that flows between these two points passes through shared resources such as routers, switches, and other network equipment that make up the public WAN. To secure VPN communication while passing through the WAN, the two participants create an IP Security (IPSec) tunnel[1].

An IPSec tunnel consists of a pair of unidirectional Security Associations (SAs)—one at each end of the tunnel—that specify the security parameter index (SPI), destination IP address, and security protocol (Authentication Header or Encapsulating Security Payload) employed.

> ***Note:*** *For more information on SPIs, see "Security Association" on page 6-8. For more about the IPSec security protocols, see "Protocols" on page 6-6.*

Through the SA, an IPSec tunnel can provide the following security functions:

- Privacy (via encryption)
- Content integrity (via data authentication)
- Sender authentication and nonrepudiation (via data origin authentication)

The security functions you employ depend on your needs. If you only need to authenticate the IP packet source and content integrity, you can authenticate the packet without applying any encryption. On the other hand, if you are only concerned with preserving privacy, you can encrypt the packet without applying any authentication mechanisms. Optionally, you can both encrypt and authenticate the packet. Most network security designers choose to encrypt, authenticate, and replay-protect their VPN traffic.

NetScreen supports IPSec technology for creating VPN tunnels with the following three kinds of key creation mechanisms:

- Manual Key
- AutoKey IKE with a preshared key
- AutoKey IKE with a certificate

1. The term "tunnel" does not denote either transport or tunnel mode (see "Modes" on page 6-4). It simply refers to the IPSec connection.

IPSEC CONCEPTS

IP Security (IPSec) is a suite of related protocols for cryptographically securing communications at the IP packet layer. IPSec consists of two modes and two main protocols:

- Transport and tunnel modes
- The Authentication Header (AH) protocol for authentication and the Encapsulating Security Payload (ESP) protocol for encryption (and authentication)

IPSec also provides methods for the manual and automatic negotiation of Security Associations (SAs) and key distribution, all the attributes for which are gathered in a Domain of Interpretation (DOI) document.

IPSec Architecture

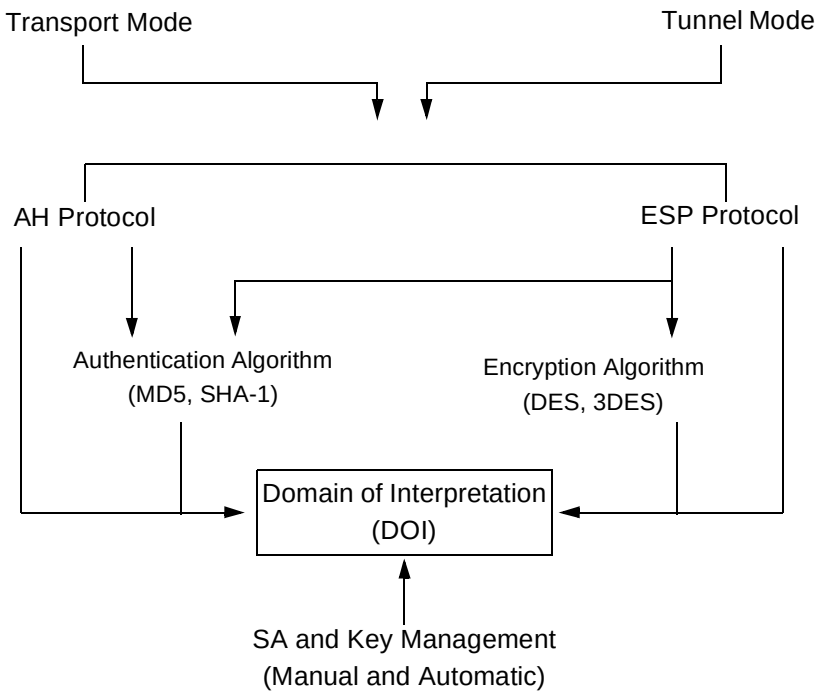

Note: *The IPSec Domain of Interpretation (DOI) is a document containing definitions for all the security parameters required for the successful negotiation of a VPN tunnel—essentially, all the attributes required for SA and IKE negotiations.*

Modes

IPSec operates in one of two modes: transport and tunnel. When both ends of the tunnel are hosts, transport mode is usually used. When at least one of the endpoints of a tunnel is a security gateway, such as a router or firewall, tunnel mode is usually used. NetScreen devices always operate in tunnel mode.

Transport mode

The original IP packet is not encapsulated within another IP packet. The entire packet can be authenticated (in AH), the payload can be encrypted (in ESP), and the original header remains in plaintext as it is sent across the WAN.

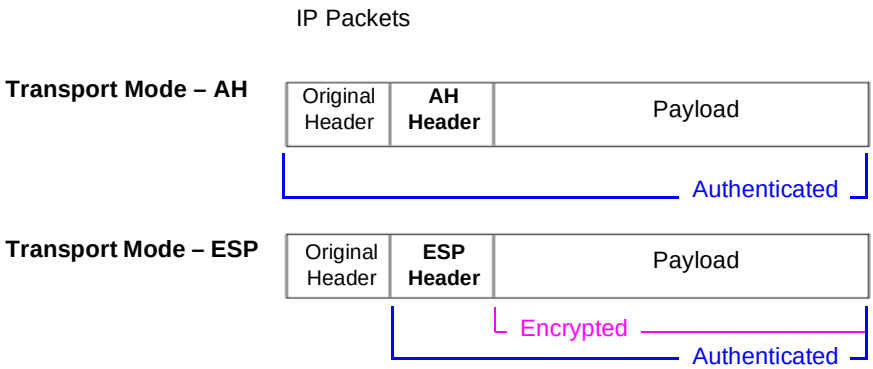

Tunnel mode

The entire original IP packet—payload and header—is encapsulated within another IP payload and a new header appended to it. The entire original packet can be encrypted, authenticated, or both. If authentication is applied, the new header is also authenticated.

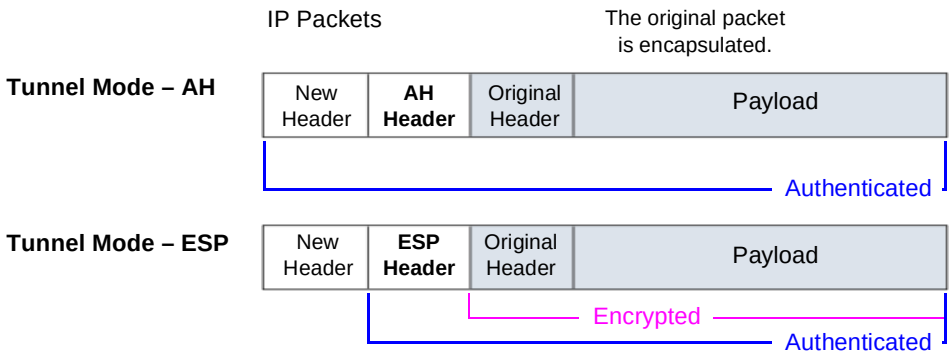

In a LAN-to-LAN VPN, the source and destination addresses used in the new header are the IP addresses of the NetScreen devices (the Untrusted IP address in NAT or Route mode; the System IP address in Transparent mode); the source and destination addresses of the encapsulated packets are the addresses of the ultimate endpoints of the connection.

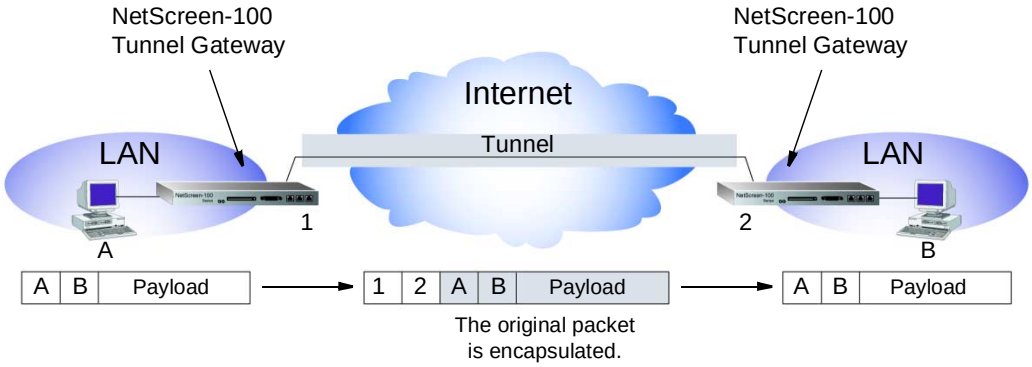

LAN-to-LAN VPN in Tunnel Mode

In a dialup-to-LAN VPN, there is no tunnel gateway on the VPN dialup client end of the tunnel; the tunnel extends directly to the client itself. In this case, on packets sent to the dialup client, both the new header and the encapsulated original header have the same IP address: that of the client's computer.

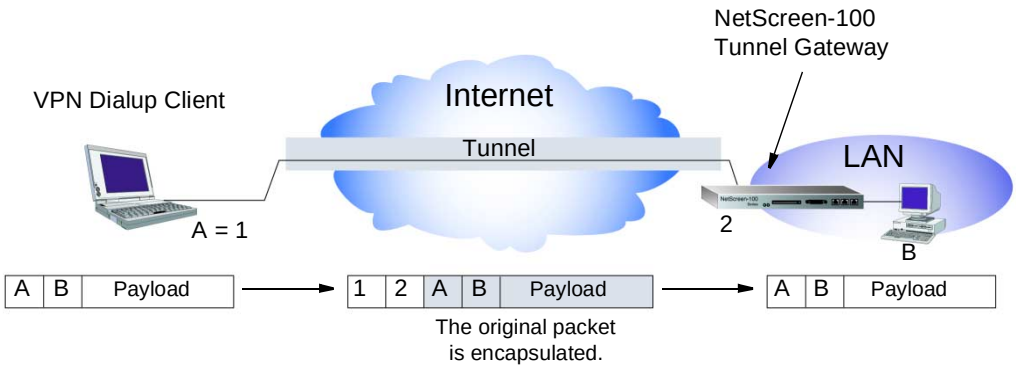

Dialup-to-LAN VPN in Tunnel Mode

Protocols

IPSec uses two protocols to secure communications at the IP layer:

- Authentication Header (AH)—A security protocol for authenticating the source of an IP packet and verifying the integrity of its content
- Encapsulating Security Payload (ESP)—A security protocol for encrypting the entire IP packet (and authenticating its source and content)

AH

The Authentication Header (AH) protocol provides a means to verify the authenticity/integrity of a packet's content and origin. You can authenticate the packet by the checksum calculated via a hash-based message authentication code (HMAC) using a secret key and either MD5 or SHA-1 hash functions.

Message Digest version 5 (MD5)—An algorithm that produces a 128-bit hash (also called a digital signature or message digest) from a message of arbitrary length and a 16-byte key. The resulting hash is used, like a fingerprint of the input, to verify content and source authenticity and integrity.

Secure Hash Algorithm-1 (SHA-1)—An algorithm that produces a 160-bit hash from a message of arbitrary length and a 20-byte key. It is generally regarded as more secure than MD5 because of the larger hashes it produces, and because the computational processing is done in the ASIC, the performance cost is negligible.

Note: *For more information on MD5 and SHA-1 hashing algorithms, see the following RFCs: (MD5) 1321, 2403; (SHA-1) 2404. For information on HMAC, see RFC 2104.*

ESP

The Encapsulating Security Payload (ESP) protocol provides a means to ensure privacy (encryption) and verify the source and content of data (authentication). ESP in tunnel mode encapsulates the entire IP packet (header and payload), and then appends a new IP header to the now encrypted packet. This new IP header contains the destination address needed to route the protected data through the network.

With ESP, you can encrypt and authenticate, encrypt only, or authenticate only. For encryption, you can choose either of the following encryption algorithms:

Data Encryption Standard (DES)—A cryptographic block algorithm with a 56-bit key.

Triple DES—A more powerful version of DES in which the original DES algorithm is applied in three rounds, using a 168-bit key. DES provides a significant performance savings but is considered unacceptable for many classified or sensitive material transfers.

For authentication, you can use either MD5 or SHA-1 algorithms.

For either the encryption or authentication algorithm you can select **NULL**; however, you cannot select **NULL** for both simultaneously.

Key Management

The distribution and management of keys are critical to successfully using VPNs. IPSec supports both manual and automatic key distribution methods. NetScreen offers three key management methods:

- Manual Key
- AutoKey IKE with preshared keys
- AutoKey IKE with certificates

Manual Key

With Manual Keys, administrators at both ends of a tunnel configure all of the security parameters. This is a viable technique for small, static networks where the distribution, maintenance, and tracking of keys is not difficult. However, safely distributing Manual Key configurations across great distances does pose security issues. Aside from passing the keys face-to-face, you cannot be completely sure that the keys have not been compromised. Also, whenever you want to change the key, you are faced with the same security issues as when you initially distributed it.

AutoKey IKE with Preshared Keys

When you need to create and manage numerous tunnels, you need a method that does not require you to configure every element manually. IPSec supports the automated generation and negotiation of keys and security associations using the Internet Key Exchange (IKE) protocol. NetScreen refers to such automated tunnel negotiation as AutoKey IKE and supports AutoKey IKE with preshared keys and AutoKey IKE with certificates.

With AutoKey IKE using preshared keys to authenticate the participants in an IKE session, each side must configure and securely exchange the preshared key[2] in advance. In this regard, the issue of secure key distribution is the same as that with Manual Keys. However, once distributed, an AutoKey, unlike a Manual Key, can automatically change its keys at predetermined intervals using the IKE

protocol. Frequently changing keys greatly improves security, and automatically doing so greatly reduces key-management responsibilities. However, changing keys increases traffic overhead; therefore, doing so too often can reduce data transmission efficiency.

AutoKey IKE with Certificates

When using certificates to authenticate the participants during an AutoKey IKE negotiation, each side generates a public/private key pair (see "Public Key Cryptography" on page 6-14) and acquires a certificate (see "Certificates and CRLs" on page 6-17). As long as the issuing certificate authority (CA) is trusted by both sides, the participants can retrieve the peer's public key and verify the peer's signature. There is no need to keep track of the keys and SAs; IKE does it automatically.

> ***Note:*** *For examples of both Manual Key and AutoKey IKE tunnels, see "Example VPN Scenarios" on page 6-24.*

Security Association

A security association (SA) is a unidirectional agreement between the VPN participants regarding the methods and parameters to use in securing a communication channel. Full bidirectional communication requires at least two SAs, one at each end of the tunnel.

An SA groups together the following components for securing communications:

- Security algorithms and keys
- Protocol mode (transport or tunnel)
- Key management method (Manual Key or AutoKey IKE)
- SA lifetime

For outbound VPN traffic, the Access Policy invokes the SA associated with the VPN tunnel. For inbound traffic, the NetScreen device looks up the SA by using the following triplet: destination IP, protocol, and security parameter index (SPI) value.

---

2. A preshared key is a key for both encryption and decryption that both participants must have before initiating communication.

TUNNEL NEGOTIATION

The establishment of an IPSec tunnel can take place in either one or two phases. For a Manual Key tunnel, because all of the security association (SA) parameters have been previously defined, there is no need to negotiate which SAs to use. In essence, the tunnel has already been established. When traffic matches an Access Policy using that Manual Key tunnel, the NetScreen devices simply encrypts and authenticates the data, as you determined, and sends it off to the destination gateway.

To establish an AutoKey IKE tunnel, two phases of negotiation are required:

- In Phase 1, the participants establish a secure channel in which to negotiate the SAs.
- In Phase 2, the participants negotiate the SAs for encrypting and authenticating the ensuing traffic.

Phase 1

Phase 1 of an AutoKey IKE tunnel negotiation consists of the exchange of proposals for how to authenticate and secure the channel in which the session encryption key will be created. The exchange can be in one of two modes: Aggressive mode or Main mode (see below). Using either mode, the participants exchange proposals for acceptable security services such as:

- Encryption algorithms (DES and 3DES) and authentication algorithms (MD5 and SHA-1). For more information about these algorithms, see “Protocols” on page 6-6.
- A Diffie-Hellman Group (See “The Diffie-Hellman Exchange” on page 6-11.)
- Preshared Key or RSA/DSA certificates (see “AutoKey IKE with Preshared Keys” on page 6-7)

A successful Phase 1 negotiation concludes when both ends of the tunnel agree to accept at least one set of the Phase 1 security parameters proposed, and then process them. NetScreen devices support up to four proposals for Phase 1 (and up to four proposals for Phase 2) negotiations, allowing you to define how restrictive a range of tunnel parameters you will accept.

Main Mode and Aggressive Mode

Phase 1 can take place in either Main mode or Aggressive mode. The two modes are described below.

Main Mode: The initiator and recipient send three two-way exchanges (six messages total) to accomplish the following services:

- First exchange, (messages 1 and 2): Propose and accept the encryption and authentication algorithms.
- Second exchange, (messages 3 and 4): Execute a Diffie-Hellman exchange, and provide a nonce (randomly generated number) for the other to sign and return to authenticate his or her identity.
- Third exchange, (messages 5 and 6): Send and verify their identities.

The information transmitted in the third exchange of messages is protected by the encryption algorithm established in the first two exchanges. Thus, the participants' identities are not transmitted in the clear.

Aggressive Mode: The initiator and recipient accomplish the same objectives, but only in two exchanges, and a total of three messages:

- First message: The initiator proposes the SA, and sends a nonce, identity, and, if using certificates, the initiator's certificate.
- Second message: The recipient accepts the SA, authenticates the initiator, and sends a nonce, identity, and, if using certificates, the recipient's certificate.
- Third message: The initiator authenticates the recipient and confirms the exchange.

Because the participants' identities are exchanged in the clear (in the first two messages), Aggressive mode does not provide identity protection.

> ***Note:*** *When a dialup VPN user negotiates an AutoKey IKE tunnel with a preshared key, Aggressive mode must be used.*

The Diffie-Hellman Exchange

A Diffie-Hellman exchange allows the participants to produce a shared secret value. The strength of the technique is that it allows the participants to create the secret value over an unsecured medium without passing the secret value through the wire. There are five Diffie-Hellman (DH) groups (NetScreen supports groups 1, 2, and 5). The size of the prime modulus used in each group's calculation differs as follows:

- DH Group 1: 768-bit modulus
- DH Group 2: 1024-bit modulus
- DH Group 5: 1536-bit modulus

The larger the modulus, the more secure the generated key is considered to be; however, the larger the modulus, the longer the key-generation process takes. Because the modulus for each group is a different size, the participants must agree to use the same group.

Phase 2

After the participants have established a secure and authenticated channel, they proceed through Phase 2, in which they negotiate the SAs to secure the data to be communicated through the IPSec tunnel.

Similarly to the process for Phase 1, the participants exchange proposals to determine which security parameters to employ in the SA. A Phase 2 proposal also includes a security protocol—either Encapsulating Security Payload (ESP) or Authentication Header (AH), and selected encryption and authentication algorithms. The proposal can also specify a Diffie-Hellman group, if Perfect Forward Secrecy (PFS) is desired.

> ***Note:*** *For more about Diffie-Hellman groups, see "The Diffie-Hellman Exchange" on page 6-11. For more about PFS, see "Perfect Forward Secrecy" on page 6-12.*

NetScreen devices support up to four proposals for Phase 2 negotiations (and up to four proposals for Phase 1). NetScreen also provides a replay protection feature. Use of this feature does not require negotiation because packets are always sent with sequence numbers. You simply have the option of checking the sequence numbers or not. (For more information, see below.)

Perfect Forward Secrecy

Perfect Forward Secrecy (PFS) is a method for deriving Phase 2 keys independent from and unrelated to the preceding keys. Ordinarily, the Phase 1 proposal creates the key (the SKEYID_d key) from which all Phase 2 keys are derived. The SKEYID_d key can generate Phase 2 keys with a minimum of CPU processing. Unfortunately, if an unauthorized party gains access to the SKEYID_d key, all your encryption keys are compromised.

PFS addresses this security risk by forcing a new Diffie-Hellman key exchange to occur for each Phase 2 tunnel. Using PFS is thus more secure. Although the first rekeying procedure in Phase 2 might take slightly longer with PFS enabled, PFS does not slow subsequent rekeying because the computation is done before the lifetime of the currently active key expires.

Replay Protection

A replay attack occurs when somebody intercepts a series of packets and uses them later either to flood the system, causing a denial-of-service (DoS), or to gain entry to the Trusted network. The replay protection feature enables NetScreen devices to check every IPSec packet to see if it has been received before. If packets arrive outside a specified sequence range, the NetScreen device rejects them.

Packet Flow: LAN-to-LAN

To see the various components comprising the creation of an IPSec tunnel in relation to each other, the following example illustrates how a packet flows through a tunnel.

A company based in Tokyo has just opened a branch office in Paris and needs to connect the two sites through an IPSec tunnel. The tunnel uses Manual key, the ESP protocol, 3DES for encryption and SHA-1 for authentication.

The NetScreen devices protecting each site are in NAT mode. The addresses are as follows:

- Tokyo:
  - Untrusted interface: 201.22.3.14
  - Trusted LAN: 192.168.10.0/24
- Paris:
  - Untrusted interface: 203.3.3.10
  - Trusted LAN: 172.16.5.0/24

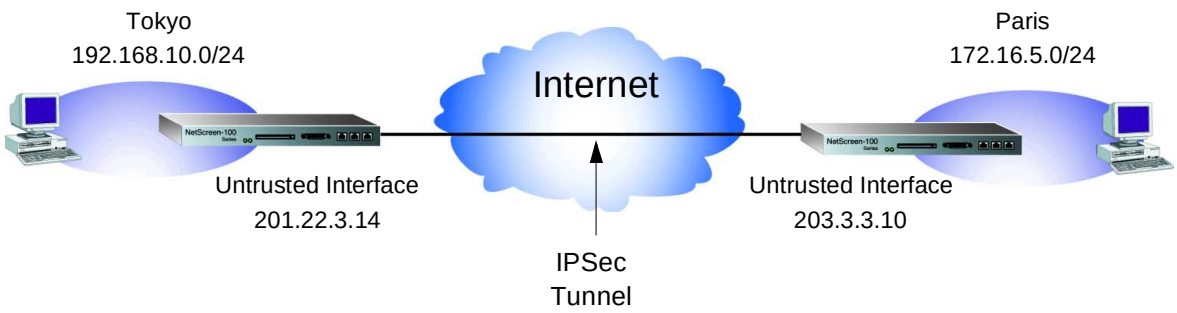

The path of a packet coming from 192.168.10.10 and going to 172.16.5.20 through an IPSec tunnel proceeds as follows:

1. The host on the 192.168.10.0/24 subnet sends a packet to a server on the 172.16.5.0 subnet.
2. The packet reaches its gateway; that is, the Netscreen device in Tokyo.
3. The Netscreen device performs the following operations:
   - It checks its Access Control List (ACL) and (using the source and destination address) determines to send the packet through the VPN tunnel to the Paris office.
   - It encrypts the entire packet (including the header) and puts a new header on the packet—changing the source IP address to 201.22.3.14, and the destination IP address to 203.3.3.10.
   - It sends the packet to 203.3.3.10; that is, the Untrusted IP address of the NetScreen device in Paris.
4. The remote Netscreen device performs the following operations:
   - Using the SPI, destination IP address, and IPSec protocol contained in the outer packet header, it locates the SA and keys.
   - It decrypts the packet, uncovering its ultimate destination.
   - It checks the Access Control List (ACL) and, finding an Access Policy that grants access.

PUBLIC KEY CRYPTOGRAPHY

In public key cryptography, a public/private key pair is used to encrypt and decrypt data. Data encrypted with a public key, which the owner makes available to the public, can only be decrypted with the corresponding private key, which the owner keeps secret and protected. For example, if Alice wants to send Bob an encrypted message, Alice can encrypt it with Bob's public key and send it to him. Bob then decrypts the message with his private key.

The reverse is also useful; that is, encrypting data with a private key and decrypting it with the corresponding public key. This is known as creating a digital signature. For example, if Alice wants to verify her identity as the sender of a message, she can encrypt the message with her private key and send the message to Bob. Bob then decrypts the message with Alice's public key, thus verifying the authenticity of Alice as the sender.

Public/private key pairs also play an important role in the use of digital certificates. The Certificate Authority (CA) that issues a certificate "signs" a portion of data by encrypting the data with its private key. The recipient then uses the CA's public key to decrypt the signed portion of data. If it decrypts properly, the recipient can confirm the integrity of the CA's signature and, by extension, the integrity of the accompanying certificate.

PKI

The term Public Key Infrastructure (PKI) refers to the hierarchical structure of trust required for the successful implementation of public key cryptography. To verify the trustworthiness of a certificate, you must be able to track a path of certified CAs from the one issuing your local certificate back to a root authority of a CA domain.

PKI Hierarchy of Trust – CA Domain

The root level CA validates subordinate CAs.

Subordinate CAs validate local certificates and other CAs.

Local certificates contain the user's public key.

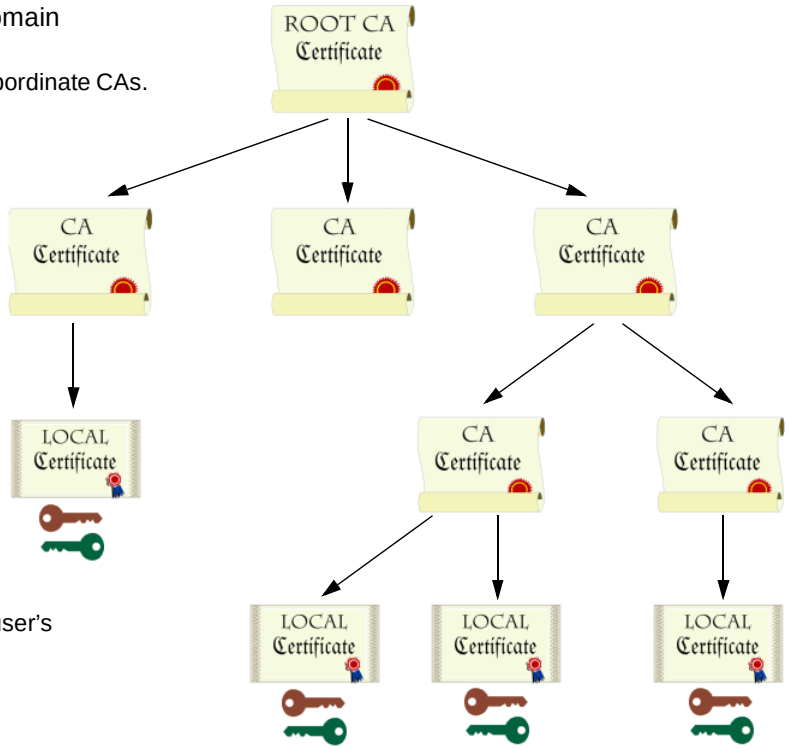

If certificates are used solely within an organization, that organization can have its own CA domain within which a company CA issues and validates certificates among its employees. If that organization later wants its employees to be able to exchange their certificates with those from another CA domain (for example, with employees at another organization that also has its own CA domain), the two CAs can develop cross-certification; that is, they can agree to trust the authority of each other. In this case, the PKI structure does not extend vertically but horizontally.

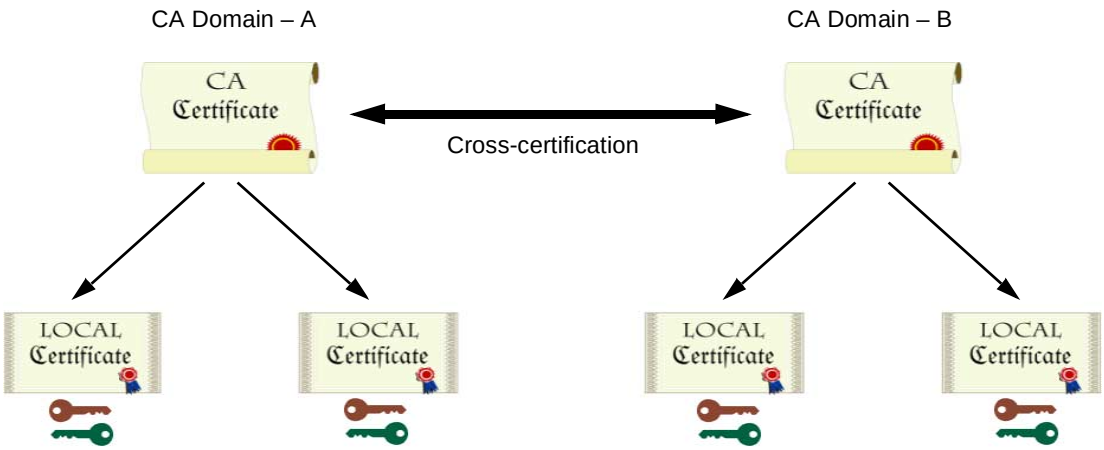

Users in CA domain A can use their certificates and key pairs with users in CA domain B because the CAs have cross-certified each other.

For convenience and practicality, PKI must be transparently managed and implemented. Toward this goal, the NetScreen ScreenOS does the following:

1. Generates a public/private key pair when you create a certificate request.
2. Supplies that public key as part of the certificate request in the form of a text file for transmission to a Certificate Authority (CA) for certificate enrollment.
3. Supports loading the local certificate, the CA certificate, and the certificate revocation list (CRL)[3] into the unit.

   You can also specify an interval for refreshing the CRL online. For more information on CRLs, see “Certificates and CRLs” on page 6-17.

4. Provides public key and certificate delivery when establishing an IPSec tunnel.
5. Supports certificate path validation upward through eight levels of CA authorities in the PKI hierarchy.
6. Supports the PKCS #7 cryptographic standard, which means the NetScreen device can accept X.509 certificates and CRLs packaged within a PKCS #7 envelope. PKCS #7 support allows you to submit multiple X.509 certificates within a single PKI request. You can now configure PKI to validate all the submitted certificates from the issuing CA at one time.

Certificates and CRLs

A digital certificate is an electronic means for verifying your identity through the word of a trusted third party, known as a Certificate Authority (CA). The CA server you use can be owned and operated by an independent CA[4], or by your own organization, in which case you become your own CA. If you use an independent CA, you must contact them for the addresses of their CA and CRL servers (for obtaining certificates and certificate revocation lists), and for the information they require when submitting personal certificate requests. When you are your own CA, you make the rules.

To use a digital certificate to authenticate your identity when establishing a secure VPN connection, you must first do the following:

- Obtain a personal certificate from a CA, and load the certificate in the NetScreen device
- Obtain a CA certificate for the CA that issued the personal certificate (basically verifying the identity of the CA verifying you), and load the CA certificate in the NetScreen device
- Obtain a CRL, and load that in the NetScreen device

3. The Certificate Authority usually provides a CRL. Although you can load a CRL into the NetScreen device, you cannot view it once loaded.

4. NetScreen supports the following CAs: Baltimore®, Entrust®, Microsoft, Netscape, RSA Keon®, and Verisign®.

During the course of business, there are several events that make it necessary to revoke a certificate. You might wish to revoke a certificate if you suspect that it has been compromised or when a certificate holder leaves a company. Managing certificate revocations and validation can be accomplished locally (which is a limited solution) or by referencing a Certificate Authority's CRL. You access a CA's CRL online, at given intervals.

Obtaining a Certificate

To obtain a signed digital certificate, you must complete several tasks in the following order:

1. Configure default server settings.
2. Generate a public/private key pair.
3. Fill out the Certificate Request.
4. Submit your request to your CA of choice.
5. After you receive your signed certificate, you must load it into the NetScreen device along with the CA certificate and the CRL.

You now have the following items for the following uses:

- A local certificate for your machine, to validate your identity with each tunnel connection
- A CA Certificate (their public key), to be used to verify the peer's certificate
- A Certificate Revocation List (CRL), to identify invalid certificates

When you receive these three files (the certificate files have the extension .cer, and the CRL has the extension .crl), load them into your NetScreen using the procedure described in the following section.

> ***Note:*** *If you are planning to use e-mail to submit a PKCS10 file to obtain your certificates, you must properly configure your NetScreen settings so that you can send e-mail to your system administrator. You have to modify the route table, correctly set your primary and secondary domain servers, and specify the SMTP server and e-mail address settings.*

Example: Configuring Default Server Settings for a CRL

In Phase 1 negotiations, participants check the CRL list to see if certificates received during an IKE exchange are still valid. It is not mandatory to load the CRL. If a CRL is not loaded, the NetScreen device tries to retrieve the CRL through the LDAP or HTTP[5] CRL location defined within the certificate itself or in the server settings field.

> ***Note:*** *With ScreenOS 2.5 and later, you can disable CRL checking during X.509 certificate validation. While this feature does make it easier to perform X.509 certificate validation, disabling CRL checking compromises the security of your NetScreen device.*

This example shows how to configure the default server to check the CRL daily by going to the LDAP server at 2.2.2.121 and locating the CRL file.

WebUI

VPN >> Certificates >> Default Server Settings: Enter the following, and the click **OK**:

LDAP Server: 2.2.2.121

CRL URL: ldap:///CN=Netscreen, CN=ns2001,CN=Public KeyServices, CN=Services,CN=Configuration,DC=NS2001,DC=com?CertificateRevocationList?base?objectclass=CRLDistribution Point

CRL Refresh Frequency: Daily

X509 Cert_Path Validation Level: Full

CLI

1. set pki ldap server-name 2.2.2.121
2. set pki ldap crl-url ldap:///CN=Netscreen,CN=ns2001,CN=PublicKey Services,CN=Services,CN=Configuration,DC=NS2000,DC=com?Certificate RevocationList?base?objectclass=CRLDistributionPoint
3. set pki x509 default crl-refresh daily
4. set pki x509 default cert-path full

5. The CRL distribution point extension within an X509 certificate can be either an HTTP URL or an LDAP URL.

Example: Requesting a Certificate

When you request a certificate, the NetScreen device generates a key pair. The public key becomes incorporated in the request itself and, eventually, in the digitally signed local certificate you receive from the CA.

In the following example, the security administrator is making a certificate request for Michael Zhang in the Development department at NetScreen Technologies in Santa Clara, California. The certificate is going to be used for a NetScreen-100 at IP address 10.10.5.44. The administrator instructs the NetScreen device to write the request to a file, which he then copies and pastes in the certificate request text field at the CA's certificate enrollment site. After the enrollment process is complete, the CA usually sends the certificates via e-mail to the security administrator.

WebUI

1. VPN >> Certificates >> Certificate Request: Enter the following, and then click **Generate**:

   Name: Michael Zhang

   Phone: (408) 330-7800

   Unit/Department: Development

   Organization: NetScreen Technologies

   County/Locality: Santa Clara

   State: CA

   Country: US

   Email: (leave blank; some CAs do not support this field.)

   IP Address: 10.10.5.44

   Write to file: (select)

   Create new key pair of *1024*[6] length: (select)

   The NetScreen generates a PKCS #10 file and prompts you to open the file or save it to disk.

6. The value 1024 indicates the bit length of the key pair. If you are using the certificate for SSL (see "Secure Sockets Layer" on page 3-3), be sure to use a bit length that your Web browser also supports.

2. Open the file, and copy its contents, taking care to copy the entire text but not any blank spaces before or after the text. (Start at "-----BEGIN CERTIFICATE REQUEST-----", and end at "-----END CERTIFICATE REQUEST-----".)
3. Follow the certificate request directions at the CA's Web site, pasting the PKCS #10 file in the appropriate field when required.
4. When you receive the certificate from the CA via e-mail, copy it to a text file, and save it to your workstation (to be loaded to the NetScreen device later through the WebUI) or to a TFTP server (to be loaded later through the CLI).

CLI

1. set pki x509 dn country-name US
2. set pki x509 dn state-name CA
3. set pki x509 dn local-name "Santa Clara"
4. set pki x509 dn org-name "NetScreen Technologies"
5. set pki x509 dn org-unit-name Development
6. set pki x509 dn name "Michael Zhang"
7. set pki x509 dn ip 10.10.5.44
8. set pki x509 default send-to mzhang@netscreen.com[7]
9. exec pki rsa new-key 1024

   The certificate request is sent via e-mail to mzhang@netscreen.com.
10. Copy the contents of the request, taking care to copy the entire text but not any blank spaces before or after the text. (Start at "-----BEGIN CERTIFICATE REQUEST-----", and end at "-----END CERTIFICATE REQUEST-----".)
11. Follow the certificate request directions at the CA's Web site, pasting the PKCS #10 file in the appropriate field when required.

    When you receive the certificate from the CA via e-mail, copy it to a text file, and save it to your workstation (to be loaded to the NetScreen device later through the WebUI) or to a TFTP server (to be loaded later through the CLI).

7. Using the e-mail address assumes that you have already configured the IP address for your SMTP server (**set admin mail server-name <IP address or server name>**).

Example: Loading Certificates

The CA returns the following three files to you for loading into the NetScreen device:

- A CA certificate, which contains the CA's public key
- A local certificate that identifies your local machine (your public key)
- A CRL, which lists any certificates revoked by the CA

For the WebUI example, you have downloaded the files to the administrator's workstation. For the CLI example, you have downloaded them to a TFTP server with IP address 198.168.1.5.

> ***Note:*** *NetScreen devices, including virtual systems, configured with ScreenOS 2.5 or later support loading multiple local certificates from different CAs.*

This example illustrates how to load two certificate files named auth.cer (CA certificate) and local.cer (your public key), and the CRL file named distrust.crl.

WebUI

1. VPN >> Certificates >> Load: Select **Certificate**, and then click **Browse**.
2. Navigate to the certificate file directory, select auth.cer, and then click **Open**.

   The directory path and file name appear in the Browse field in the Load Certificate or CRL dialog box.
3. Click **Load**.

   The certificate file loads, and then the Load Certificate or CRL dialog box reappears.
4. Click **Browse** again, navigate to the certificate file directory, select local.cer, and then click **Open**.

   The directory path and file name for local.cer appear in the Browse field in the Load Certificate or CRL dialog box.
5. Click **Load**.

   The certificate file loads, and then the Load Certificate or CRL dialog box reappears.
6. Select **CRL**, click **Browse** once more.

7. Navigate to the certificate file directory, select distrust.crl, and then click **Open**.
8. Click **Load**.

   The CRL file loads.

CLI

1. exec pki x509 tftp 198.168.1.5 cert-name auth.cer
2. exec pki x509 tftp 198.168.1.5 cert-name local.cer
3. exec pki x509 tftp 198.168.1.5 crl-name distrust.crl

Example VPN Scenarios

This section illustrates the following VPN designs:

- LAN-to-LAN VPNs
- Dialup-to-LAN VPNs
- Hub-and-Spoke VPN

LAN-to-LAN VPNs

This section describes the procedures for setting up three types of LAN-to-LAN VPNs. When both gateways have static Untrusted IP addresses, you can configure the following:

- LAN-to-LAN VPN, Manual Key tunnel
- LAN-to-LAN VPN, AutoKey IKE tunnel (with a preshared key or certificates)

When one gateway has a static and one has a dynamic Untrusted IP address, you can configure

- Dynamic Peer LAN-to-LAN VPN, AutoKey IKE tunnel (with a preshared key or certificates)

As used here, a static LAN-to-LAN VPN involves an IPSec tunnel connecting two LANs, each with a NetScreen device operating as a secure gateway with a fixed Untrusted IP address (or to the System IP address if in Transparent mode).

If one of the NetScreen devices has a dynamically assigned Untrusted IP address (such as a NetScreen-5 using DHCP or PPPoE, or a NetScreen-10 using DHCP), that device is termed a dynamic peer and the VPN is configured differently.

Example: LAN-to-LAN VPN, Manual Key

In this example, a Manual Key tunnel provides a secure communication channel between the offices in Tokyo and Paris, using ESP with 3DES encryption and SHA-1 authentication. The NetScreen devices protecting each site are in NAT mode. The addresses are as follows:

- Tokyo:
  - Untrusted interface: 201.22.3.14
  - Trusted LAN: 192.168.10.0/24
- Paris:
  - Untrusted interface: 203.3.3.10
  - Trusted LAN: 172.16.5.0/24

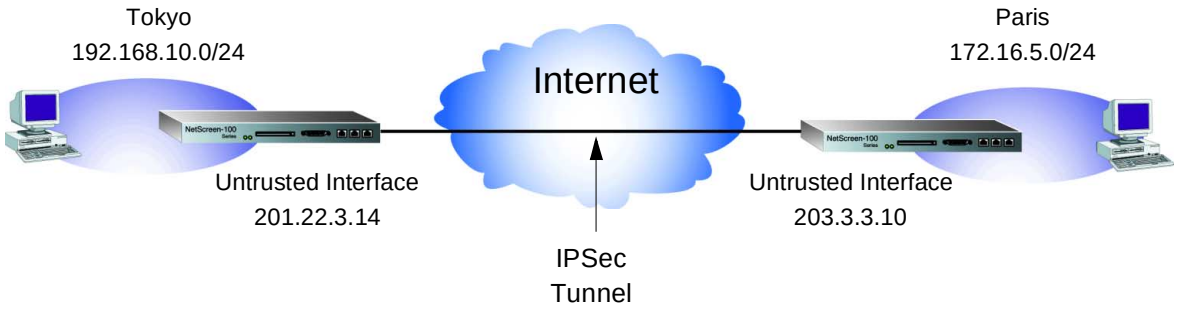

To set up the tunnel, perform the following three steps:

1. Enter the address for the remote endpoint of the tunnel in the Untrusted address book.
2. Configure the devices at each end of the tunnel with the same options and parameters.
3. Set up an outgoing Access Policy for VPN traffic to pass through the tunnel.

WebUI (Tokyo)

1. Address >> Untrusted >> New Address: Enter the following, and then click **OK**:

   Name: Paris_office

   IP Address/Domain Name: 172.16.5.0

   Netmask: 255.255.255.0

2. VPN >> Manual Key >> New VPN Entry: Enter the following, and then click **OK**:

   VPN Tunnel Name: Tokyo_Paris
   Gateway IP: 203.3.3.10
   Security Index: 3020 (Local), 3030 (Remote)
   ESP-CBC Algorithm: 3DES-CBC
   Generate Key by Password: asdlk24234
   Authentication Algorithm: SHA-1
   Generate Key by Password: PNas134a

3. Policy >> New Policy: Enter the following, and then click **OK**:

   Name: Tokyo to Paris
   Source Address: Inside any
   Destination Address: Paris_Office
   Service: ANY
   Action: Tunnel
   VPN Tunnel: Tokyo_Paris

Note: *Select* ***Authenticate*** *to require the user to enter his login name and password. Selecting* ***Authenticate*** *at both ends of the tunnel forces the user to log in twice.*

WebUI (Paris)

1. Address >> Untrusted >> New Address: Enter the following, and then click **OK**:

   Name: Tokyo_office
   IP Address/Domain Name: 192.168.10.0
   Netmask: 255.255.255.0

2. VPN >> Manual Key >> New VPN Entry: Enter the following, and then click **OK**:

   VPN Tunnel Name: Paris_Tokyo
   Gateway IP: 201.22.3.14
   Security Index: 3030 (Local), 3020 (Remote)
   ESP-CBC Algorithm: 3DES-CBC
   Generate Key by Password: asdlk24234
   Authentication Algorithm: SHA-1
   Generate Key by Password: PNas134a

3. Policy >> New Policy: Enter the following, and then click **OK**:

   Name: Paris to Tokyo
   Source Address: Inside any
   Destination Address: Tokyo_office
   Service: ANY
   Action: Tunnel
   VPN Tunnel: Paris_Tokyo

CLI (Tokyo)

1. set address untrust paris_office 172.16.5.0 255.255.255.0
2. set vpn tokyo_paris manual 3020 3030 gateway 203.3.3.10 esp 3des password asdlk24234 auth sha-1 password PNas134a
3. set policy outgoing “Tokyo to Paris” “inside any” paris_office any tunnel vpn tokyo_paris

CLI (Paris)

1. set address untrust tokyo_office 192.168.10.0 255.255.255.0
2. set vpn paris_tokyo manual 3030 3020 gateway 201.22.3.14 esp 3des password asdlk24234 auth sha-1 password PNas134a
3. set policy outgoing “Paris to Tokyo” “inside any” tokyo_office any tunnel vpn paris_tokyo

Example: LAN-to-LAN VPN, AutoKey IKE (Preshared Key or Certificates)

In this example, an AutoKey IKE tunnel using either a preshared secret or a pair of certificates (one at each end of the tunnel) provides the secure connection between the Tokyo and Paris offices. The tunnel again uses ESP with 3DES encryption and SHA-1 authentication.

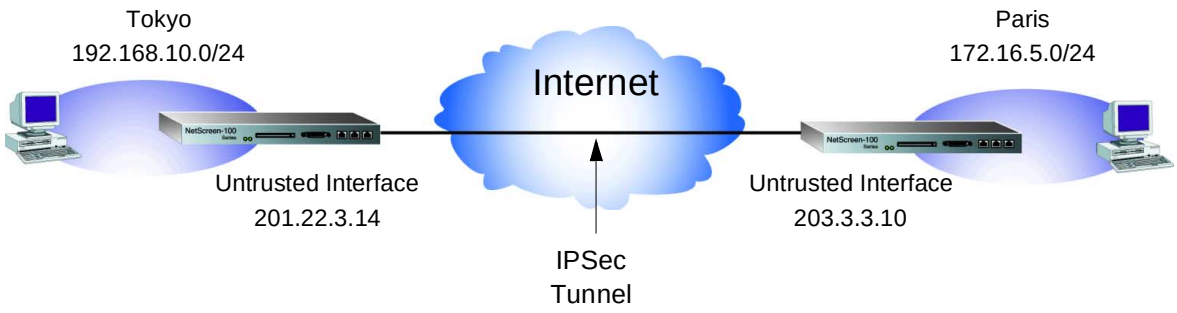

Setting up the AutoKey IKE tunnel using AutoKey IKE with either a preshared secret or certificates involves the following two steps:

1. Define the remote gateway and key exchange mode, and specify either a preshared secret or a certificate.
2. Create the Autokey IKE VPN entry.

> ***Note:*** *You also need to create an address book entry for the remote endpoint of the tunnel and configure an outgoing Access Policy. However, because those steps are the same as those explained in "Example: LAN-to-LAN VPN, Manual Key" on page 6-25, they are omitted here.*

In the following examples, the preshared key is "hopsalong3." It is assumed that both participants already have certificates. (For more information about obtaining and loading certificates, see "Certificates and CRLs" on page 6-17.)

WebUI (Tokyo)

1. VPN >>Gateway >>New Remote Gateway: Enter the following, and then click **OK**:

   Name: To_Paris

   Remote Gateway Static IP Address: 203.3.3.10

   Mode: Main

   For preshared key:

   Phase 1 Proposal: pre-g2-3des-sha

   Preshared Key: hopsalong3

   For certificate:

   Phase 1 Proposal: rsa-g2-3des-sha

2. VPN >> AutoKey IKE >> New AutoKey IKE Entry: Enter the following, and then click **OK**:

   Name: Tokyo_Paris

   Remote Gateway Tunnel Name: To_Paris

   Phase 2 Proposal: nopfs-esp-3des-sha1

WebUI (Paris)

1. VPN >> Gateway >> New Remote Gateway: Enter the following, and then click **OK**:

   Name: To_Tokyo

   Remote Gateway Static IP Address: 201.22.3.14

   Mode: Main

   Phase 1 Proposal: pre-g2-3des-sha1

   For preshared key:

   Phase 1 Proposal: pre-g2-3des-sha

   Preshared Key: hopsalong3

   For certificate:

   Phase 1 Proposal: rsa-g2-3des-sha

   Preferred certificate Peer CA: Entrust

   Preferred certificate Type: X509-SIG

2. VPN >> AutoKey IKE >> New AutoKey IKE Entry: Enter the following, and then click **OK**:

   Name: Paris_Tokyo

   Remote Gateway Tunnel Name: To_Tokyo

   Phase 2 Proposal: nopfs-esp-3des-sha1

CLI (Tokyo)

Preshare Key:

1. set ike gateway to_paris ip 203.3.3.10 main preshare hopsalong3 proposal pre-g2-3des-sha-1
2. set vpn tokyo_paris gateway to_paris tunnel proposal nopfs-esp-3des-sha-1

Certificate:

1. set ike gateway to_paris ip 203.3.3.10 main proposal rsa-g2-3des-sha
2. set vpn tokyo_paris gateway to_paris tunnel proposal nopfs-esp-3des-sha-1

CLI (Paris)

Preshared Key:

1. set ike gateway to_tokyo ip 201.22.3.14 main preshare hopsalong3 proposal pre-g2-3des-sha-1
2. set vpn paris_tokyo gateway to_tokyo tunnel proposal nopfs-esp-3des-sha-1

Certificate:

1. set ike gateway to_tokyo ip 201.22.3.14 main proposal rsa-g2-3des-sha
2. set vpn paris_tokyo gateway to_tokyo tunnel proposal nopfs-esp-3des-sha-1

Example: LAN-to-LAN VPN, Dynamic Peer

In this example, an IPSec tunnel securely connects the users on the Trusted LAN behind the NetScreen-10 to the mail server on the company's DMZ, behind the NetScreen-100. The NetScreen-100 has a static Untrusted IP address. The ISP for the NetScreen-10 assigns its Untrusted IP address dynamically via DHCP.

Although the NetScreen-10 is a DHCP client, to allow traffic to originate from the NetScreen-100 side of the tunnel, the NetScreen-10 cannot function as a DHCP server. Its Trusted LAN must contain static IP addresses that are known to the NetScreen-100 administrator, who can then add them to the NetScreen-100 address book for use in Access Policies to tunnel traffic to those destination addresses. After the NetScreen-10 establishes the tunnel, traffic through the tunnel can originate from either end[8].

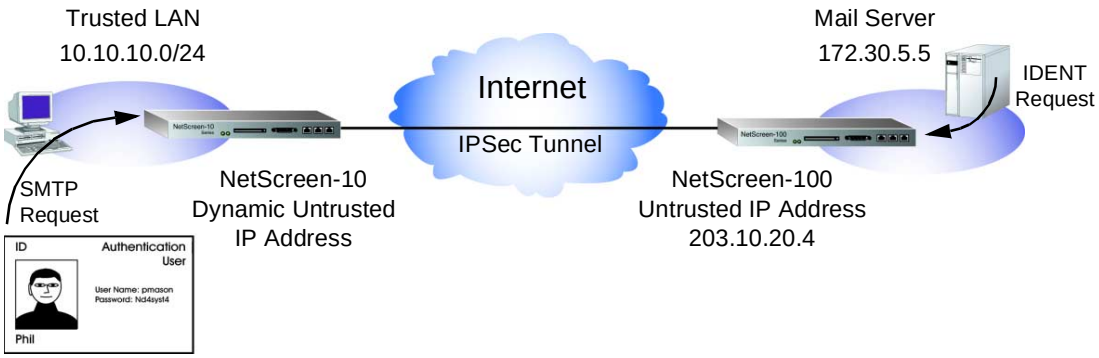

In this example, authentication user Phil (login name: pmason; password: Nd4syst4) wants to get his e-mail from the mail server at the company site. When he attempts to do so, he is authenticated twice: first, the NetScreen-10 authenticates him locally before allowing traffic from him through the tunnel; second, the mail server program authenticates him, sending the IDENT request through the tunnel.

> ***Note:*** *The mail server can send the IDENT request through the tunnel only because the NetScreen-100 administrator can set up an Access Policy directing that traffic through the tunnel to the 10.10.10.0/24 subnet.*

8. The use of dynamic peers, instead of VPN dialup users, offers the inclusion of NetScreen devices with dynamically assigned Untrusted IP addresses in a hub-and-spoke VPN architecture. For more information, see "Hub-and-Spoke VPNs" on page 6-47.

WebUI (NetScreen-10)

1. Users >> Users >> New User: Enter the following, and then click **OK**:
   User Name: pmason
   Authentication User: (select)
   Authentication Password: Nd4syst4
   Confirm Password: Nd4syst4
2. Address >> Untrusted: Enter the following, and then click **OK**:
   Name: Mail Server
   IP Address/Domain Name: 172.30.5.5
   Netmask: 255.255.255.255
3. VPN >> Gateway >> New Remote Gateway: Enter the following, and then click **OK**:
   Name: To_Mail
   Remote Gateway Static IP Address: (select)
   IP Address: 203.10.20.4
   Mode: Aggressive
   For preshared key:
   Phase 1 Proposal: pre-g2-3des-sha
   Preshare Key: 12345678
   For certificate:
   Phase 1 Proposal: rsa-g2-3des-sha
   Local ID: name@company.com
4. VPN >> AutoKey IKE >> New AutoKey IKE Entry: Enter the following, and then click **OK**:
   Name: branch_corp
   Remote Gateway Tunnel Name: To_Mail
   Phase 2 Proposal: nopfs-esp-3des-sha
5. Policy >> Outgoing >> New Policy: Enter the following, and then click **OK**:
   Source Address: Inside Any
   Destination Address: Mail Server
   Service: Mail
   Action: Tunnel
   VPN Tunnel: branch_corp
   Authentication: (select)

WebUI (NetScreen-100)

1. Address >> Untrusted: Enter the following, and then click **OK**:

   Name: branch office

   IP Address/Domain Name: 10.10.10.0

   Netmask: 255.255.255.0

2. VPN >> Gateway >> New Remote Gateway: Enter the following, and then click **OK**:

   Name: To_branch

   Remote Gateway Dynamic IP Address: (select)

   Peer ID: name@company.com

   Mode: Aggressive

   Phase 1 Proposal For preshared key:

   Phase 1 Proposal: pre-g2-3des-sha

   Preshare Key: 12345678

   For certificate:

   Phase 1 Proposal: rsa-g2-3des-sha

3. VPN >> AutoKey IKE >> New AutoKey IKE Entry: Enter the following, and then click **OK**:

   Name: corp_branch

   Remote Gateway Tunnel Name: To_branch

   Phase 2 Proposal: nopfs-esp-3des-sha

4. Policy >> Outgoing >> New Policy: Enter the following, and then click **OK**:

   Source Address: Mail Server

   Destination Address: branch office

   Service: Any

   Action: Tunnel

   VPN Tunnel: corp_branch

CLI (NetScreen-10)

1. set user pmason password Nd4syst4
2. set address untrust “mail server” 172.30.5.5 255.255.255.255

Preshared Key:

3. set ike gateway to_mail ip 203.10.20.4 aggressive preshare 12345678 proposal pre-g2-3des-sha-1
4. set vpn branch_corp gateway to_mail tunnel proposal nopfs-esp-3des-sha
5. set policy outgoing “inside any” “mail server” mail tunnel branch_corp auth

Certificate:

3. set ike gateway to_mail ip 203.10.20.4 aggressive local-id name@company.com proposal rsa-g2-3des-sha
4. set vpn branch_corp gateway to_mail tunnel proposal nopfs-esp-3des-sha
5. set policy outgoing “inside any” “mail server” mail tunnel branch_corp auth

CLI (NetScreen-100)

1. set address trust “mail server” 172.30.5.5 255.255.255.255
2. set address untrust “branch office” 10.10.10.0 255.255.255.0

Preshared Key:

3. set ike gateway to_branch dynamic name@company.com aggressive preshare 12345678 proposal pre-g2-3des-sha
4. set vpn corp_branch gateway to_branch tunnel proposal nopfs-esp-3des-sha
5. set policy outgoing “mail server” “branch office” any tunnel vpn corp_branch

Certificate:

3. set ike gateway to_branch dynamic name@company.com aggressive proposal rsa-g2-3des-sha
4. set vpn corp_branch gateway to_branch tunnel proposal nopfs-esp-3des-sha
5. set policy outgoing “mail server” “branch office” any tunnel vpn corp_branch

Dialup-to-LAN VPNs

NetScreen devices also support VPN dialup connections. You can configure a NetScreen security gateway with a static IP address to secure an IPSec tunnel with a NetScreen-Remote client or with another NetScreen device with a dynamic IP address, such as the NetScreen-5 (acting as a DHCP or PPPoE client) or the NetScreen-10 (acting as a DHCP client).

You can configure tunnels for VPN dialup users on a per-user basis or form users into a VPN dialup group for which you need only configure one tunnel. For more information on creating VPN dialup groups, see “Dialup User Groups” on page 4-22.

This section describes the procedures for setting up two types of Dialup-to-LAN VPNs:

- Dialup-to-LAN VPN, Manual Key tunnel
- Dialup-to-LAN VPN, AutoKey IKE tunnel (with a preshared secret or certificates)
- Dynamic Peer Dialup-to-LAN VPN, AutoKey IKE tunnel (with a preshared secret or certificates)

Example: Dialup-to-LAN VPN, Manual Key

In this example, a remote user (Wendy) needs to access a UNIX server at the corporate site. The tunnel uses 3DES for encryption and MD5 for authentication.

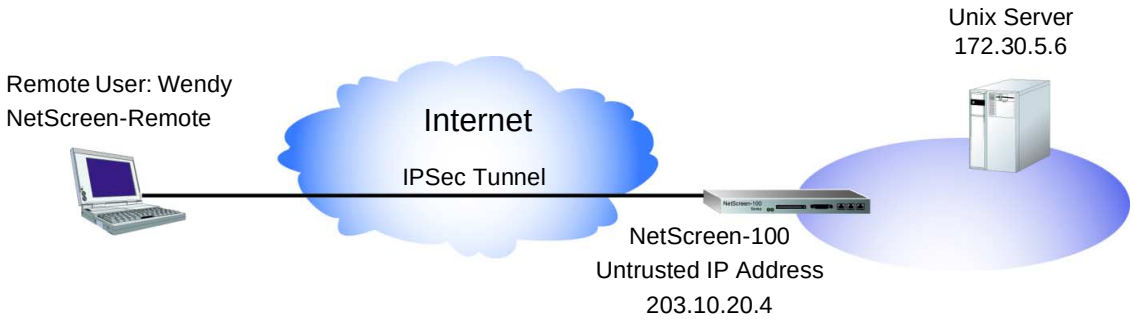

WebUI (NetScreen-100)

1. Address >> Trusted >> New Address: Enter the following, and then click **OK**:

   Name: UNIX

   IP Address/Domain Name: 172.30.5.6

   Netmask: 255.255.255.255

2. Users >> New User: Enter the following, and then click **OK**:

   User Name: Wendy

   VPN Dialup User – Manual Key Only: (select)

   Security Index: 4556 (Local), 4556 (Remote)

   ESP: (select)

   Encryption Algorithm: 3DES-CBC

   Generate Key by Password[9]: asdlk24234

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

Authentication Algorithm: SHA-1

Generate Key by Password: PNas134a

3. Policy >> Outgoing >> New Policy: Enter the following, and then click **OK**:

   Source Address: UNIX

   Destination Address: Dial-up VPN

   Service: ANY

   Action: Tunnel

   VPN Tunnel: Dialup User–Wendy

CLI

1. set address trust unix 172.30.5.6 255.255.255.255
2. set user wendy dialup 4556 4556 esp 3des password asdlk24234 auth sha-1 password PNas134a
3. set policy outgoing unix “dial-up vpn” any tunnel vpn-dialup wendy

NetScreen-Remote Security Policy Editor

1. Click **Options >> Secure >> Specified Connections**.
2. Click the **Add a new connection** button, and type **Unix** next to the new connection icon that appears.
3. Configure the connection options:

   Connection Security: Secure

   Remote Party ID Type: IP Address

   IP Address: 172.30.5.6

   Connect using Secure Gateway Tunnel: (select)

   ID Type: IP Address; 203.10.20.4

4. Click the **PLUS** symbol, located to the left of the unix icon, to expand the connection policy.
5. Click **Security Policy**, and select **Use Manual Keys**.
6. 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.

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

   Encapsulation Protocol (ESP): (select)

   Encrypt Alg: 3DES

   Hash Alg: SHA-1

   Encapsulation: Tunnel

8. Click **Inbound Keys**, and in the Security Parameters Index field, type **4556**.
9. Click **Enter Key**, enter the following[10], and then click **OK**:

   Choose key format: Binary

   ESP Encryption Key: dccbee96c7e546bc

   ESP Authentication Key: dccbe9e6c7e546bcb0b667794ab7290c

10. Click **Outbound Keys**, and in the Security Parameters Index field, type **4556**.
11. Click **Enter Key**, enter the following, and then click **OK**:

    Choose key format: Binary

    ESP Encryption Key: dccbee96c7e546bc

    ESP Authentication Key: dccbe9e6c7e546bcb0b667794ab7290c

12. Click the **Save** button.

10. These are the two generated keys that you copied after configuring the NetScreen device.

Example: Dialup-to-LAN VPN, AutoKey IKE (Preshared Key or Certificates)

In this example, an AutoKey IKE tunnel using either a preshared secret or a pair of certificates (one at each end of the tunnel[11]) provides the secure connection between Wendy and the UNIX server. The tunnel again uses ESP with 3DES encryption and SHA-1 authentication.

Setting up the AutoKey IKE tunnel using AutoKey IKE with either a preshared secret or certificates involves the following four steps:

1. Define Wendy as an IKE dynamic peer.
2. Configure the remote gateway and key exchange mode, and specify either a preshared secret or a certificate.
3. Configure an AutoKey IKE VPN.
4. Configure an Access Policy.

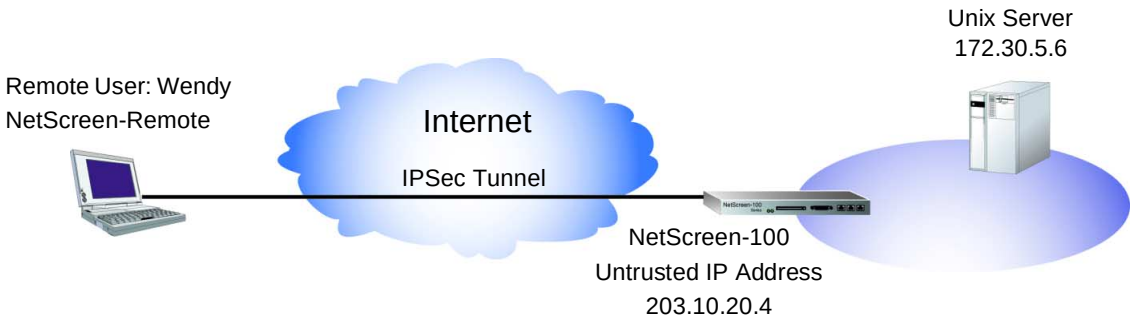

WebUI (NetScreen-100)

1. Address >> Trusted >> New Address: Enter the following, and then click **OK**:

   Name: UNIX

   IP Address/Domain Name: 172.30.5.6

   Netmask: 255.255.255.255

11. The preshared key is 12341234. It is assumed that both participants already have certificates. For more information about certificates, see “Certificates and CRLs” on page 6-17.

2. Users >> Users >> New User: Enter the following, and then click **OK**:

   User Name: Wendy

   IKE Dynamic Peer: (select)

   Identity: wparker@email.com

3. VPN >> Gateway >> New Remote Gateway: Enter the following, and then click **OK**:

   Name: Wendy_NSR

   Remote Gateway Dialup User: (select)

   User/Group: Dialup User–Wendy

   Mode: Aggressive

   Phase 1 Proposal For preshared key:

   Phase 1 Proposal: pre-g2-3des-sha

   Preshared Key: 12341234

   For certificate:

   Phase 1 Proposal: rsa-g2-3des-sha

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

   Name: Wendy_UNIX

   Remote Gateway Tunnel Name: Wendy_NSR

   Phase 2 Proposal: nopfs-esp-3des-sha

5. Policy >> Outgoing >> New Policy: Enter the following, and then click **OK**:

   Source Address: UNIX

   Destination Address: Dial-Up VPN

   Service: ANY

   Action: Tunnel

   VPN Tunnel: Wendy_UNIX

CLI

1. set address trust unix 172.30.5.6 255.255.255.255
2. set user wendy ike-id wparker@email.com

Preshared Key:

3. set ike gateway wendy_nsr dialup wendy aggressive preshare 12341234 proposal pre-g2-3des-sha-1
4. set vpn wendy_unix gateway wendy_nsr tunnel proposal nopfs-esp-3des-sha-1
5. set policy outgoing unix “dial-up vpn” any tunnel vpn wendy_unix

Certificates:

3. set ike gateway wendy_nsr dialup wendy aggressive proposal rsa-g2-3des-sha-1
4. set vpn wendy_unix gateway wendy_nsr tunnel proposal nopfs-esp-3des-sha-1
5. set policy outgoing unix “dial-up vpn” any tunnel vpn wendy_unix

NetScreen-Remote Security Policy Editor

1. Click **Options >> Secure >> Specified Connections**.
2. Click the **Add a new connection** button, and type **Unix** next to the new connection icon that appears.
3. Configure the connection options:

   Connection Security: Secure

   Remote Party ID Type: IP Address

   IP Address: 172.30.5.6

   Connect using Secure Gateway Tunnel: (select)

   ID Type: IP Address; 203.10.20.4
4. Click the **PLUS** symbol, located to the left of the unix icon, to expand the connection policy.

5. Click **My Identity**: Do either of the following:

   Click **Pre-shared Key** >> **Enter Key**: Type **12341234**, and then click **OK**.

   ID Type: (select **E-mail Address**), and type **wparker@email.com**.

   or

   Select a certificate from the Select Certificate drop-down list.

   ID Type: (select **E-mail Address**)[12]

6. Click the **Security Policy** icon, and select **Aggressive Mode**.
7. Click the **PLUS** symbol, located to the left of the Security Policy icon, and then the **PLUS** symbol to the left of Authentication (Phase 1) and Key Exchange (Phase 2) to expand the policy further.
8. Click **Authentication (Phase 1)** >> **Proposal 1**: Select the following Encryption and Data Integrity Algorithms:

   Encrypt Alg: Triple DES

   Hash Alg: SHA-1

   Key Group: Diffie-Hellman Group 2

9. Click **Key Exchange (Phase 2)** >> **Proposal 1**: Select the following IPSec Protocols:

   Encapsulation Protocol (ESP): (select)

   Encrypt Alg: 3DES

   Hash Alg: SHA-1

   Encapsulation: Tunnel

10. Click the **Save** button.

12. The e-mail address from the certificate appears in the identifier field automatically.

Example: Dialup-to-LAN VPN, Dynamic Peer

In this example, an IPSec tunnel securely connects the user behind the NetScreen-Remote to the mail server on the company's DMZ, behind the NetScreen-100. The NetScreen-100 has a static Untrusted IP address. The NetScreen-Remote client has a dynamically assigned external IP address and a static (virtual) internal IP address, which must be known to the NetScreen-100 administrator so that he can add it to the NetScreen-100 address book for use in Access Policies to tunnel traffic to that destination. After the NetScreen-Remote client establishes the tunnel, traffic through the tunnel can originate from either end.

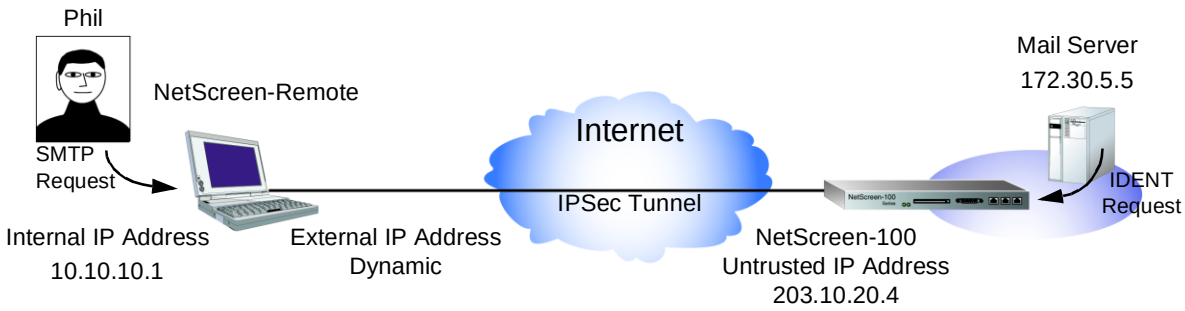

In this example, Phil wants to get his e-mail from the mail server at the company site. When he attempts to do so, he is authenticated by the mail server program which sends him an IDENT request through the tunnel.

> ***Note:*** *The mail server can send the IDENT request through the tunnel only because the NetScreen-100 administrator can set up an Access Policy directing that traffic through the tunnel to 10.10.10.1.*

WebUI (NetScreen-100)

1. Address >> Untrusted: Enter the following, and then click **OK**:

   Name: Phil

   IP Address/Domain Name: 10.10.10.1

   Netmask: 255.255.255.255

2. VPN >> Gateway >> New Remote Gateway: Enter the following, and then click **OK**:

   Name: To_Phil

   Remote Gateway Dynamic IP Address: (select)

   Peer ID: 10.10.10.1

   Mode: Aggressive

   Phase 1 Proposal For preshared key:

   Phase 1 Proposal: pre-g2-3des-sha

   Preshare Key: 12345678

   For certificate:

   Phase 1 Proposal: rsa-g2-3des-sha

3. VPN >> AutoKey IKE >> New AutoKey IKE Entry: Enter the following, and then click **OK**:

   Name: corp_Phil

   Remote Gateway Tunnel Name: To_Phil

   Phase 2 Proposal: nopfs-esp-3des-sha

4. Policy >> Outgoing >> New Policy: Enter the following, and then click **OK**:

   Source Address: Mail Server

   Destination Address: branch office

   Service: Any

   Action: Tunnel

   VPN Tunnel: corp_Phil

CLI (NetScreen-100)

1. set address trust “mail server” 172.30.5.5 255.255.255.255
2. set address untrust phil 10.10.10.1 255.255.255.255

Preshared Key:

3. set ike gateway to_phil dynamic 10.10.10.1 aggressive preshare 12345678 proposal pre-g2-3des-sha
4. set vpn corp_phil gateway to_phil tunnel proposal nopfs-esp-3des-sha
5. set policy outgoing “mail server” phil any tunnel vpn corp_branch

Certificate:

3. set ike gateway to_phil dynamic 10.10.10.1 aggressive proposal rsa-g2-3des-sha
4. set vpn corp_phil gateway to_phil tunnel proposal nopfs-esp-3des-sha
5. set policy outgoing “mail server” phil any tunnel vpn corp_phil

NetScreen-Remote

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

   Connection Security: Secure

   Remote Party ID Type: IP Address

   IP Address: 172.30.5.5

   Connect using Secure Gateway Tunnel: (select)

   ID Type: IP Address; 203.10.20.4
5. Click the **PLUS** symbol, located to the left of the unix icon, to expand the connection policy.

6. Click **My Identity**: Do either of the following:

   Click **Pre-shared Key** >> **Enter Key**: Type **12341234**, and then click **OK**.

   ID Type: Internal Network IP Address: Type **10.10.10.1**.

   or

   Select a certificate from the Select Certificate drop-down list.

   ID Type: (select **IP Address**); Internal Network IP Address: Type **10.10.10.1**.

7. Click the **Security Policy** icon, and select **Aggressive Mode**.
8. Click the **PLUS** symbol, located to the left of the Security Policy icon, and then the **PLUS** symbol to the left of Authentication (Phase 1) and Key Exchange (Phase 2) to expand the policy further.
9. Click **Authentication (Phase 1)** >> **Proposal 1**: Select the following Encryption and Data Integrity Algorithms:

   Encrypt Alg: Triple DES

   Hash Alg: SHA-1

   Key Group: Diffie-Hellman Group 2

10. Click **Key Exchange (Phase 2)** >> **Proposal 1**: Select the following IPSec Protocols:

    Encapsulation Protocol (ESP): (select)

    Encrypt Alg: 3DES

    Hash Alg: SHA-1

    Encapsulation: Tunnel

11. Click the **Save** button.

Hub-and-Spoke VPNs

Initially, most VPNs are deployed in a configuration allowing remote users to reach a central site to access key resources. However, as networks grow, individuals might need to create tunnels to more than just the central network site.

A hub-and-spoke design addresses this by allowing a central, powerful site to act as a hub with a series of VPN tunnels branching out from it like spokes to perimeter sites. Consequently, each remote site need only maintain a single tunnel for all VPN communications.

By establishing a VPN connection to the hub site, a perimeter site makes secure VPN connections with every other perimeter site—without needing to create any additional VPN tunnels. That is, the perimeter site uses the existing tunnels that all converge at the hub.

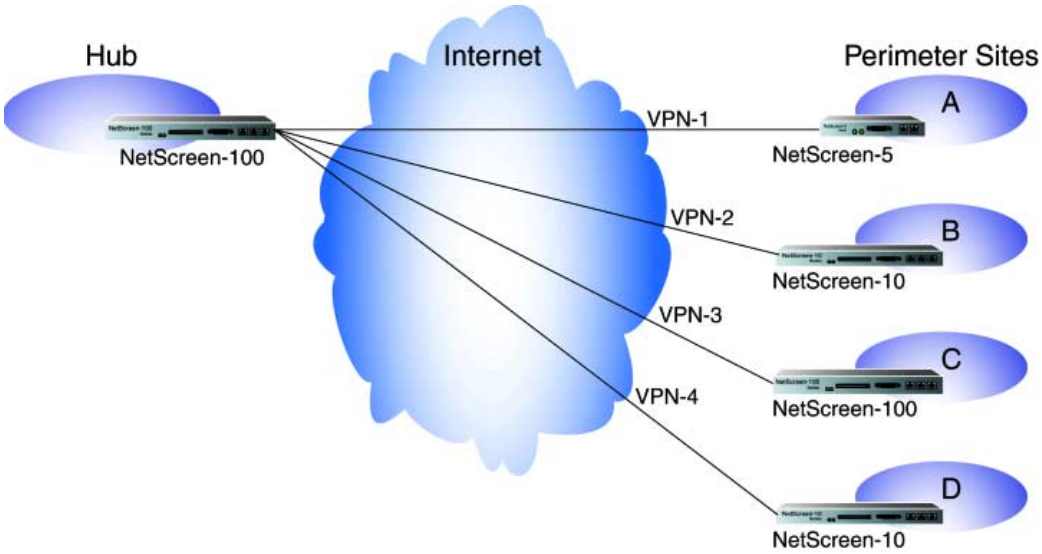

A few benefits of hub-and-spoke VPNs:

- You can conserve the number of VPNs you need to create. For example, perimeter site A can link to the hub, and perimeter sites B, C, D..., but A only has to set up one VPN tunnel. Especially for NetScreen-5 users, who can use a maximum of just ten VPN tunnels concurrently, applying the hub-and-spoke method dramatically increases their VPN options and capabilities.

- The administrator at the hub device can completely control VPN traffic between perimeter sites. For example,
  - He or she might permit only HTTP traffic to flow from sites A to B, but allow any kind of traffic to flow from B to A.
  - He or she can allow traffic originating from A to reach C, but deny traffic originating from C to reach A.
  - He or she can allow a specific host at A to reach the entire D network, while allowing only a specific host at D to reach a different host at A.
- The administrator at the hub device can completely control outbound traffic from all perimeter networks. At each perimeter site, there must first be an Access Policy that funnels all outbound traffic through the spoke VPNs to the hub; for example: set policy outgoing “inside any” “outside any” any encrypt VPN-x (where “x” defines the specific VPN tunnel from each perimeter site). At the hub, the administrator can control Internet access, allowing certain kinds of traffic (such as HTTP only), performing URL blocking on undesirable Web sites, and so on.
- Regional hubs can be used and interconnected via spoke tunnels, allowing spoke sites in one region to reach spoke sites in another.

Setting Up a Hub-and-Spoke VPN

Assuming that VPN tunnels from each perimeter site to the hub are already in place, creating a hub-and-spoke VPN involves the following three steps:

1. Defining addresses
2. Creating Access Policies
3. Entering fixed routes in the route table of the hub device

Defining Addresses

Before you can create Access Policies to Untrusted locations, you must first define the addresses for those locations in the address book. Also, because of the unique position of the device acting as the hub, it must define some addresses that exist on its Untrusted side as if they were Trusted addresses. This is substantially simplified by carefully designing the subnets of the spoke networks and that of the hub network within two subgroups of the overall network. Using subnet masking, you can include or exclude networks. For example, in a simple hub-and-spoke design consisting of a hub and two perimeter sites A and B, you can address the sites as follows:

- If the Trusted networks for A, B, and the hub are defined as follows:

  A = 10.10.1.0/24

  B = 10.10.3.0/24

  Hub = 10.2.2.0/24

  then addresses at A, B, and the hub are seen as being in three separate networks.

- If the Trusted networks for A, B, and the hub are defined as follows:

  A = 10.10.0.0/16

  B = 10.10.0.0/16

  Hub = 10.2.2.0/24

  then A and B are seen as being in the same network, apart from the hub's network.

Note: *By subnetting all perimeter sites within the 10.10.0.0/16 range, you do not need to create individual routes to each subnet in the hub's route table; one statement suffices for routing to all perimeter sites.*

- If the Trusted networks for A, B, and the hub are defined as follows:

  A = 10.0.0.0/8

  B = 10.0.0.0/8

  Hub = 10.0.0.0/8

  then addresses at A, B, and the hub are seen as being in the same network.

Note: *By putting all networks within the 10.0.0.0/8 range, each perimeter network needs to define only one Untrusted address and create only one outgoing Access Policy to permit access to all other networks in the hub-and-spoke domain.*

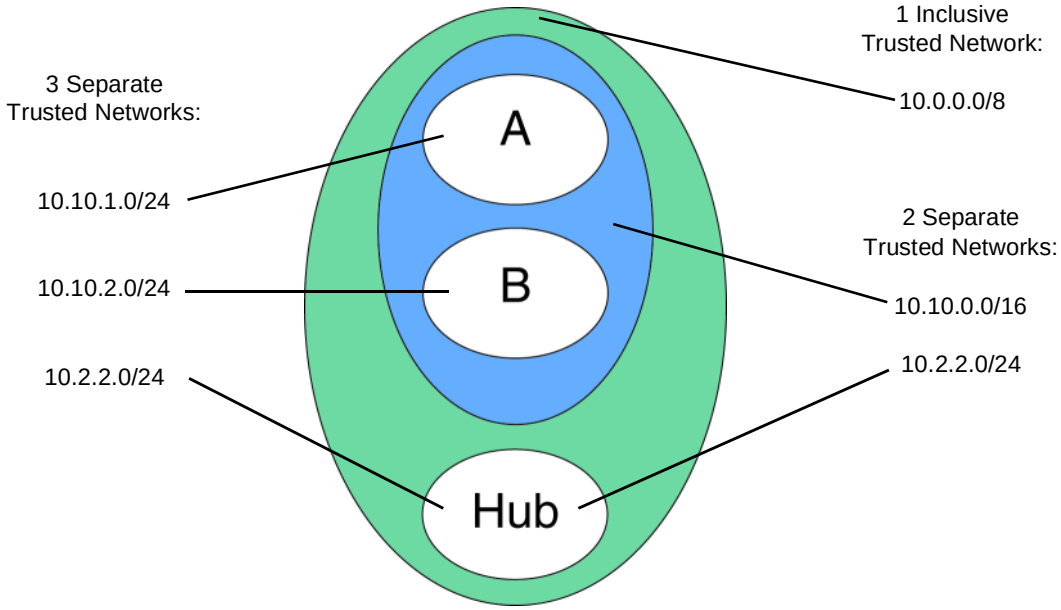

Because of the addressing demands of a hub-and-spoke VPN, the Trusted hosts on all participating sites must have fixed IP addresses[13].

Creating Access Policies

You must create Access Policies to encrypt/decrypt traffic and permit it to cross the firewalls. Because bidirectional traffic is assumed for VPNs, the outgoing Access Policy that you create functions both as an outgoing and incoming policy.

Routing at the Hub

The final step is to define the route that the NetScreen device at the hub uses when receiving incoming traffic bound for the Untrusted interface. You must define a static route in the route table so that the device knows to send packets it receives for perimeter networks to the Untrusted interface.

13. Although Trusted hosts must have fixed IP addresses, the Untrusted interfaces can be assigned dynamically.

Hub-and-Spoke VPN Packet Flow

The packet flow from a host (Host A) on the Trusted network behind the NetScreen-5 (A) via the NetScreen (hub) to a server (Server B) behind the NetScreen-10 (B) proceeds as follows:

1. Receiving an outbound packet destined for 10.10.3.1, A checks its Access Control List (ACL) for an Access Policy permitting outbound encrypted traffic from Host A (10.10.1.1).
2. Finding permission granted, it encrypts the packet and sends it through tunnel VPN-1.
3. The hub receives the encrypted packet and checks the security parameter index (SPI) number on the IP packet header to locate the security association (SA) required to decrypt the packet.
4. After decrypting the packet, the hub checks the associated Access Policy associated with the SA to verify that this traffic is allowed to cross the firewall to the Trusted side.
5. Noting that it is allowed, the hub then refers to the route table to determine where to send the packet.
6. Learning that it is destined for the Untrusted interface, the hub then consults the ACL for an Access Policy permitting this.
7. Finding one that also specifies that the packet must be encrypted and sent through VPN-2 to B, the hub encrypts the packet and sends it to B through VPN-2.
8. B, after decrypting and inspecting the packet, forwards it to the server on its Trusted interface.

The packet flow proceeds similarly in the reverse direction.

Example: Hub-and-Spoke VPN

Using VPN-1 and VPN-2, secure VPN communications can be established between hosts on network A (protected by a NetScreen-5) and B (protected by a NetScreen-10) by routing the VPN traffic through the hub (protected by a NetScreen-100).

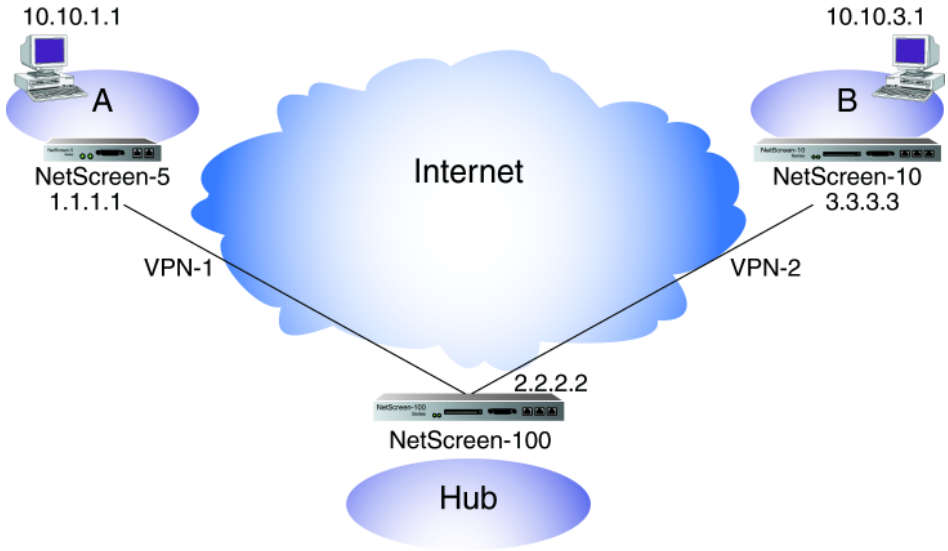

WebUI: NetScreen-5 (Network A)

1. Address >> Untrusted >> New Address: Enter the following, and then click **OK**:

   Address Name: All networks
   IP Address/Domain Name: 10.0.0.0
   Netmask: 255.0.0.0
   Comment: Encompasses all addresses

2. Policy >> Outgoing >> New Policy: Enter the following, and then click **OK**:

   Source Address: Network-A
   Destination Address: All networks
   Service: Any
   Action: Tunnel
   VPN Tunnel: VPN-1

NetScreen-10 (Network B)

1. Address >> Untrusted >> New Address: Enter the following, and then click **OK**:

   Address Name: All networks
   IP Address/Domain Name: 10.0.0.0
   Netmask: 255.0.0.0
   Comment: Encompasses all addresses

2. Policy >> Outgoing >> New Policy: Enter the following, and then click **OK**:

   Source Address: Network-B
   Destination Address: All networks
   Service: Any
   Action: Tunnel
   VPN Tunnel: VPN-2

NetScreen-100 (The Hub)

1. Address >> Untrusted >> New Address: Enter the following, and then click **OK**:

   Address Name: Network A
   IP Address/Domain Name: 10.10.1.0
   Netmask: 255.255.255.0
   Comment: A's addresses

2. Address >> Untrusted >> New Address: Enter the following, and then click **OK**:

   Address Name: Network C
   IP Address/Domain Name: 10.10.3.0
   Netmask: 255.255.255.0
   Comment: C's addresses

3. Address >> Trusted >> New Address: Enter the following, and then click **OK**:

   Address Name: All Networks
   IP Address/Domain Name: 10.0.0.0
   Netmask: 255.0.0.0
   Comment: Encompasses all addresses

4. Policy >> Outgoing >> New Policy: Enter the following, and then click **OK**:
   Source Address: All Networks
   Destination Address: Network A
   Service: Any
   Action: Tunnel
   VPN Tunnel: VPN-1
5. Policy >> Outgoing >> New Policy: Enter the following, and then click **OK**:
   Source Address: All Networks
   Destination Address: Network C
   Service: Any
   Action: Tunnel
   VPN Tunnel: VPN-2
6. Configure >> Route Table >> New Entry:
   Network Address: 10.10.0.0
   Netmask: 255.255.0.0
   Interface: Untrusted

CLI

1. ns5-> set address untrust “all networks” 10.0.0.0 255.0.0.0 “encompasses all addresses”
2. ns5-> set policy outgoing “inside any” “all networks” any tunnel vpn-tunnel VPN-1

1. ns10-> set address untrust all 10.0.0.0 255.0.0.0 “encompasses all addresses”
2. ns10-> set policy outgoing “inside any” “all networks” any tunnel vpn-tunnel VPN-2

1. ns100-> set address untrust “network A” 10.10.1.0 255.255.255.0 “A’s addresses”
2. ns100-> set address untrust “network C” 10.10.3.0 255.255.255.0 “C’s addresses”
3. ns100-> set address trust “all networks” 10.0.0.0 255.0.0.0 “encompasses all addresses”
4. ns100-> set route 10.10.0.0 255.255.0.0 interface untrust

Traffic Shaping 7

This chapter discusses the various ways you can use your NetScreen device to manage limited bandwidth without compromising quality and availability of the network to all of your users.

The topics discussed include:

APPLYING TRAFFIC SHAPING

Traffic shaping is the allocation of the appropriate amount of network bandwidth to every user and application on an interface. The appropriate amount of bandwidth is defined as cost-effective carrying capacity at a guaranteed Quality of Service.

You can use NetScreen-5, -10, and -100 devices to shape traffic by creating policies and by applying appropriate rate controls to each class of traffic going through your NetScreen devices.[1]

Managing Bandwidth at the Access Policy Level

To classify traffic, you create a policy which specifies the amount of guaranteed bandwidth, the maximum bandwidth, and the priority for each class of traffic.

The physical bandwidth of every interface is allocated to the guaranteed bandwidth parameter for all policies. If there is any bandwidth left over, it is sharable by any other traffic.

In other words, each policy gets its guaranteed bandwidth and shares whatever is left over on a priority basis (up to the limit of its maximum bandwidth specification).

---

1. Traffic shaping is not supported on the NetScreen-1000.

The traffic shaping function applies to traffic from all policies. If you turn off traffic shaping for a policy, but traffic shaping is still turned on for other policies, the system applies a default traffic shaping policy to this policy assigning 0 guaranteed bandwidth, unlimited maximum bandwidth, and a priority of 7 (the lowest priority setting). If you do not want the system to assign this default traffic shaping policy to policies for which you have turned off traffic shaping, then turn off traffic shaping system wide.

You can set traffic shaping to automatic. This allows the system to turn on traffic shaping when a policy requires it, and turn off traffic shaping when policies do not require it.

Example: Traffic Shaping

This example shows how you could partition 45Mbps of bandwidth on a T3 interface among three departments on the same subnet.

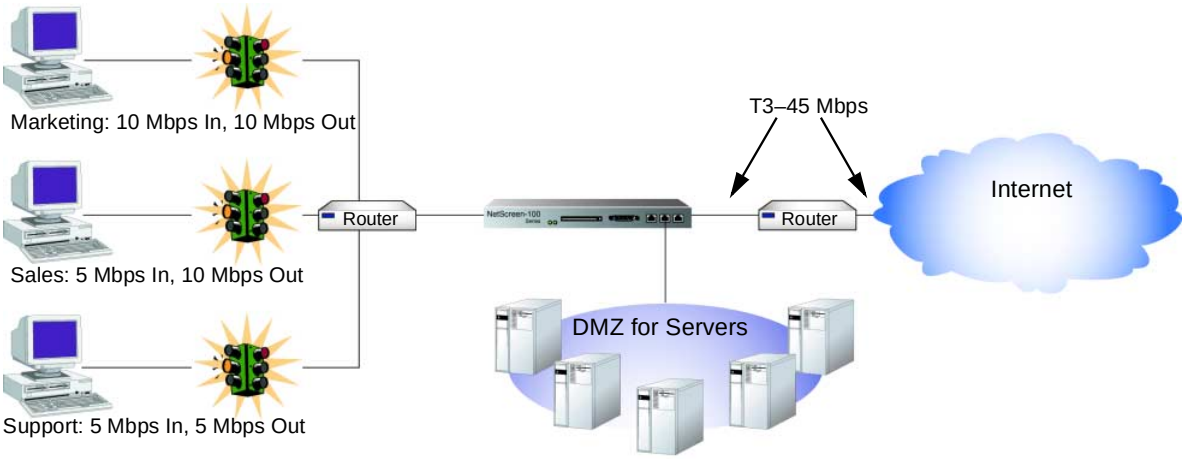

Web UI

1. Interface >> Trusted >> Edit: Enter the following, and then click **Save**:
   Traffic Bandwidth: 45000
2. Interface >> Untrusted >> Edit: Enter the following, and then click **Save**:
   Traffic Bandwidth: 45000
3. Policy >> Outgoing >> New Policy: Enter the following, and then click **OK**:
   Name: Marketing Traffic Shaping Policy
   Source Address: Marketing
   Destination Address: Outside Any

Service: Any
Action: Permit
VPN Tunnel: None[2]
Logging: Enable
Counting: Enable
Alarm Threshold: Accept Defaults
Schedule: Accept Defaults
Guaranteed Bandwidth: 10000
Maximum Bandwidth: 15000

4. Policy >> Outgoing >> New Policy: Enter the following, and then click **OK**:
   Name: Sales Traffic Shaping Policy
   Source Address: Sales
   Destination Address: Outside Any
   Service: Any
   Action: Permit
   VPN Tunnel: None
   Logging: Enable
   Counting: Enable
   Alarm Threshold: Accept Defaults
   Schedule: Accept Defaults
   Guaranteed Bandwidth: 10000
   Maximum Bandwidth: 10000
5. Policy >> Outgoing >> New Policy: Enter the following, and then click **OK**:
   Name: Support Traffic Shaping Policy
   Source Address: Support
   Destination Address: Outside Any
   Service: Any
   Action: Permit
   VPN Tunnel: None

2. You can also do traffic shaping on VPN tunnels.

Logging: Enable
Counting: Enable
Alarm Threshold: Accept Defaults
Schedule: Accept Defaults
Guaranteed Bandwidth: 5000.
Maximum Bandwidth: 10000.

6. Policy >> Incoming >> New Policy: Enter the following, and then click **OK**:
   Name: Allow Incoming Access to Marketing
   Source Address: Outside Any
   Destination Address: Marketing
   Service: Any
   Action: Permit
   VPN Tunnel: None
   Logging: Enable
   Counting: Enable
   Alarm Threshold: Accept Defaults
   Schedule: Accept Defaults
   Guaranteed Bandwidth: 10000.
   Maximum Bandwidth: 10000.
7. Policy >> Incoming >> New Policy: Enter the following, and then click **OK**:
   Name: Allow Incoming Access to Sales
   Source Address: Outside Any
   Destination Address: Sales
   Service: Any
   Action: Permit
   VPN Tunnel: None
   Logging: Enable
   Counting: Enable
   Alarm Threshold: Accept Defaults
   Schedule: Accept Defaults
   Guaranteed Bandwidth: 5000.
   Maximum Bandwidth: 10000.

8. Policy >> Incoming >> New Policy: Enter the following, and then click **OK**:

   Name: Allow Incoming Access to Support
   Source Address: Outside Any
   Destination Address: Support
   Service: Any
   Action: Permit
   VPN Tunnel: None
   Logging: Enable
   Counting: Enable
   Alarm Threshold: Accept Defaults
   Schedule: Accept Defaults
   Guaranteed Bandwidth: 5000.
   Maximum Bandwidth: 5000.

CLI

To enable traffic shaping by Access Policy:

1. set interface trust bandwidth 45000
2. set interface untrust bandwidth 45000
3. set policy outgoing marketing “outside any” any permit log count traffic gbw 10000 mbw 15000
4. set policy outgoing sales “outside any” any permit log count traffic gbw 10000 mbw 10000
5. set policy outgoing support “outside any” any permit log count traffic gbw 5000 mbw 10000
6. set policy incoming marketing “outside any” any permit log count traffic gbw 10000 mbw 10000
7. set policy incoming sales “outside any” any permit log count traffic gbw 5000 mbw 10000
8. set policy incoming support “outside any” any permit log count traffic gbw 5000 mbw 5000

or

To enable traffic shaping at the interface:

1. set interface trust bandwidth 45000
2. set interface untrust bandwidth 45000
3. set traffic-shaping mode auto

SETTING SERVICE PRIORITIES

The traffic shaping feature supported on NetScreen devices allows you to perform priority queuing on the bandwidth that is not allocated to guaranteed bandwidth, or is guaranteed but not used.

Priority queuing is a feature that lets all your users and applications have access to available bandwidth as they need it, while ensuring that important traffic can get through, if necessary at the expense of less important traffic.

Queuing allows NetScreen to buffer traffic in up to eight different priority queues. These seven queues are:

- High priority
- 2nd priority
- 3rd priority
- 4th priority
- 5th priority
- 6th priority
- 7th priority
- Low priority (default)

The priority setting for a policy means that the bandwidth not already guaranteed to other policies is queued on the basis of High priority first and Low priority last. Policies with the same priority setting compete for bandwidth in a Round Robin fashion. The NetScreen device processes all of the traffic from all of the policies with High priority before processing any traffic from policies with the next lower priority setting, and so on, until all traffic requests have been processed. If traffic requests exceed available bandwidth, the lowest priority traffic is dropped.

Caution *Be careful not to allocate more bandwidth than the interface can support. The policy configuration process does not prevent you from creating unsupported policy configurations, but you can lose data if the guaranteed bandwidth on contending policies surpasses the traffic bandwidth set on the interface.*

If you do not allocate any guaranteed bandwidth, then you can use priority queuing to manage all of traffic on your network. That is, all High priority traffic is sent before any 2nd priority traffic is sent, and so on. The NetScreen device processes Low priority traffic only after all other traffic has been processed.

Example: Priority Queuing

In this example, Marketing is guaranteed 3 Mbps of inbound bandwidth and 5 Mbps of outbound bandwidth and can use up to 10 Mbps of the total 45 Mbps of available bandwidth. Sales is guaranteed 7 Mbps of inbound bandwidth and 5 Mbps of outbound bandwidth and can use up to 15 Mbps of the total 45 Mbps of available bandwidth. Support is guaranteed 10 Mbps of inbound bandwidth and 10 Mbps of outbound bandwidth and can use up to 20 Mbps of the total 45 Mbps of available bandwidth. If all three computers send traffic through the router to the NetScreen firewall, the NetScreen device must allocate 40 Mbps of bandwidth to fulfill the guaranteed Access Policy requirements. Of the 5 Mbps of available bandwidth remaining, 3 Mbps of available bandwidth goes to Sales to fulfill its maximum bandwidth allocation. Marketing traffic, having the lowest priority, uses the 8 Mbps of guaranteed traffic, and gets the remaining 2 Mbps towards its maximum bandwidth allocation.

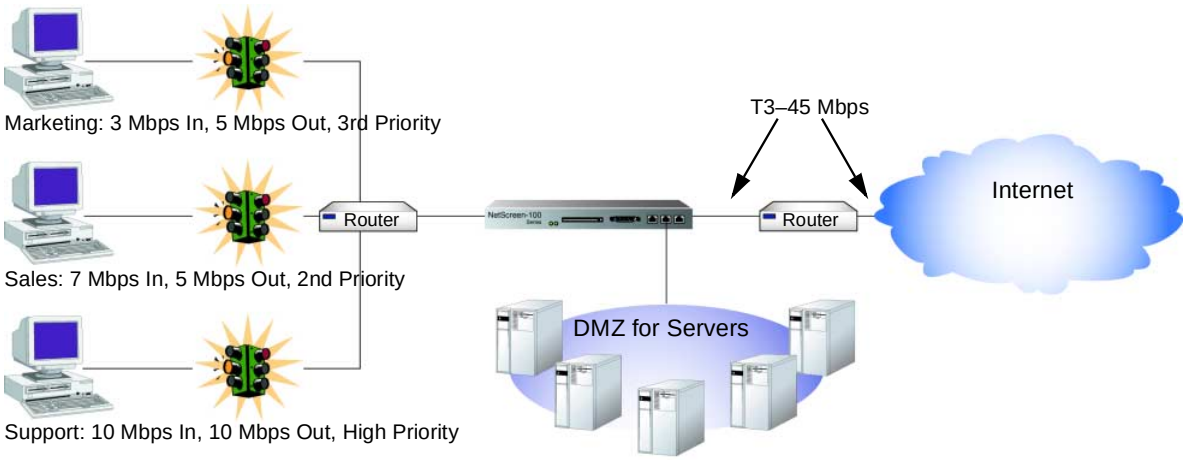

Web UI

1. Policy >> Outgoing >> New Policy: Enter the following, and then click **OK**:

   Name: Priority Policy 1

   Source Address: Marketing

   Destination Address: Outside Any

   Service: Any

   Action: Permit

   VPN Tunnel: None[3]

   Logging: Enable

   Counting: Enable

Alarm Threshold: Accept Defaults
Schedule: Accept Defaults
Guaranteed Bandwidth: 5000.
Maximum Bandwidth: 5000.
Traffic Priority: 3rd priority
DS Codepoint Marking: Enable

2. Policy >> Outgoing >> New Policy: Enter the following, and then click **OK**:

   Name: Priority Policy 2
   Source Address: Sales
   Destination Address: Outside Any
   Service: Any
   Action: Permit
   VPN Tunnel: None
   Logging: Enable
   Counting: Enable
   Alarm Threshold: Accept Defaults
   Schedule: Accept Defaults
   Guaranteed Bandwidth: 5000.
   Maximum Bandwidth: 5000.
   Traffic Priority: 2nd priority
   DS Codepoint Marking[4]: Enable

3. Policy >> Outgoing >> New Policy: Enter the following, and then click **OK**:

   Name: Priority Policy 3
   Source Address: Support
   Destination Address: Outside Any
   Service: Any
   Action: Permit

3. You can also do traffic shaping on VPN tunnels.

4. Differentiated Services (DiffServ) is a system for tagging (or "marking") traffic at a position within a hierarchy of priority. DS Codepoint Marking maps the NetScreen priority level of the Access Policy to the first three bits of codepoint in the DS field in the IP packet header. For more information about DS Codepoint Marking, see "Traffic Shaping" on page 5-5.

VPN Tunnel: None
Logging: Enable
Counting: Enable
Alarm Threshold: Accept Defaults
Schedule: Accept Defaults
Guaranteed Bandwidth: 10000.
Maximum Bandwidth: 10000.
Traffic Priority: High priority
DS Codepoint Marking: Enable

4. >>Policy >> Incoming >> New Policy: Enter the following, and then click **OK**:
   Name: Priority Policy 1
   Source Address: Marketing
   Destination Address: Outside Any
   Service: Any
   Action: Permit
   VPN Tunnel: None[5]
   Logging: Enable
   Counting: Enable
   Alarm Threshold: Accept Defaults
   Schedule: Accept Defaults
   Guaranteed Bandwidth: 3000.
   Maximum Bandwidth: 5000.
   Traffic Priority: 3rd priority
   DS Codepoint Marking: Enable
5. Policy >> Incoming >> New Policy: Enter the following, and then click **OK**:
   Name: Priority Policy 2
   Source Address: Sales
   Destination Address: Outside Any
   Service: Any
   Action: Permit

5. You can also do traffic shaping on VPN tunnels.

VPN Tunnel: None
Logging: Enable
Counting: Enable
Alarm Threshold: Accept Defaults
Schedule: Accept Defaults
Guaranteed Bandwidth: 7000.
Maximum Bandwidth: 10000.
Traffic Priority: 2nd priority
DS Codepoint Marking: Enable

6. Policy >> Incoming >> New Policy: Enter the following, and then click **OK**:

Name: Priority Policy 3
Source Address: Support
Destination Address: Outside Any
Service: Any
Action: Permit
VPN Tunnel: None
Logging: Enable
Counting: Enable
Alarm Threshold: Accept Defaults
Schedule: Accept Defaults
Guaranteed Bandwidth: 10000.
Maximum Bandwidth: 10000.
Traffic Priority: High priority
DS Codepoint Marking: Enable

CLI

1. set interface trust bandwidth 45000
2. set interface untrust bandwidth 45000
3. set policy outgoing marketing “outside any” any permit log count traffic gbw 5000 prio 3 mbw 5000 dscp enable
4. set policy outgoing sales “outside any” any permit log count traffic gbw 5000 prio 2 mbw 5000 dscp enable
5. set policy outgoing support “outside any” any permit log count traffic gbw 10000 prio High mbw 10000[6] dscp enable
6. set policy incoming marketing “outside any” any permit log count traffic gbw 3000 prio 3 mbw 5000 dscp enable
7. set policy incoming sales “outside any” any permit log count traffic gbw 7000 prio 2 mbw 10000 dscp enable
8. set policy incoming support “outside any” any permit log count traffic gbw 10000 prio High mbw 10000[7] dscp enable

6. In the outgoing policy examples in “Managing Bandwidth at the Access Policy Level” on page 7-1, not specifying a priority setting automatically defaults the policy priority setting to 7th priority.

7. In the outgoing policy examples in “Managing Bandwidth at the Access Policy Level” on page 7-1, not specifying a priority setting automatically defaults the policy priority setting to 7th priority.

LOAD BALANCING

Balancing the loads on NetScreen-100 servers allows traffic intended for a server (one IP address on the Untrusted network) to be mapped to several different IP addresses on the DMZ or Trusted network. There can be up to eight servers in one load balancing group, which appears as a single virtual server.

> ***Note:*** *You must configure the Virtual IP servers before you can configure the load balancing feature. The NetScreen-100 checks to see if the servers are up and responding by pinging each server.*

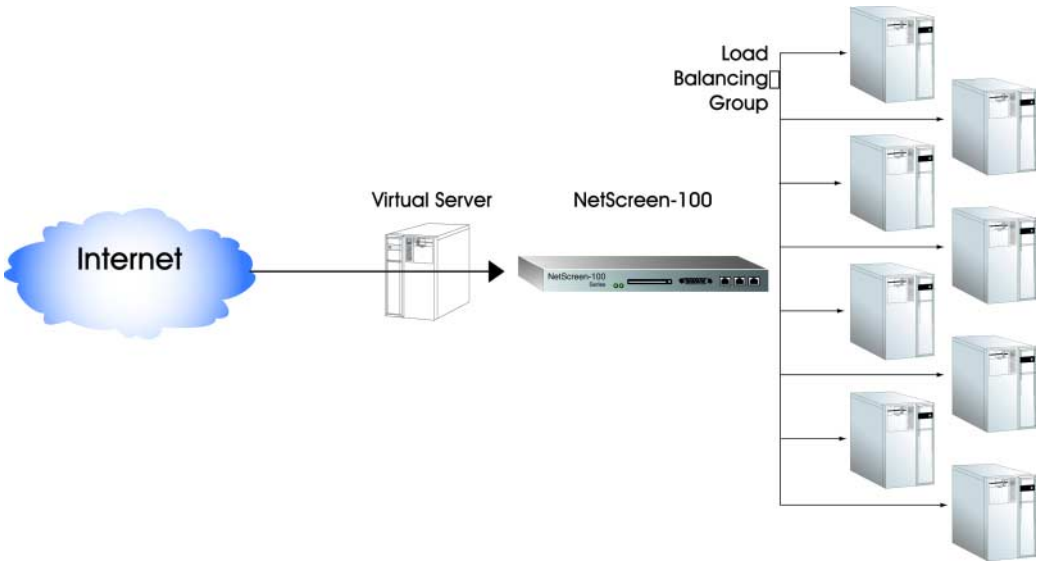

You can balance loads in one of the following ways:

Round Robin—Each new connection request is sent to each physical server in turn. Over time, each server gets the same amount of connection requests, although connection time may vary. This is one of the most commonly used load balancing methods when the servers are approximately equal in their capacity.

In a round robin scenario, there is a list of servers (at least two) capable of handling the request. The first request comes in and is handled by the first server that is on the list. The second request is handled by the second server on the list. This continues until the last server on the list is used. The next request is sent to the first server again. This rotation continues indefinitely.

However, the Load Balancing algorithm dictates specific characteristics of what constitutes the next session (determined by stickiness, session timeouts, and the like). For example, if the next request comes from the same client as the previous request, it is very likely that the request will go to the same server as the previous request.

Weighted Round Robin—Each server is given a static weighting function based on the capacity of each server. Servers are presented connection requests in proportion to their weighting relative to the total capacity of the system.

Least Connections—The number of active connections supported by each server is tracked. As new connections are received, they are forwarded to the server with the fewest active connections.

Weighted Least Connections—Each server is given a static weighting function based on the capacity of each server. The idea is that servers with greater inherent capacity should support a larger number of active connections.

Since the server-cluster appears as a single entity, users only need a single IP address or URL to access the services.

Weighted Round Robin Example

As an example, suppose you have three servers: A at 10.200.100.1, B at 10.200.100.2, and C at 10.200.100.3. They are assigned the following weights:

A = 10, B = 20, C = 30

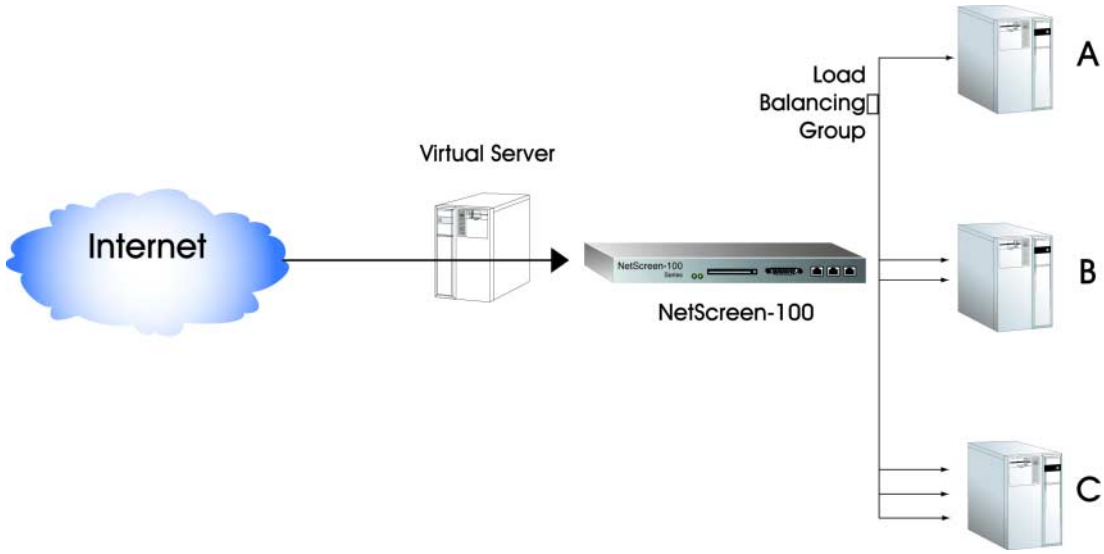

Using the Weighted Round Robin method, server A receives 1/6th of the incoming connection requests, Server B receives 1/3rd of the incoming connection requests, and Server C receives 1/2 of the incoming connection requests. This means the first 10 connections that come in are assigned to server A. The next 20 are assigned to server B. The next 30 are assigned to server C. Over many connections, this distributes the connections across all servers according to the ratios defined by the weights.

Web UI

Virtual IP >> Virtual IP tab >> New Service: Enter the following, and then click **OK**:

Virtual IP: 10.200.100.0

Virtual Port: 21

Service: FTP

Load Balance: Weighted Round Robin

Server IP: 10.200.100.1
10.200.100.2
10.200.100.3

Server Weight:
10
20
30

CLI

set vip 10.200.100.0 21 ftp weighted-round-robin 10.200.100.1/10

set vip 10.200.100.0 21 ftp weighted-round-robin 10.200.100.2/20

set vip 10.200.100.0 21 ftp weighted-round-robin 10.200.100.3/30

Weighted Least Conns Example

In this algorithm, the total sum of the weights should add up to 100 (as in 100%). Suppose we have two servers, A and B with the following weights:

A=80, B=20

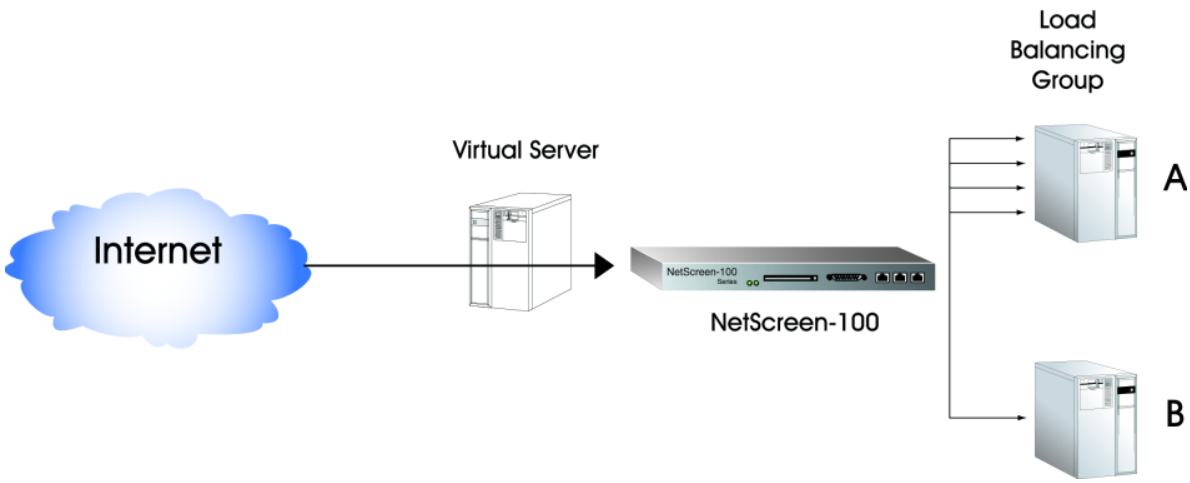

This means that server A can handle four connections for every one connection of server B. So, for example, if server A has two connections and server B has none, the next connection still goes to server A. The next connection also goes to server A. Once server A gets four connections, then server B gets one. This continues indefinitely.

When server A has 16 connections and server B has 3 connections, Server B gets the next connection.

Web UI

Virtual IP >> Virtual IP tab >> New Service: Enter the following, and then click **OK**:

Virtual IP: 10.200.100.0

Virtual Port: 21

Service: FTP

Load Balance: Weighted Least Conns

Server IP: 10.200.100.1
10.200.100.2

Server Weight:
80
20

CLI

1. set vip 10.200.100.0 21 ftp weighted-least-conns 10.200.100.1/80
2. set vip 10.200.100.0 21 ftp weighted-least-conns 10.200.100.2/20

High Availability

8

This chapter explains the following aspects of setting up two or more NetScreen-100 devices or two or more NetScreen-1000 devices for high availability (HA):

Note: *You can only set NetScreen-100 devices for HA with other NetScreen-100 devices, and NetScreen-1000 devices with other NetScreen-1000 devices.*

By cabling together two or more NetScreen-100 devices or two or more NetScreen-1000 devices, you can configure the units for HA. If one unit fails, the second unit can assume its functions with no service interruption[1], all sessions and Security Associations (SAs) are maintained so that VPN traffic is not interrupted, and because configurations and Access Policies are updated from the Master unit to the Slave unit(s), there is no loss of enforcement capability.

In each redundant group, or cluster, of NetScreen-100 or NetScreen-1000 devices, one device is the Master and the others are the Slaves. You can configure a specific unit as the Master, or you can allow the system to negotiate a Master dynamically based on the lowest Media Access Control (MAC) address.

If a Master unit detects a Slave unit failure, it generates a system alarm. If the Master unit fails, the remaining Slave units elect a new Master unit to which traffic is seamlessly redirected. Furthermore, during a fail-over event, NetScreen devices also notify the adjacent networking devices upstream and downstream so that the entire network can quickly converge to the new data path. When changing state between Master and Slave, both NetScreen units generate alarms, which can take the form of SNMP vendor-specific traps, e-mail alerts from the Master unit, or any other configured logging mechanism.

1. Because of special requirements, sessions generated from H.323-related applications (such as NetMeeting) are not maintained on the Slave unit.

You can enable authentication and encryption for HA communication involving configuration and keying material and control protocol material between the Master and Slave units to increase security.

> ***Note:*** *The session sync material is not encrypted.*

To ensure the integrity of the data paths between the NetScreen interfaces and the adjacent networking devices, you can configure and enable path monitoring, in which each device in the redundant group continually monitors the accessibility of specified devices on connected networks. Each Slave unit continually compares its ability to access these addresses with that of the Master unit. Should a Slave unit achieve better results than the Master, it then becomes the Master.

To implement HA configuration, you need to perform the following steps:

1. Cable the NetScreen-100 or NetScreen-1000 units together, as shown in “Cabling Options” on page 8-3.
2. Create redundant groups.
3. Specify device priority. (This task is optional).
4. Enter the password(s), if you want to encrypt or authenticate HA communications between members of a redundant group.
5. Set up IP addresses with which the units can perform path monitoring tests. (This task is optional.)
6. Synchronize configurations.

CABLING OPTIONS

Cabling for HA differs on the NetScreen-100 and NetScreen-1000. With the NetScreen-100, you can use any of the physical ports—Trusted, Untrusted, or DMZ—and specify it as a Virtual HA port. The NetScreen-1000 has a dedicated physical port—the HA port—used solely for cabling NetScreen-1000 devices together into redundant groups.

NetScreen-100 Diagrams and Directions

The NetScreen-100 can accommodate the HA connection on any one of its interfaces—Trusted, Untrusted, or DMZ[2]—without affecting the normal traffic flow passing over that interface unless resources become limited. In that case, the normal flow takes priority and HA communications are dropped, which can result in unnecessary fail-over events.

> ***Note:*** *Because HA communication contains sensitive information, NetScreen does not recommend enabling a Virtual HA interface on a physical interface handling normal network traffic. For both performance and security, NetScreen recommends dedicating one interface solely to HA communications.*

Whichever interface you choose, you must use the same one on all NetScreen devices in the same redundant group.

You can designate any of the three interfaces as the Virtual HA Interface as follows:

WebUI

Admin >> HA: Select **DMZ**, **TRUST**, or **UNTRUST** from the Link Port drop-down list.

CLI

set ha interface <dmz | trust | untrust>

2. A Virtual HA interface is a physical interface (Trusted, Untrusted, or DMZ) that, in addition to its regular duties passing network traffic, also functioning as a virtual interface for HA communications. This contrasts with the NetScreen-1000, which has a dedicated physical HA interface.

Connection Diagrams

Two HA connection diagrams for the NetScreen-100 are provided below.

The figure below shows a basic HA configuration using the DMZ port to link the two NetScreen-100 devices. Note also that the redundant switches and routers on both the Trusted and Untrusted interfaces add further protection against inaccessibility due to a device failure. Should either a primary switch, router or Master NetScreen device fail, the devices signal each other to reroute traffic through the functioning secondary switch, router or Slave unit.

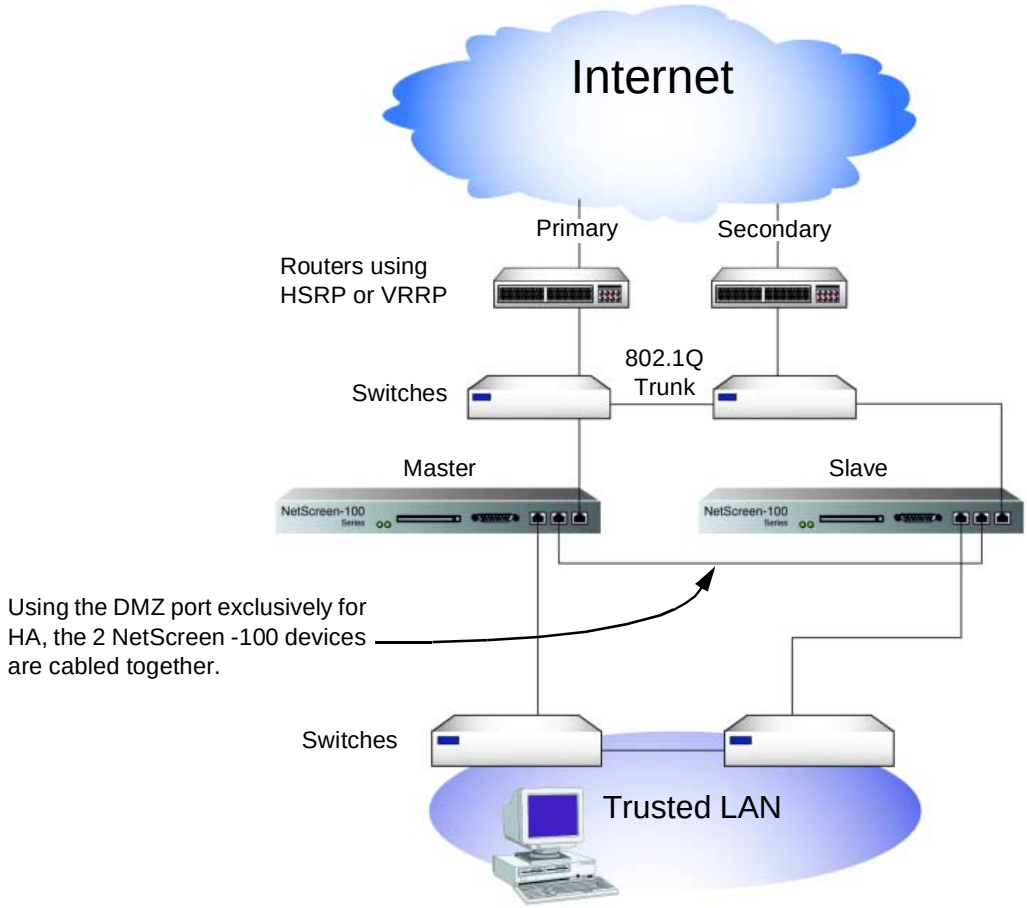

The following figure shows two NetScreen-100 devices using the DMZ port as a Virtual HA interface. Network traffic flows to and from the DMZ while the control messages and synchronization state messages required for HA also pass across the DMZ interface. The redundant switches and routers on all interfaces again provide an extra level of high availability in the event of a device failure.

> ***Note:*** *In this scenario, the DMZ switch becomes the single point of failure for the DMZ network.*

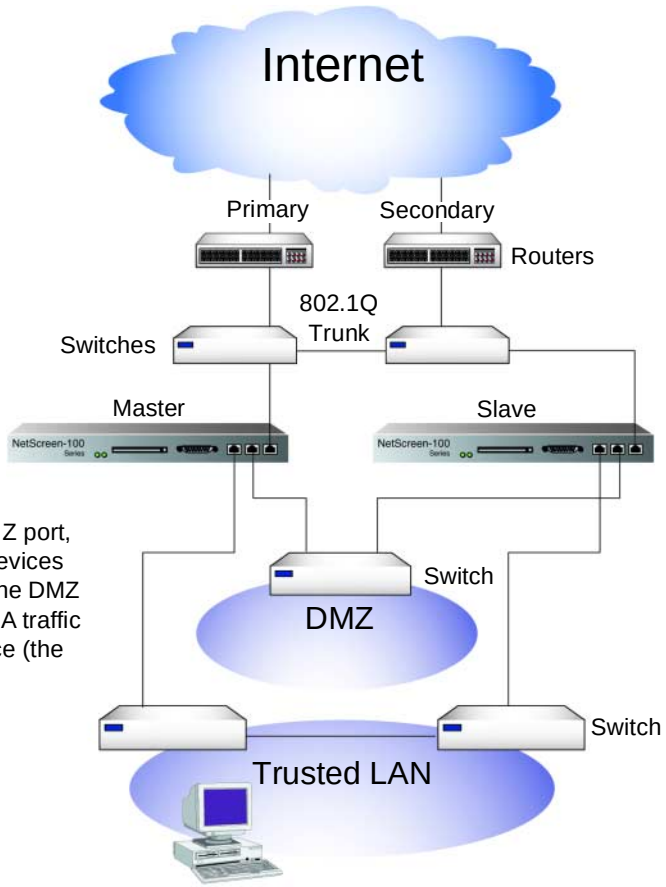

Assigning HA to the DMZ port, the 2 NetScreen -100 devices route network traffic to the DMZ while also maintaining HA traffic across the DMZ interface (the Virtual HA interface).

Basic Connection Procedure

To cable two or more NetScreen-100 units for HA, do the following:

1. For two NetScreen-100 units protecting only a Trusted network, connect a cross-over 10/100BaseTx cable to the DMZ port on each unit.

   Or

   For two NetScreen-100 units protecting both Trusted and DMZ networks, connect a cross-over cable from the DMZ port on each unit to a hub or switch, and then connect that hub or switch to the DMZ network.

2. Connect a cross-over cable from the Trusted port on the Master unit to one switch.
3. Connect a cross-over cable from the Trusted port on the Slave unit to another switch.
4. Connect a straight-through cable from the Untrusted port on the Master unit to a switch.
5. Connect a straight-through cable from the Untrusted port on the Slave unit to another switch.
6. Cable the switches on the Untrusted side together.
7. Cable the switches on the Trusted side together.
8. Connect the switches on the Untrusted side to routers.

Note: *See the switch and router vendor's documentation for the best configuration method to use.*

NetScreen-1000 Diagram and Directions

The NetScreen-1000 has a dedicated physical port for its high-availability interface.

If you are connecting two NetScreen-1000 devices for high-availability (HA), use a cross-over Fast Ethernet cable for the connection between the two devices. The maximum cable length you can use is 100 meters. If the two NetScreen devices are located in separate buildings, they must be connected through a switch.

> ***Note:*** *HA communications are proprietary layer-two broadcasts and cannot be routed.*

Depending on your network setup, you might use a gigabit-only or a gigabit and Fast Ethernet combination switch.

Connection Diagram

The following figure shows two NetScreen-1000 devices connected with a cable between the HA port of each device. Note that the redundant sets of switches and routers add an extra level of redundancy in the event of a device failure.

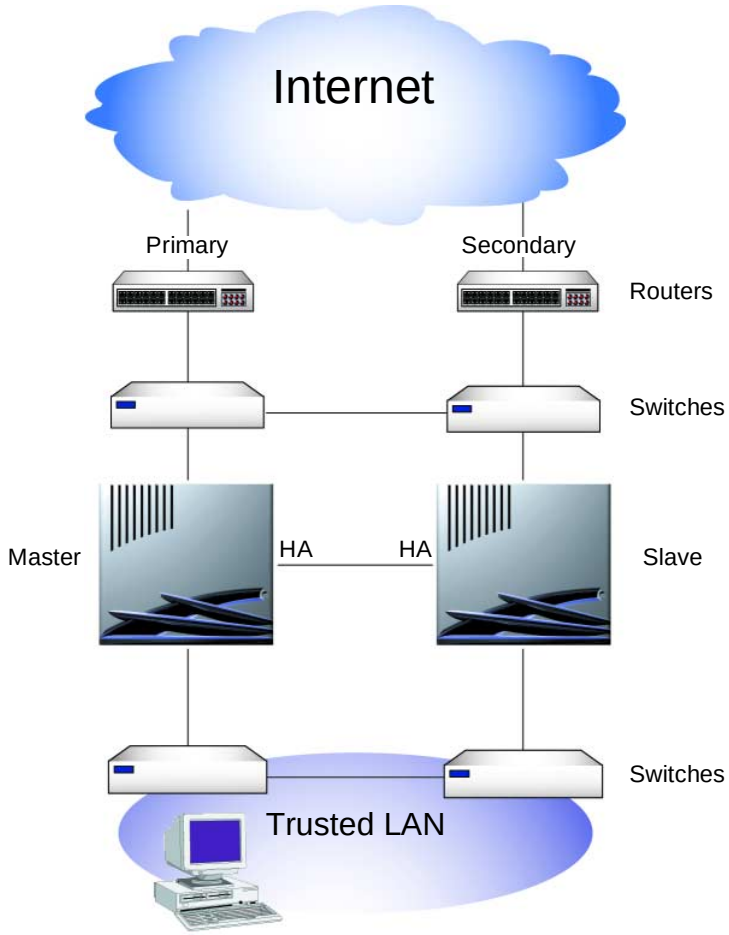

Basic Connection Procedure

To cable two NetScreen-1000 units for HA, do the following:

1. Connect the 10/100 Fast Ethernet cable to the HA port in the base of the Auxiliary board in each NetScreen-1000 unit.
2. Connect a 1-gigabit fiber optic cable from the Trusted port on the Master unit to one switch.
3. Connect a 1-gigabit fiber optic cable from the Trusted port on the Slave unit to another switch.
4. Connect a 1-gigabit fiber optic cable from the Untrusted port on the Master unit to a switch.

5. Connect a 1-gigabit fiber optic cable from the Untrusted port on the Slave unit(s) to another switch.
6. Cable the switches on the Untrusted side together.
7. Cable the switches on the Trusted side together.
8. Connect the switches on the Untrusted side to routers.

Note: *See the switch and router vendor's documentation for the best configuration method to use.*

REDUNDANT GROUPS

Before two NetScreen-100 or -1000 devices can interact with each other in HA mode, you must configure them in a redundant group. In a redundant group, one of the NetScreen units acts as the Master unit, performing firewall, VPN, and traffic management. The Slave unit acts as a hot stand-by, automatically keeping up with all configuration and session state information. If the Master unit fails, one of the Slave units automatically becomes the new Master unit. When the failed unit is repaired and placed back in service, it becomes a Slave unit for the existing Master unit.

You can also specify device priority in the redundant group hierarchy. For the NetScreen-100, the group priority numbers can be from 1 to 255. For the NetScreen-1000, the group priority numbers can be from 1 to 65,535. The device with the number closest to 1 is the Master unit. A value of 0 disables high availability and, for the NetScreen-1000, shuts down the HA port.

Alternatively, you can allow the devices to elect a Master and determine priority numbers automatically during startup. The device that boots up first becomes the Master. If you have more than two units in the redundant group, the order in which each of the subsequent units boots up determines their priority within the group.

Note: *If two units boot up simultaneously, the unit with the lower MAC address becomes the Master.*

For the NetScreen-100, the status of the unit is indicated by the Status LEDs. The Status LED on the Master unit glows green. The LED on the Slave unit glows yellow. For the NetScreen-1000, the Status HA LEDs on the Auxiliary board function identically—the Master glows green; the Slave glows yellow.

When forming a redundant group, it is important to configure the Manage IP address on each NetScreen device for the interface(s) through which you intend to manage the device and from which you plan to perform path monitoring. The Manage IP remains associated with the unique MAC address for its interface, while virtual MAC addresses, associated with the interface IP addresses, are identical among all the members of the redundant group. The virtual MAC addresses are created when you set a group ID number. The virtual MACs are formed as follows: 0010dbff<group_number><interface_value[3]>.

Example: Forming a Redundant Group

The administrator places two NetScreen-1000 devices in a redundant group with the ID number 3. He assigns each device with the same group ID number and determines the Master and Slave units by assigning the Master priority number 1. (The Slave receives the default priority number 100.) He also configures Manage IP addresses for the Trusted and Untrusted interfaces on both devices.

Through the CLI, the administrator also sets up the Slave unit in a link-up state. During a failover, when a Slave unit in a link-up state becomes the active device, it does not need to perform 802.1Q trunk negotiations (NetScreen-1000) nor execute the Spanning Tree Protocol (STP), which can take anywhere from 30 to 50 seconds to complete. The new Master unit simply sends out a series of ARPs, alerting other network devices to its presence.

Also, through the CLI, the administrator indicates a secondary path for the Slave to communicate with the Master. In case the Slave does not receive a heartbeat from the Master along the primary HA path, the Slave sends a message along the secondary path to confirm that the Master unit is truly down. If the Master unit is still up (perhaps only the cable carrying HA traffic has been damaged, causing the heartbeat loss) and the Slave promotes itself to Master, two devices would be acting as the Master, resulting in a condition known a split-brain scenario. Use of a secondary path to confirm the Master's condition can help avoid such a problem.

> ***Note:*** *With NetScreen-100 devices, you also need to specify a physical interface (Trusted, Untrusted, or DMZ) to act as the Virtual HA interface.*

3. Values for the interfaces are as follows: Trusted = 0, Untrusted = 1, DMZ = 2.

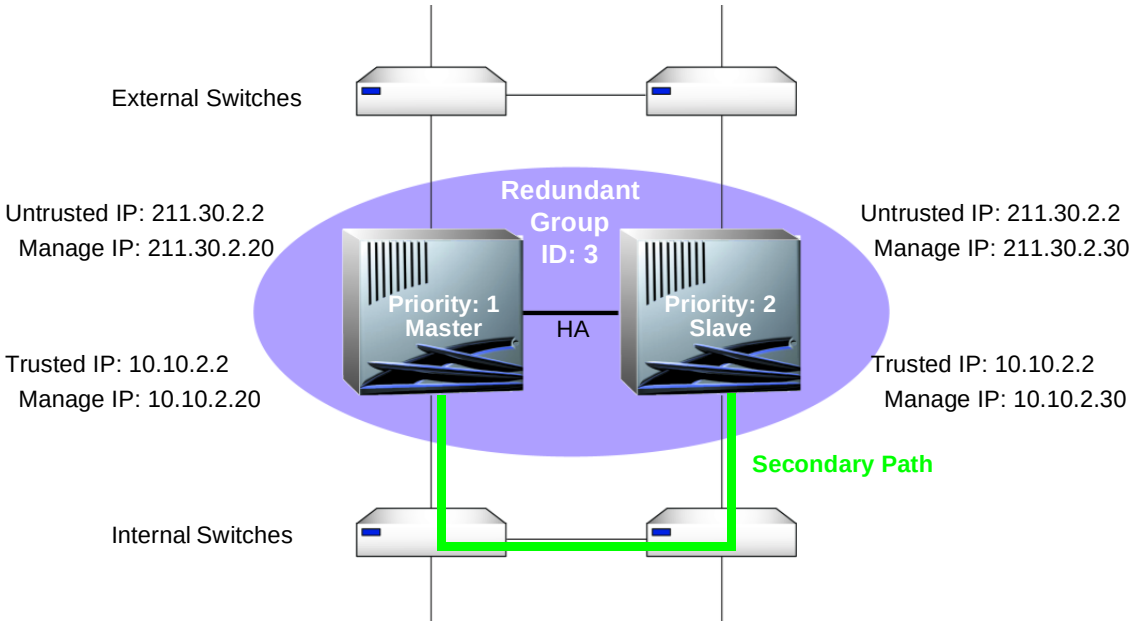

WebUI

Master Unit

1. Interface >> Trusted >> Edit: Enter the following, and then click **Save**:
   IP Address: 10.10.2.2
   Netmask: 255.255.255.0
   Manage IP: 10.10.2.20
2. Interface >> Untrusted >> Edit: Enter the following, and then click **Save**:
   IP Address: 211.30.2.2
   Netmask: 255.255.255.0
   Manage IP: 211.30.2.20
3. Admin >> HA: Enter the following, and then click **Apply**:
   (NetScreen-100) HA Port: DMZ
   Group ID: 3
   Priority: 1

Slave Unit

1. Interface >> Trusted >> Edit: Enter the following, and then click **Save**:
   IP Address: 10.10.2.2
   Netmask: 255.255.255.0
   Manage IP: 10.10.2.30
2. Interface >> Untrusted >> Edit: Enter the following, and then click **Save**:
   IP Address: 211.30.2.2
   Netmask: 255.255.255.0
   Manage IP: 211.30.2.30
3. Admin >> HA: Enter the following, and then click **Apply**:
   (NetScreen-100) HA Port: DMZ
   Group ID: 3

CLI

Master Unit

1. set interface trust 10.10.2.2 255.255.255.0
2. set interface trust manage-ip 10.10.2.20
3. set interface untrust 211.30.2.2 255.255.255.0
4. set interface untrust manage-ip 211.30.2.20
5. set ha group 3
6. set ha priority 1
7. set ha link-up-on-slave
8. set ha second-path trust

Slave Unit

1. set interface trust 10.10.2.2 255.255.255.0
2. set interface trust manage-ip 10.10.2.30
3. set interface untrust 211.30.2.2 255.255.255.0
4. set interface untrust manage-ip 211.30.2.30
5. set ha group 3
6. save config ha-master
7. reset

   Configuration modified, save? [y]/n (Type **n**.)

   System reset, are you sure? y/[n] (Type **y**.)

Disabling HA

You can disable HA functionality by removing a device from its redundant group. You can do this by setting the group ID to 0 (in the WebUI or CLI), or by using the CLI command **unset ha group**.

SECURING HA COMMUNICATIONS

All communications between the Master and all Slaves can be authenticated and encrypted. To secure HA communications, you must first enable authentication and encryption, and then specify passwords. The passwords are used as seed material for the device to generate the preshared key.

All Master and Slave units in the same redundant group must have the same authorization password and the same encryption password, although the authorization password can be different from the encryption password.

Example: Enabling Authentication and Encryption

The two devices in the redundant group authenticate and encrypt HA communications. The authentication password is p6NE9mKq and the encryption password is 45Zg8HB1.

WebUI

1. Admin >> HA: Select **HA Authentication Password**, and type **p6NE9mKq**.
2. Admin >> HA: Select **HA Encryption Password**, type **45Zg8HB1**, and then click **Apply**.

CLI

1. set ha authentication password p6NE9mKq
2. set ha encryption password 45Zg8HB1
3. save

PATH MONITORING

Path monitoring checks the network connection between a NetScreen interface and the interface on another device. Path monitoring can occur at both layer 2 (Ethernet) and layer 3 (IP). Path monitoring for the Ethernet link state occurs automatically. Layer 3 (IP) path monitoring functions by sending a ping request at determined intervals and then monitoring if the target replies with a ping reply. A successful attempt indicates that layer 3 connectivity is possible over the network. A failed attempt indicates that the two devices cannot make a layer 3 connection.

Path tracking is a useful tool for devices within a redundant group to determine whether the network connectivity of a Master unit is better than or as good as a Slave unit. If Slave unit has a better success rate than the Master unit (for example, the Slave unit gets 5 ping replies and the Master gets only 4), then it automatically is promoted to Master, and the Master unit is demoted to Slave.

> ***Note:*** *The NetScreen-100 and -1000 can monitor up to 16 IP addresses.*

When using path monitoring, you must set a Manage IP address for the physical interface, or interfaces, from which the NetScreen devices ping. The Slave unit uses the Manage IP address to distinguish itself from the Master unit, which uses the physical IP address shared by all the devices in the redundant group.

> ***Note:*** *The Manage IP address is also the address with which you can configure a Slave unit.*

Example: Enabling Path Tracking

Two NetScreen-100 devices are in a redundant group. Both are set to ping the IP addresses of four Web servers every 6 seconds. The alarm threshold is set to 3 failures. Not receiving a ping response after 5 consecutive attempts is considered a failed attempt. (The default threshold for failure is 3 failed ping attempts.) In this example, the Master unit has a 100% success rate, while the Slave unit succeeds only 50% of the time.

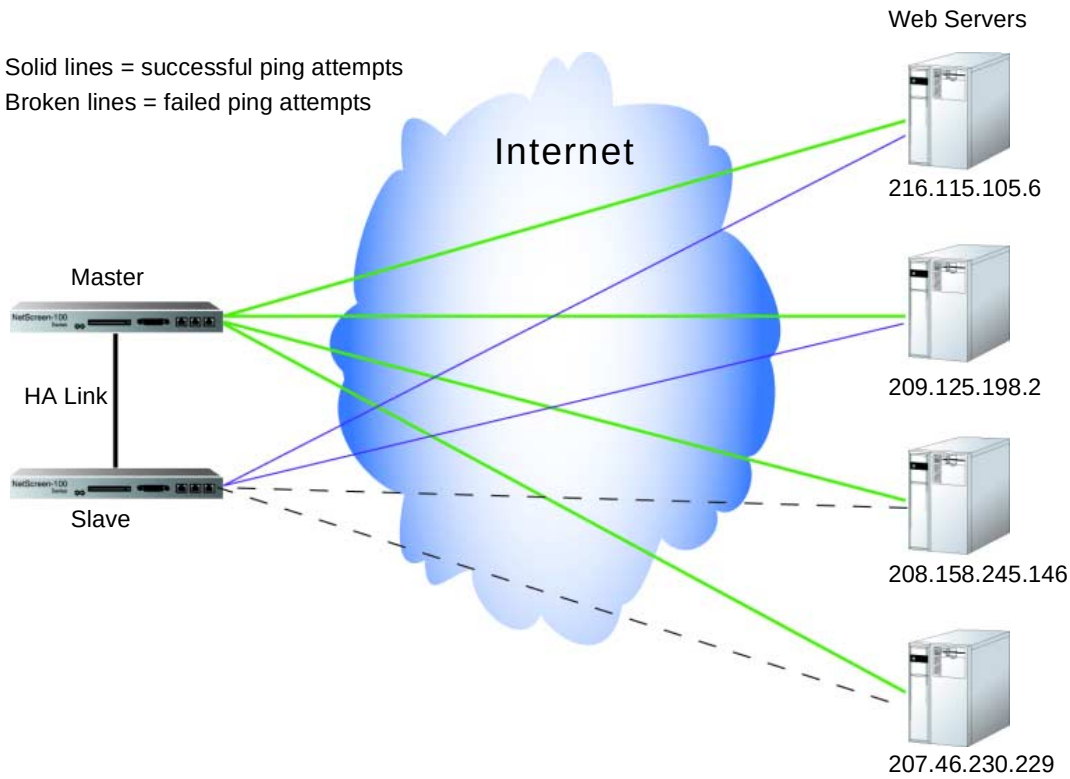

WebUI

1. Admin >> HA >> New Path: Enter the following, and then click **OK**:
   Track IP: 216.11.105.6
   Seconds: 6
   Threshold: 5
   Interface: Untrusted
2. Admin >> HA >> New Path: Enter the following, and then click **OK**:
   Track IP: 209.125.198.2
   Seconds: 6
   Threshold: 5
   Interface: Untrusted

3. Admin >> HA >> New Path: Enter the following, and then click **OK**:
   Track IP: 208.158.245.146
   Seconds: 6
   Threshold: 5
   Interface: Untrusted
4. Admin >> HA >> New Path: Enter the following, and then click **OK**:
   Track IP: 207.46.230.229
   Seconds: 6
   Threshold: 5
   Interface: Untrusted
5. Admin >> HA: Select **Path Tracking on**, and then click **Apply**.

CLI

1. set ha track ip 216.115.105.6 interface untrust
2. set ha track ip 216.115.105.6 interval 6
3. set ha track ip 216.115.105.6 threshold 5
4. set ha track ip 209.125.198.2 interface untrust
5. set ha track ip 209.125.198.2 interval 6
6. set ha track ip 209.125.198.2 threshold 5
7. set ha track ip 208.158.245.146 interface untrust
8. set ha track ip 208.158.245.146 interval 6
9. set ha track ip 208.158.245.146 threshold 5
10. set ha track ip 207.46.230.229 interface untrust
11. set ha track ip 207.46.230.229 interval 6
12. set ha track ip 207.46.230.229 threshold 5
13. set ha track ip
14. save

Monitoring NetScreen Devices 9

This chapter discusses the following topics that can be used when monitoring NetScreen devices:

SYSLOG

Syslog is facility that enables the logging of system events to a single file for later review. A NetScreen device generates syslog messages for system events at user-defined priority levels and sends these messages via UDP (port 514) to a syslog host, which runs on a UNIX system. Syslog messages can be used by the syslog host to create e-mail alerts for the system administrator, or be displayed 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).*

You can also send syslog messages through a VPN tunnel. In the WebUI, select the **Enable Syslog VPN encryption**. In the CLI, use the **set syslog vpn** command. (For more about this option, see "Virtual Private Networks" on page 3-24.)

Logging Priority Levels

The syslog message levels are as follows:

- **EMERGENCY:** Generates messages on SYN attacks, Tear Drop attacks, and Ping of Death attacks. For more information on these types of attacks, see "Firewall Protection" on page 2-2.
- **ALERT:** Generates messages for multiple user authentication failures and other firewall attacks not included in the emergency category.
- **CRITICAL:** Generates messages for URL blocks, high availability (HA) status changes, and global communications.
- **ERROR:** Generates messages for admin name and password changes.
- **WARNING:** Generates messages for admin logins and logouts, failures to log in and log out, and user authentication failures, successes, and timeouts.
- **NOTICE:** Generates messages for link status changes, load balancing server status changes, and traffic logs.
- **INFO:** Generates any kind of message not specified in other categories.
- **DEBUG:** Generates all messages.

Syslog messages are organized hierarchically, so that setting a level includes that level and all levels above it. For example, an alert setting generates messages for alert and emergency messages, whereas a debug setting generates messages for all levels.

> ***Note:*** *You can also send traffic logs with the syslog messages.*

WebTrends

WebTrends® offers a product called the WebTrends Firewall Suite that allows you to customize syslog reports to display the information you want in the format you specify. You can create reports that focus on areas such as firewall activity, network traffic flow, or event alarms.

> ***Note:*** *The WebTrends Syslog Server and the WebTrends Firewall Suite must run on the same Windows NT system. You must have administrator rights to configure it.*

You can also send WebTrends messages through a VPN tunnel. In the WebUI, select the **Enable WebTrends VPN encryption**. In the CLI, use the **set webtrends vpn** command. (For more about this option, see "Virtual Private Networks" on page 3-24.)

Example: Enabling Syslog and WebTrends

The following example illustrates how to set up the syslog facility to send messages with a priority of DEBUG or higher to port 514 on a WebTrends Syslog Server at 172.10.16.25. The security and facility levels are set to Local0. Traffic logs are included with the system event messages.

WebUI

Admin >> Syslog: Enter the following, then click **Apply**.

Syslog Host Name: 172.10.16.25

Syslog Host Port: 514[1]

Security Facility: Local0

Facility: Local0

Only log messages with a priority level of "x" or higher: DEBUG

1. The syslog host port number must match the WebTrends® port number.

Enable syslog messages: Check this box to enable this feature.

Send traffic log messages: (select)

Enable WebTrends Messages: (select)

WebTrends Host Name: 172.10.16.25

WebTrends Host Port: 514

Note: *When you enable syslog and WebTrends on a NetScreen-5, 10, 100, or 1000 running in Transparent mode, you must set up a static route on the Route Table. For more information, see "Route Table Configuration" on page 2-12.*

CLI

1. set syslog config 172.10.16.25 auth/sec Local0
2. set syslog auth/sec local0
3. set syslog enable
4. set syslog traffic
5. set syslog webtrends enable ip 172.10.16.00 port 514
6. save

SNMP

The Simple Network Management Protocol (SNMP) agent for the your NetScreen 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.

The NetScreen devices support the SNMPv1 protocol (described in RFC-1157) and all relevant Management Information Base II (MIB II) groups defined in RFC-1213.

Accordingly, the NetScreen SNMP agent generates the following traps, or notifications, when specified events or conditions occur:

- **Cold Start Trap:** The cold start trap is generated when the NetScreen device becomes operational after you power it on.
- **Trap for SNMP Authentication Failure:** The authentication failure trap is triggered if the SNMP manager sends the incorrect community string.
- **Traps for System Alarms:** System alarms are triggered by firewall conditions and NetScreen device error conditions. Three NetScreen enterprise traps are defined to cover alarms related to hardware, security, and software. (For more information on firewall settings and alarms, see “Firewall Protection” on page 2-2, and “Alarms” on page 9-15.)
- **Traps for Traffic Alarms:** Traffic alarms are triggered when traffic exceeds the alarm thresholds set in Access Policies. (For more information on configuring Access Policies, see “Access Policies” on page 5-1.

The following table list possible alarm types and their associated trap number:

Trap Enterprise ID	Description
100	Hardware problems
200	Firewall problems
300	Software problems
400	Traffic problems
500	VPN problems

For more information on types of problems, see “Logs” on page 9-12.

> ***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. There are also several shareware and freeware SNMP manager applications available from the Internet.*

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

Implementation Overview

The following points summarize how SNMP is implemented in NetScreen devices:

- The network administrator can create up to three communities, each containing up to eight hosts. Hosts must be listed individually; they cannot be specified as a range.
- Each community has either read-only or read-write permission for the MIB II data.
- Each community can be allowed or denied receiving traps.
- Access to the MIB II data and traps is available through any physical interface.
- Each system alarm generates a single NetScreen enterprise SNMP trap to each of the hosts in each community that is set to receive traps.
- Cold Start / Link Up / Link Down traps will be sent to all hosts in communities that are set to receive traps.
- If you specify trap-on for a community, you also have the option to allow traffic alarms.

You can also send SNMP messages through a VPN tunnel. In the WebUI, select the **Enable SNMP VPN encryption**. In the CLI, use the **set snmp vpn** command. (For more about this option, see “Virtual Private Networks” on page 3-24.)

Example: Setting Up SNMP Communities

In the example below, we configure SNMP for two communities, named "JCarney" and "TCooper." In the first community, its members can read MIB II data and receive traps. In the second community, its members can read and write MIB II data, receive traps, and traffic alarms. The contact person is "John Fisher" in "Miami." The JCarney community host IP addresses are 172.16.20.181, 172.16.40.245, and 172.16.40.55. The TCooper community host is 172.16.20.250.

> ***Note:*** *The MIB II system group variables sysContact, sysName (which is the same as the host name), sysLocation, and sysServices are read/write objects. All other variables are read-only.*

WebUI

1. Admin >> SNMP: Enter the following settings, and then click **Apply**:

   System Contact: John Fisher

   Location: Miami

2. New Community: Enter the following settings, and then click **OK**:

   Community Name: JCarney

   Trap: (select)

   Hosts: 172.16.20.181
   172.16.40.245
   172.16.40.55

3. New Community: Enter the following settings, and then click **OK**:

   Community Name: TCooper

   Community Name: Private

   Write: (select)

   Trap: (select)

   Traffic Alarms: (select)

   Hosts: 172.16.20.250

CLI

1. set snmp contact John Fisher
2. set snmp location Miami
3. set snmp community JCarney read-only trap-on
4. set snmp host JCarney 172.16.20.181
5. set snmp host JCarney 172.16.40.245

6. set snmp host JCarney 172.16.40.55
7. set snmp community TCooper read-write trap-on `traffic`
8. set snmp host TCooper 172.16.20.250
9. reset

VPN Monitoring

The NetScreen ScreenOS provides the ability to determine the status and condition of active VPNs through the use of SNMP VPN monitoring objects and traps.

> ***Note:*** *To enable your SNMP manager application to recognize the VPN monitoring MIBs, you must import the NetScreen-specific MIB extension files into the application. You can find the MIB extension files on the NetScreen documentation CD that shipped with your NetScreen device.*

By enabling the VPN monitoring feature on a Manual Key or AutoKey IKE VPN tunnel, the NetScreen device activates its SNMP VPN monitoring objects, which note data on the following:

- The total number of active VPN sessions
- The time each session started
- The Security Association (SA) elements for each session:
  - ESP encryption (DES or 3DES) and authentication algorithm (MD5 or SHA-1) types
  - AH algorithm type (MD5 or SHA-1)
  - Key exchange protocol (AutoKey IKE or Manual Key)
  - Phase 1 authentication method (Preshared Key or certificates)
  - VPN type (dialup or peer-to-peer)
  - Peer and local gateway IP addresses
  - Peer and local gateway IDs
  - Security Parameter Index (SPI) numbers
- Session status parameters
  - VPN monitoring status (up or down)
  - Tunnel status (up or down)
  - Phase 1 and 2 status (inactive or active)
  - Phase 1 and 2 lifetime (time in seconds before rekeying; Phase 2 lifetime is also reported in remaining bytes before rekeying)

When VPN monitoring is enabled, the NetScreen device also pings the remote gateway at specified intervals (in 10-second increments) to monitor network connectivity between the two VPN gateways.[2] The pings are sent through the VPN tunnel between the two end points of the tunnel—the Untrusted interface IP addresses (NAT or Route mode) or the System IP addresses (Transparent mode). The VPN monitoring MIB notes whether the ping elicited a response, a running average of successful responses, the latency of the response, and the average latency over the last 30 attempts.

An SNMP trap is triggered when the ping activity indicates that the VPN status has changed:

- **Up to Down:** The state of the VPN tunnel is up, but the ping request has not elicited a response after a specified number of attempts.
- **Down to Up:** The state of the VPN tunnel is down, but the ping request elicits a response.

To enable VPN monitoring, do the following:

WebUI

VPN >> Manual Key >> New Manual Key Entry: Configure the VPN, select the **VPN Monitor Enable** check box, and then click **OK**.

Or

VPN >> AutoKey IKE >> New AutoKey IKE Entry: Configure the VPN, select the **VPN Monitor Enable** check box, and then click **OK**.

CLI

1. Configure the Manual Key or AutoKey IKE VPN.
2. set vpn <vpn_name> monitor
3. set vpnmonitor frequency <number (in 10-second increments)>
4. save

2. You must use the following CLI command to change the ping interval: **set vpnmonitor**. The default is 10 seconds.

NETSCREEN-GLOBAL MANAGER

You can use NetScreen-Global Manager to monitor the network traffic, resource utilization, and system events on multiple NetScreen devices. For example, you can view a graph of network traffic in real-time for any NetScreen device that you have configured to communicate with NetScreen-Global Manager. Each NetScreen device displays its own graph of network traffic, and you can view more than one graph at a time.

NetScreen-Global Manager provides four reports for the NetScreen devices that it manages:

- Network activity—Data on the amount of traffic that flows through each of the NetScreen device interfaces: Trusted, Untrusted, and DMZ (NetScreen-10 and -100).
- Resource utilization—Percentage of CPU, flash card, and memory used every second.
- Event log—Data on the alarms triggered, changes to the NetScreen device configuration, and general system information.
- Summary report—Traffic information according to source IP addresses, destination IP addresses, services, Access Policies, and VPNs.

The NetScreen-Global Manager can provide report data for thousands of NetScreen devices, but Windows resources limit the number of reports you can view at any particular time.

NetScreen-Global PRO

NetScreen-Global Pro, working with NetScreen-Global Manager, allows you to view additional logs and reports, including protocol distribution, active-user, and interface summary reports. For more information about NetScreen-Global PRO, see the documentation that accompanies the application. (These documents are listed in “NetScreen Documentation” on page xv.)

Note: *You can also send NetScreen-Global Manager and NetScreen-Global PRO messages through a VPN tunnel. In the WebUI, select the* ***Enable Global Manager/PRO VPN encryption****. In the CLI, use the* ***set global vpn*** *command. (For more about this option, see “Virtual Private Networks” on page 3-24.)*

COUNTERS

Counters measure traffic as defined in the Access Policies section.

Example: Viewing Counters

This example illustrates how to view the amount of traffic per second handled by the first Incoming Access Policy, which has a policy ID of 3.

WebUI

1. Counters >> Incoming >> View Count Details: Click on a line in the graph to view information at that interval.

   Note: *The X-axis represents time and the Y-axis represents the number of bytes. The X-axis displays seconds, minutes, hours, days, or months, depending on which tab you select. The color of the bar appears in blue, unless an alarm threshold was set and exceeded, in that case the bar is red.*

2. Counters >> View Count Details >> Download to File: Specify the desired location of the counter data for review and analysis.
3. Counters >> View Count Details >> Update Now: The screen refreshes with the most recent data available.

CLI

```
get counter policy 3 second
```

LOGS

NetScreen devices categorize system-level events as system events, traffic events, and denied traffic. System events are discussed in “Events Log” on page 9-12, traffic events in “Traffic Log” on page 9-13, and denied traffic in “Self Log” on page 9-14.

Events Log

System events are classified into two distinct categories: information (indicated by a yellow circle in the WebUI) and configuration (indicated by a red circle).

You can view system events through the WebUI, the CLI, or you can send event logs to a syslog host, (see “Syslog” on page 9-2), a WebTrends host (“WebTrends” on page 9-3), or NetScreen-Global Manager (see “NetScreen-Global Manager” on page 9-10).

WebUI

Log >> Event Log

CLI

get log event[3]

3. The **get log** command allows you to impose parameters to determine the events displayed. For more information, see the *NetScreen CLI Reference Guide.*

Traffic Log

NetScreen devices provide comprehensive monitoring tool to monitor traffic flow in real-time. You can view all graphs of the usage data with a browser on the Internet, or download it and import it into a spreadsheet or database for numerous applications.

NetScreen devices maintain a traffic log. To view traffic logs you have defined in the Access Policies page:

WebUI

Log >> Traffic Log >> View Log Entries: Select the action you want to perform:

Click **Download to File** to save the data to a desired location for review and analysis.

The data can be saved to your local C: drive in a *.txt text format. The file contents are tab-delimited.

Click **Clear Log** to clear the log after downloading the most recent data available.

CLI

get log traffic

Self Log

In addition to the Event Log and Traffic Log, each NetScreen supports a logging function called Self Log. The Self Log records all the dropped packets detected by that NetScreen device; that is, the log shows all the traffic which terminated at the device and was denied.

In addition to viewing the Self 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.

Example: Downloading the Self Log

This example illustrates how to download a Self Log to your workstation desktop (WebUI) or to a TFTP server with the IP address 10.10.20.200 (CLI).

WebUI

1. Log >> Log Self: Select the **Download to File** option.

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

2. Select the **Save this file to disk** option.

   The File Download wizard prompts you to choose a directory.

3. Specify the desktop, name the file and give it a .txt extension, and then click **Save**.

CLI

get log self > tftp 10.10.20.200

ALARMS

Your NetScreen device maintains two types of alarms—traffic alarms, when traffic thresholds on Access Policies are exceeded, and event alarms, when system events occur. You can also configure the your NetScreen device to alert you through one or more of the following methods whenever an alarm is generated:

- E-mail
- Syslog or WebTrends
- SNMP
- NetScreen-Global Manager or NetScreen-Global PRO

Traffic Alarms

You define and enable Traffic Alarms by setting alarm thresholds in the Access Policy Configuration dialog box. They are triggered when traffic exceeds those thresholds. (For more information on monitoring traffic, see “Traffic Log” on page 9-13.)

Alarms >> Traffic Alarm >> Recent Alarm Time: View alarm details by selecting the following:

To save the data to the specified location for review and analysis, click **Download to File**.
You can save the data to your local C: drive in a *.txt text format. The file contents are tab-delimited.

To erase all the data, click **Clear Alarms**.

To move to the next/previous page, click **Next/Previous**.

Event Alarms

Event alarms, such as those for attacks against which you have configured the device to guard, are triggered by the events themselves. (For more information on monitoring events, see “Events Log” on page 9-12.)

WebUI

Alarms >> Event Alarm: View the Event Alarm details by selecting the following:

To save the data to the specified location for review and analysis, click **Download to File**.
You can save the data to your local C: drive in a *.txt text format. The file contents are tab-delimited.

Click the **Next/Previous** option to see the next/previous page.

CLI

1. get alarm traffic
2. get alarm event

Example: Sending E-mail Alerts

This example shows how you set up notification by e-mail alerts when there is an alarm. In this case, the mail server is at 172.16.10.254, the first e-mail address to notify is jharker@netscreen.com, then second address to notify is dlittle@netscreen.com.

WebUI

Admin >> Settings:[4] Enter the following information, then click **Apply**:

Enable E-Mail Alert Notification: (select)

SMTP Server Name: 172.16.10.254[5]

E-Mail Address 1: jharker@netscreen.com

E-Mail Address 2: dlittle@netscreen.com

Send Traffic Log: Check this box. If you enable logging when you create an Access Policy, then the traffic log will be sent to the e-mail account(s) at regular time intervals.

CLI

1. set admin mail alert
2. set admin mail mail-addr1 172.16.10.4
3. set admin mail mail-addr2 192.44.10.101
4. set admin mail server-name 172.16.10.254
5. set admin mail traffic-log
6. save

4. On the NS-1000, alarm notification details are on the Admin tab.

5. If you have DNS enabled, you can also use a host name for the mail server, such as mail.netscreen.com.

10 Troubleshooting NetScreen Devices

This chapter focusses on responding to NetScreen devices that have stopped working, or are working at compromised levels. Specific topics covered are:

Use this chapter as a guide to help you fix NetScreen device problems before calling technical support. For information on how to contact technical support, see "Contacting Technical Support" on page 10-45.

RESPONDING TO HARDWARE FAILURES

This section discusses hardware problems. In general, when troubleshooting, you should first check the hardware before investigating the software. Otherwise, you might spend hours trying to fix a problem by redoing configurations only to later discover that a loose wire was the cause of all the trouble.

The NetScreen Device Does Not Power On

When you power on the NetScreen-5, -10, or -100, verify that it has started successfully by confirming that the Power LED illuminates green and the Status LED is blinking green.

When you power on the NetScreen-1000, verify it has started successfully by confirming that the green Power LEDs on the Power Supply, Switching, and Processing boards glow steady green.

> ***Note:*** *The NetScreen-1000 takes about 30 seconds to boot. Please wait until the green Com LED on the Master Processing board glows steady, and the Com LED on the Slave Processing boards flash once per second.*

Power LED: Glows steady green

Status LED: Blinks green

Power LEDs and Master Processing Board COM LED: Glow steady green

COM LEDs on Slave Processing Boards: Blink green

Note: *Cables removed for clarity.*

Link LED Is Off

The Link LEDs indicate the connection status between the NetScreen device and the network. If the green link LED is not illuminated, there is a problem with the network connection.

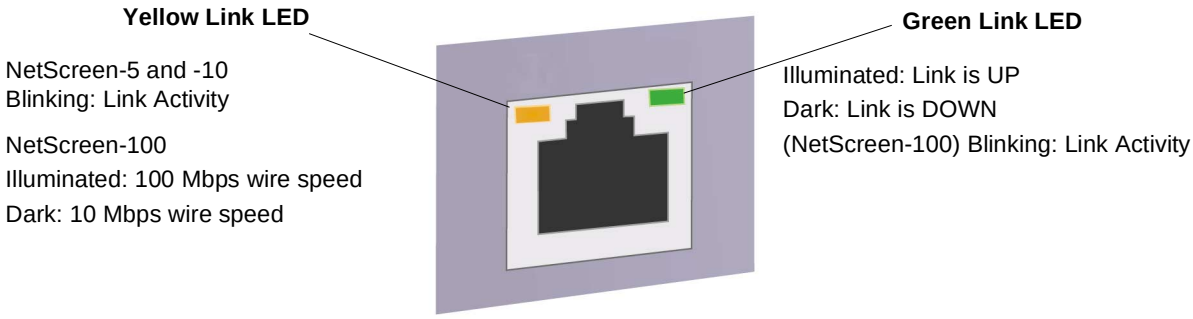

- Verify that the Ethernet cable is properly connected and the network hub is operational.
- Try plugging the Ethernet cable into a different location on the hub, or into a different hub.
- Try another Ethernet cable.
- If the link LED still does not illuminate, there may be a problem with the Ethernet interface. Contact your local NetScreen reseller representative.

Cannot Connect to the Internet

If the NetScreen device cannot access the Internet, check for the following:

- Cables are correct type and are securely connected to the proper ports.
- Link lights on NetScreen, hosts, hubs and routers are illuminated.
- Host IP and netmask are configured correctly for your configuration.
- Host gateway is defined in the host and points to the correct destination: correct router if in Transparent mode, Trusted Interface if in NAT.
- Host is on same subnet as its default gateway.
- Host has a valid DNS entry.
- DNS service is available through the firewall.
- Necessary Access Policies have been created and put in proper order.

NetScreen-1000 Hardware Failures

Power Failures

Follow the procedure below for hot-swapping the power supplies.

1. Unscrew the screws at the top and bottom of the power supply board until they are loose but not completely detached.

   For the location of the screws and to see the labelled parts of a power supply board, see below:

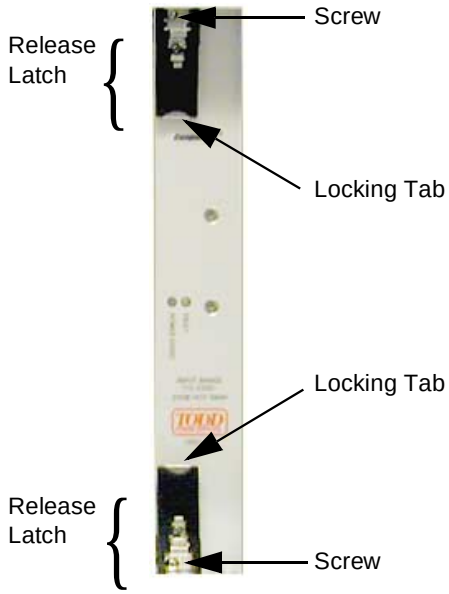

2. Press the top locking tab upward and the bottom locking tab downward.
3. Pull the release latches forward.
4. Completely slide the power supply board straight out toward yourself.

Note: *Because the power boards fit neatly in their slots, if a board becomes skewed during removal, it might rub against the track and get stuck. If this happens, you might need to jog the board slightly to loosen it.*

5. Slide the new power supply board into its slot, pressing firmly until it clicks and the release latches snap back into place.
6. Tighten the screws.

Board Failures

It is recommended to have spare Processing, Auxiliary, and Switching boards on hand. Should one fail, you can promptly replace it with a spare. You can hot swap the Auxiliary and Slave Processing boards; that is, you can remove one board and insert another one without interrupting performance. On the other hand, before changing a Switching board, you must turn off the NetScreen-1000.

To hot swap an Auxiliary or Processing board:

1. Press the upper red locking tab straight up and the lower red locking tab straight down.

   The locking tab and release latch are shown below.

Note: *When you press the red locking tabs, the Hot Swap LED on the Auxiliary board glows red.*

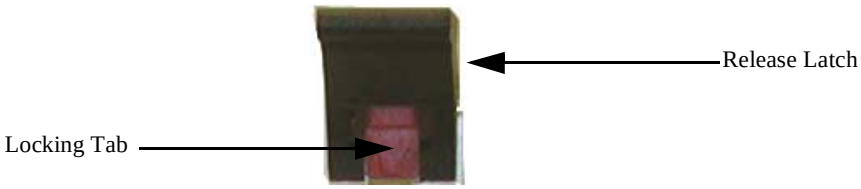

2. Press the upper black release latch upward and the lower black release latch downward.
3. Slide the board straight out until it separates from the backplane. (If the cable lengths permit, you can remove the board completely at this point.)
4. Disconnect the cables from the removed board and connect them to the board that you are about to install.
5. If you have not already done so, remove the old board completely, and slide the new board into the slot, pressing firmly until it clicks into place.

Note: *When you insert the board properly, the red Hot Swap LED goes off.*

Fan Failures

You only need to replace the fan assembly in the event of failure, indicated when the FAN LED on the Auxiliary board glows red (see the *NetScreen-1000 Installer's Guide*), and an event alarm and SNMP trap is triggered. Although the unit can operate with two of the three fans, there is a serious risk of overheating. Especially critical is the far right fan, located directly beneath the dual power supplies. Should the power supplies become too hot, the system automatically shuts down.

To obtain a replacement fan assembly while it is still protected under the one-year warranty, call NetScreen support. After that, contact NetScreen's sales department.

1. Open the fan access port by pulling the upper edge of the door forward.

2. Press the fan assembly release tabs, located at the far left and right sides, toward the center.

Danger *The fans spin at high speed and can cause physical injury. Exercise extreme caution when removing the fan assembly, and handle it only by the sides. Do not reach into any of the 3 air ducts. The following illustration shows the correct method of removing the fan assembly.*

3. With both hands, reach underneath the fan tray, and slide out the unit toward yourself.

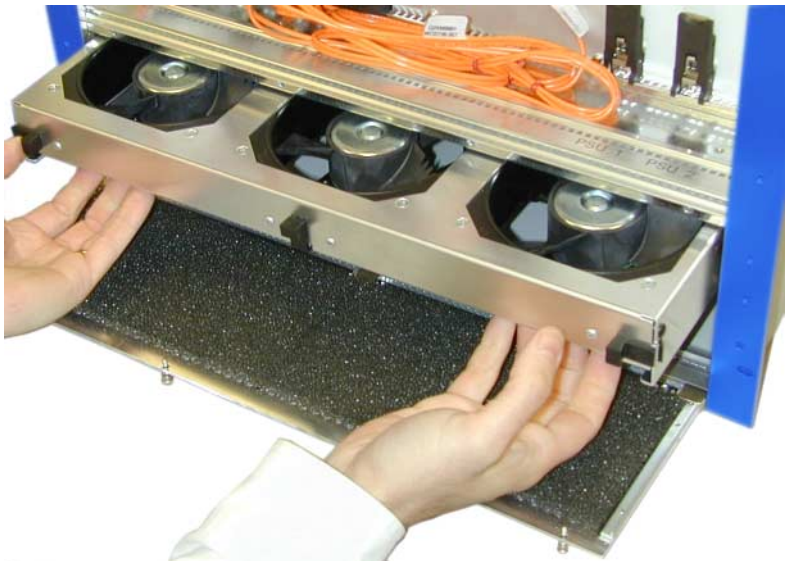

Note: *Note the orientation of the plug (see the following illustration) when you remove the fan tray. Attempting to install the fan assembly upside down will break the plug and can result in serious damage to the NetScreen-1000.*

Fan Plug on the Back of the Fan Assembly

Caution *Do not remove the fan assembly for more than five minutes or the unit might overheat and malfunction.*

If troubleshooting your NetScreen device has not solved the problem, refer to “Contacting Technical Support” on page 10-45.

RESPONDING TO SOFTWARE FAILURES

If your NetScreen device appears to be working at the hardware level (LEDs display normal status, unit has power, cables are connected properly), then you need to verify that your NetScreen device is working at the software level. Examples of software problems are misconfigurations, applications not working, Access Policies not working, and attack alarm notifications.

This section discusses typical troubleshooting strategies for responding to software types of problems with your NetScreen devices. These strategies can be characterized as being one of the following types of problems:

1. Peer-to Peer VPN Troubleshooting—Not working at all
2. Peer-to Peer VPN Troubleshooting— A specific application not working
3. Client (NetScreen-Remote)-to-LAN Troubleshooting
4. MIP/VIP Troubleshooting—Not working
5. MIP/VIP Troubleshooting—A specific application not working
6. Outbound access troubleshooting—Cannot reach the outside or the Internet
7. Access Policy Troubleshooting—Access Policies are not working correctly
8. Responding to attack alarms

Go to the section that best describes the type of problem you are having, and follow the troubleshooting procedures. If you do not know which category best describes your situation, start with the first category and proceed with those troubleshooting procedures.

Peer-to-Peer VPN Troubleshooting

This section discusses troubleshooting strategies for solving problems between NetScreen devices.

Checking Your VPN Configuration

If you are having problems with your peer-to-peer VPN, you should first check your configuration. Use the following procedures, to verify your VPN configuration.

WebUI

VPN >> Manual Key >> New VPN Entry: Verify the following:

VPN Tunnel Name: This field is arbitrary. You can name it whatever you want.

Gateway IP: This needs to be the Untrusted IP address of the remote NetScreen. This is where the NetScreen device will send its encrypted packets.

Security Index (local and remote): These can be any hexadecimal value between 3000 and 8000.

Remember that the remote NetScreen device should use the reverse values for the local and the remote Security Index. For example, if you used 3000 for the local on one side, you need to use 3000 for the remote on the other side.

If you are going to have multiple tunnels leaving one device, the Security Indexes for each tunnel need to be different. These are tunnel identifiers and need to be distinct for each tunnel.

ESP-Encryption Algorithm: Select the type of encryption you want to use. If your top priority is security, use 3DES-CBC (each packet encrypted 3 times). If your top priority is performance, use 56bit DES-CBC. Whatever you use here must be used on the remote side as well.

HEX Key VS. Generated Key By Password: A key is a hexadecimal value that is used to encrypt and decrypt the datagrams. Whether you enter a HEX Key or you use the Generated Key By Password is arbitrary. NetScreen always suggests using the Generated Key By Password. The value that you use on one side of the tunnel must be used on the other side as well.

ESP-Authentication Algorithm: Which this parameter is optional, NetScreen suggests that you use this. Authentication adds another layer of security to the VPN in that it ensures that none of the packets are tampered with in transit. The most

common selection is MD5. SHA-1 is more powerful form of authentication but can sometimes cause problems with some applications. If security is a concern, choose SHA-1. You choice must be the same on both sides of the tunnel.

HEX Key VS. Generated Key By Password: Follow the guidelines outlined in "ESP-Encryption Algorithm: HEX Key VS. Generated Key By Password."

VPN >> AutoKey IKE >> New AutoKey IKE Entry: Verify that:

VPN Tunnel Name: This is arbitrary. You can call the tunnel whatever you want.

Gateway IP: This needs to be the Untrusted IP address of the remote NetScreen. This is where the NetScreen device will send its encrypted packets.

Preshared Key: An ASCII value that will be used to encrypt and decrypt packets as well as help identify the tunnel and provide authentication. This needs to be the same on both sides of the tunnel.

ESP-Encryption Algorithm: Select the type of encryption that you want to use. If your top priority is security, use 3DES-CBC (each packet encrypted 3 times). If your top priority is performance, use 56bit DES-CBC. Whatever you use here must be used on the remote side as well.

ESP-Authentication Algorithm: Select the type of authentication you want. This part of the VPN tunnel is optional. However, NetScreen suggests that you use this. Authentication adds another layer of security to the VPN in that it ensures that none of the packets are tampered with in transit. The most common selection here is MD5. SHA-1 is more powerful form of authentication but can sometimes cause

problems with some applications. If security is a concern, choose SHA-1. You choice must be the same on both sides of the tunnel.

Key Life Time: This value determines how often the tunnel renegotiates. Unless you have a specific need, leave this value as default. If one end of your IKE VPN tunnel fails for whatever reason, you may need to reboot the NetScreen device. This is due to the nature of IKE behavior, and gateways have problems recognizing that one end of the tunnel has broken, and will not renegotiate. The resolution is to reboot the firewall.

CLI

Apply the information in the preceding WebUI section to the following CLI commands for a manual key VPN and an AutoKey IKE VPN:

Manual Key

set vpn <vpn_name> gateway {<a.b.c.d> | <domain_name>} [replay | no-replay] proposal <p2_proposal>

set vpn <vpn_name> manual <local-spi> <remote-spi> gateway <a.b.c.d> [esp {des {key <64-bit hex> | password <string>} | 3des {key <192-bit hex> | password <string>} | null} [auth {md5 {key <16-byte hex> | password <string>} | sha-1 {key <20-byte hex> | password <string>}]] [ah {md5 {key <16-byte hex> | password <string>} | sha-1 {key <20-byte hex> | password <string>}]

set vpn <vpn_name> monitor

AutoKey IKE

set ike p1-proposal <name> {preshare | rsa-sig} {group1 | group2 | group5} esp {des | 3des} {md5 | sha-1} [{seconds | minutes | hours | days} <lifetime>]

set ike p2-proposal <name> {no-pfs | group1 | group2 | group5} {ah | {esp {null | des | 3des}}} {null | md5 | sha-1} [{seconds | minutes | hours | days} <lifetime>] [kbytes <lifesize>]

set ike gateway <name> {ip <peer_ip> [id <peer_id>] | dialup <user_name>} {main | aggressive} [preshare <preshare_key>] proposal <p1_proposal>

set ike {accept-all-proposal | id-mode {ip | subnet} | policy-checking}

Check Your Address Book Entry

Check the address book entry for the remote LAN. If you have not created such an entry, do so now.

WebUI

Address >>Untrusted: Edit your previously made entry, or click **New Address**, and follow these guidelines:

Address Name: Type an easily identifiable name for the address of the remote LAN.

IP Address: If the VPN is to allow access to the entire remote subnet, then this should be a network address of the remote Trusted LAN (for example, 192.168.10.0).

Netmask: If the VPN is to allow access to the entire remote subnet, then this should be the subnet mask used on the remote LAN (for example, 255.255.255.0).

Comment: You can put whatever you want here.

Location: Untrusted

CLI

Use the following **set address** and **get address** commands to create or check the address book entry for the remote LAN (and refer to the information in the previous WebUI section):

set address {trust | untrust | dmz} <address_name> {<a.b.c.d> <A.B.C.D> | <domain name>} [<comment>]

get address [all | dmz | trust | untrust]

> ***Note:*** *If you want all outbound traffic to use the VPN tunnel, do not worry about this address book definition. You can simply use* ***Outside Any*** *as the destination in the Access Policy (as described in the next section).*

Check your Outgoing Access Policy

Check the outgoing Access Policy. (When creating a VPN, incoming Access Policies are not needed.)

WebUI

Policy >> Outgoing >> Detail: Use the following guidelines:

Source Address: An address book definition that you have defined to include the subnet you want to give access to the VPN (typically “inside any”).

Destination Address: This is the address book definition that was defined above.

Action: Encrypt for VPN policies.

VPN Tunnel: This is the tunnel name that was created.

The remaining fields are optional

CLI

Use the following **get policy** command to view the desired Access Policy (and refer to the information in the preceding WebUI section):

get policy outgoing <number>

Check Access Policy Order

One easily overlooked detail is that the order of policies matters. For example:

set policy outgoing “inside any” “outside any” any permit

set policy outgoing “inside any” “outside any” any encrypt “VPN to NY”

Given the order of these two policies, the second Access Policy will never work. The first Access Policy will be appropriate for any type of outgoing traffic. Use the Up and Down arrow icons under the Configure column in the web browser management tool to move the encrypted Access Policy (the second Access Policy) to the top of the list.

Check the Surrounding Network

Just because the NetScreen devices on both sides of the tunnel are configured correctly does not mean that the tunnel will work. There are many other things that can cause the VPN to function incorrectly, including:

- Local hosts have the wrong gateway.

  The local hosts need to use the NetScreen device as a gateway, or the NetScreen device needs to be the final way out. To the hosts (clients) on the LAN a VPN is simply just another hop to another network. If you are using another Internal router as the gateway for all hosts, make sure there is at least a route to the NetScreen device for the remote networks.

- Both sides of the VPN are using the same network address.

  If both sites have the same subnet network, the packet will naturally try to traverse to a device on the local LAN. It will not see a device on that network, and pings will return with a series of time outs. It will not go through the NetScreen device or the VPN tunnel.

- Untrusted Address is not a public address.

  One very important point is that the Untrusted address of the NetScreen device needs to be a public one. This is so that the remote NetScreen can send packets to it. In other words, the upstream (Internet) router must not be doing NAT.

- Upstream (Internet) Router or ISP cannot block protocol 50.

  If the router or the ISP is blocking or restricting the use of Protocol 50 (or IPSec, the protocol that drives VPN) the VPN will not work. Check the router filters in both directions. The packet may be allowed in one direction, but may be filtered in the other direction.

Using the Debugger

One of the last resorts is to use the Command Line Interface Debugger to gather more information. Follow the below guidelines to use the debugger as it pertains to VPN.

1. Turn on the debug buffer. This allows you to view the captured information at your leisure after it is captured. The command to do this is:

   set console dbuf

2. Turn on the debugger. The command to do this is:

   debug flow basic

3. Send some data through the tunnel. The best type of data to send is a ping from one side of the tunnel to the Trusted port of the remote NetScreen. Allow at least four pings to fail.
4. Turn off the debugger. The data has been captured to the debug buffer so you can stop the debugger.

   undebug all

5. Examine the contents of the debug buffer.

   get dbuf stream

The following is an example of a capture of data from a ping that was never returned. It looks as though the packet was sent out correctly.

```
01165.0:   => 0010db00b2f1
01165.0:   packet is encrypted
01165.0:   Send to untrust
01166.0:   trust:10.10.10.10/768->172.16.10.10/58883,1
01166.0:   route 172.16.10.10->192.168.5.2, to untrust
01166.0:   search acl set 0
01166.0:   resolve 192.168.5.2 on untrust if
01166.0:   => 0010db00b2f1
01166.0:   packet is encrypted
01166.0:   Send to untrust
01167.0:   trust:10.10.10.10/1024->172.16.10.10/58883,1
01167.0:   route 172.16.10.10->192.168.5.2, to untrust
01167.0:   search acl set 0
01167.0:   resolve 192.168.5.2 on untrust if
01167.0:   => 0010db00b2f1
01167.0:   packet is encrypted
01167.0:   Send to untrust
01168.0:   trust:10.10.10.10/1280->172.16.10.10/58883,1
01168.0:   route 172.16.10.10->192.168.5.2, to untrust
01168.0:   search acl set 0
```

The following is a copy of the debug flow output from a remote NetScreen during a ping.

```
00852.0:   trust:172.16.10.144/1481->172.16.10.10/23,6

00852.0:   copy a packet to stack

00861.0:   Dec: SPI=00003001, Data=136

00861.0:   packet dropped, SA not found
```

Note the line that says:

```
Packet dropped, SA not found
```

The line before this identifies the `SPI=00003001` as being the culprit. There is a problem with Security Indexes in this example. If you see this type of message, you should double check the SPI settings in the VPN configuration.

Suppose the second NetScreen had nothing in the debug buffer. What could be determined from this? We can definitely determine that it didn't get any packets. Possible problems include (but are not limited to):

- Wrong VPN gateway
- Router blocking protocol 50
- ISP blocking protocol 50

If you have difficulty deciphering the debug flow output, please send it to support along with the information outlined in "Contacting Technical Support" on page 10-45.

Peer-to-Peer VPN Application Troubleshooting

A VPN must be able to run required applications through the tunnel. This is why IPSec is fast becoming a popular protocol. It is very flexible and allows for any IP application (virtually) to be run over it. However, IPSec can introduce some difficulties with some applications. These difficulties stem from several misconceptions in the area of packet sizes and fragmentation.

Packet Size and Fragmentation

Ethernet has a maximum packet size of 1518 bytes. Of this, 18 bytes are used for the MAC header and trailer, and 8 bytes may be used for the Logical Link Control (LLC) and Sub Network Access Protocol (SNAP) header. This leaves 1492 bytes available for the IP datagram (packet).

In a typical wide area network link, the MTU (maximum transmit unit or maximum packet size) is about 1006 bytes. If we try to send an Ethernet packet through this link (1492 bytes, after the Ethernet headers are stripped into a 1006 byte transmit unit), the Ethernet packet will not fit. The routing device must fragment the Ethernet packet into two packets.

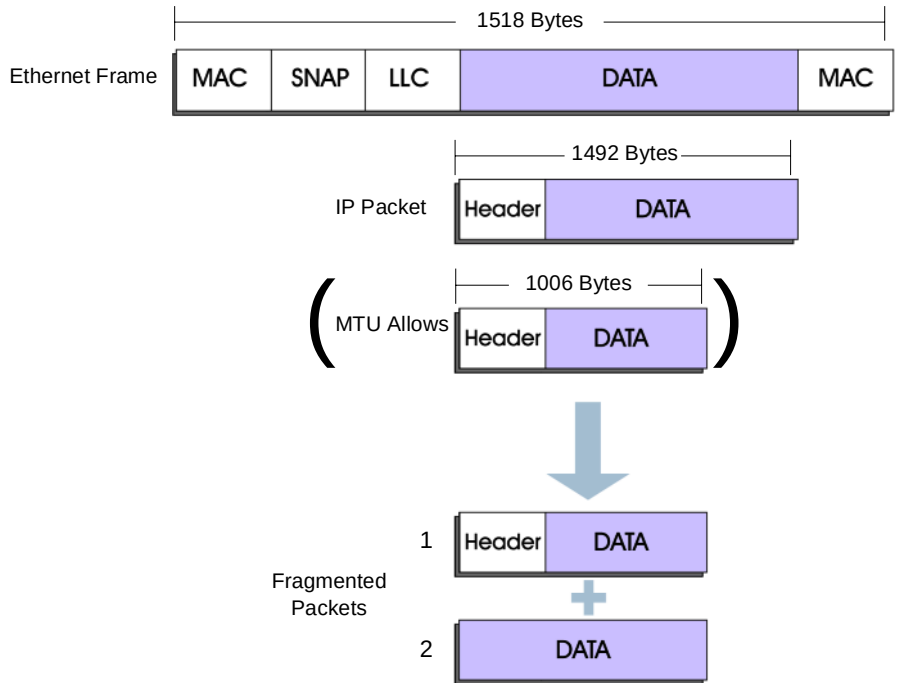

Because IPSec encrypts the entire packet, it adds an additional IP header on the packet (about another 40 bytes). Some applications use the entire Ethernet packet size to start off with. Because VPNs rely on encryption for security, adding encryption headers to an already too small packet payload capacity guarantees that there will always be fragmentation over a VPN connection.

Consider some of the pitfalls of this fragmentation scheme. If one of the two fragments is not received, the sender will resend the entire initial packet. This, again, results in the fragmentation of the packet. If for some reason, the receiver always drops the second packet, the session will find itself in an endless loop. At any rate, dropped fragments lead to retransmission of the original packet, leading to more fragmentation. This of course leads to more overhead because each packet has a header.

A Common Misconception

Other problems occur as a result of common misconceptions regarding VPNs. The main misconception is that VPNs compete at leased-line speed. However, VPNs do not offer the performance that a leased line offers. This is because the transport medium is the Internet. Even with a T1 line (1.54 MB/Sec) you will not necessarily get T1 performance between your offices. It is widely understood that the Internet can introduce high latency and is often the bottleneck. When you purchase a T1 line from your ISP, you are guaranteed (depending upon your service) a certain level of performance. With an IPSec tunnel through the Internet, there are no guarantees. There isn't even a guarantee that the connection from end-to-end will always be there.

Therefore, bandwidth-intensive applications that worked when your offices were connected through a leased T1 line may not work as well through the VPN. While IPSec offers flexibility and value, performance is only as good as the performance of the hops along the way.

Configuration Check

Since the tunnel is working, there is little checking that needs to be done. There are some things that can be done to help alleviate some of the common problems with some of the applications.

To eliminate fragmentation for TCP communications through the tunnel, issue the following CLI commands:

```
set flow tcp-mss

save
```

During the three-way handshake, the two devices on the endpoints of the connection (server and client) negotiate a maximum packet size. When the command above is used, the NetScreen device changes the maximum segment size (MSS) to reflect the change in the MTU created by the application. This modification makes sure no fragmentation occurs. This is especially important with applications that require a do not fragment (DF) bit set in the IP packet. The most common applications helped this command are PC Anywhere or an application that uses encryption (such as Secure Shell or SSH).

Check the Surrounding Network

When an application is not working properly over the VPN, verify that:

- The server is working properly by testing it locally.
- That the server or the client has the NetScreen device as a gateway, otherwise the packets will not make it across the network.
- That the server and application are configured correctly. It is a good idea to double check the configuration of the application itself before assuming the VPN is the culprit. Also, check the application vendor Website for information on running the application through a VPN.

Using the Debugger

Follow the debug instructions in “Using the Debugger” on page 10-15.

Client (NetScreen-Remote)-to-Device Troubleshooting

The main difference between a peer-to-peer VPN and a client-to-device VPN, is the client-to-device VPN is uni-directional. That is, the only direction the VPN will work is from the client to the NetScreen device. Since the client is often connected to the Internet via a dial-up (modem) connection, the IP address of the client can be different every time a connection is made. When the Access Policy is created on the NetScreen device, the Destination IP is set to **Dial-up User**. Since the device does not know the address of the client, it does not know where to send the encrypted packets. However, the configuration of the client does include the address of the NetScreen device and therefore it can send encrypted packets through the tunnel.

Checking the Configuration

When the NetScreen Remote-to-Device VPN is not working, check your configuration using the recommended settings described below.

WebUI

Users >> Users >>:Verify the following settings:

Make sure you have selected **Dialup VPN User** instead of **Authentication User**.

User Group: This is optional. However, if you add a group under the Dial-Up Group tab you can select the group here. Multiple users can belong to one group. This allows for the creation of one Outgoing VPN Access Policy for multiple users.

Security Index (local and remote): These can be any hexadecimal value between 3000 and 8000.

Remember that the remote NetScreen should use the reverse values when configuring the client. For example, if you used 3000 for the local on one side, you need to use 3000 for the Outgoing Keys on the NetScreen Remote client.

Also, if you are going to have multiple tunnels from one device, make sure that the Security Index for each tunnel is unique. These are tunnel identifiers and need to be distinct for each tunnel.

ESP-Encryption Algorithm: Select the type of encryption that you want to use. If your top priority is security, use 3DES-CBC (each packet encrypted 3 times). If your top priority is performance, use 56bit DES-CBC. Whatever you use here must be used on the remote side as well.

HEX Key and Generate Key By Password: A key is a hexadecimal value that is used to encrypt and decrypt the datagrams. Whether you enter a HEX Key or you use

the Generated Key By Password is arbitrary. NetScreen always suggests using the Generated Key By Password. Just remember that the HEX value that you use (or generate) on the device must be used on the client as well.

ESP-Authentication Algorithm: Select the type of authentication you want. This part of the VPN tunnel is optional. However, NetScreen suggests that you use this. Authentication adds another layer of security to the VPN in that it ensures that none of the packets are tampered with in transit. The most common selection here is MD5. SHA-1 is a more powerful form of authentication but can sometimes cause problems with some applications. If security is a concern, choose SHA-1.You choice must be the same on the client as well.

HEX Key and Generate Key By Password: Follow the guidelines outlined above.

CLI

Use the following **get user** command to view all user profiles or a specific profile (and refer to the information in the preceding WebUI section):

get user [all | id <number>]

Check your Outgoing Access Policy

When you create a VPN, you do not need Incoming policies. You do however, need an Outgoing Access Policy.

WebUI

Policy >> Outgoing >> Detail: Verify the following settings:

Source Address: Inside Any (or the address book definition that you defined to include the subnet you want to give access to the VPN).

Destination Address: Dial-Up VPN.

Service: Any.

Action: Encrypt.

VPN Tunnel: Select the name that you created when you set up the dialup user.

CLI

Use the following **get policy** command to view the desired outgoing Access Policy (and refer to the information in the preceding WebUI section):

get policy outgoing <number>

Check your VPN Client Settings

Check the configuration setting on the NetScreen Remote Client, using the following guidelines.

WebUI

1. NetScreen Remote >> Security Policy Editor:

   Check that a new connection has been added. If the only entry is **Other Connections**, then a secure connection for the connection to the NetScreen device has not been created. In this case, click File >> New Connection to create the connection. The name you give it is arbitrary but should be something descriptive like "VPN to London."

2. Security Policy Editor >> New Connection: Check the following settings:

   Connection Security: Secure

   Remote Party Identity and Addressing: The selection you make here depends on what you are trying to access on the Trusted side of the NetScreen device. Most users typically want to access the entire subnet. In that case, select **IP Subnet** for the ID Type. Then enter the subnet address and mask (for example, 192.168.10.0; 255.255.255.0).

   Protocol: All

Connect using Secure Gateway Tunnel: (select) For ID Type, select **IP Address**. This is where the client sends the encrypted packets. The address you enter must be the Untrusted IP address of the NetScreen device.

3. Security Policy Editor >> New Connection: Double-click the **New Connection** icon to expand it.

   The My Identity and Security Policy icons appear.

4. Security Policy Editor >> New Connection >> My Identity: Check the following settings:

   Internet Interface Name: Select the adapter used for the encrypted connection. If you are unsure of this name, select **Any**.

5. Security Policy Editor >> New Connection >> Security Policy: Select **Use Manual Keys** for the Select Phase 1 Negotiation Mode.
6. Security Policy Editor >> New Connection >> Security Policy: Double-click the **Security Policy** icon to expand it.
7. Security Policy Editor >> New Connection >> Security Policy >> Key Exchange (Phase 2) >> Proposal 1: Check the following settings:

   SA Life: Unspecified

   Encapsulation Protocol: (select)

   Encrypt Alg: Select the same algorithm as that on the NetScreen device for this user.

   Hash Alg: Select the same algorithm as that on the NetScreen device for this user.

   Encapsulation: Tunnel

8. Security Policy Editor >> New Connection >> Security Policy >> Key Exchange (Phase 2) >> Proposal 1 >> Inbound Keys: Check the following settings:

   Make sure the Security Parameters Index is the same as that on the NetScreen device for this user. It should match the Security Index marked **Remote**.

   Make sure the ESP Encryption Key and the ESP Authentication Key are the same as those on the NetScreen device for this user. Enter binary values—not ASCII values. The algorithm used to generate the

keys in NetScreen device is not the same as the algorithm on the client. Therefore, you must copy the hexadecimal keys directly.

9. Security Policy Editor >> New Connection >> Security Policy >> Key Exchange (Phase 2) >> Proposal 1 >> Outbound Keys: Check the following settings:

   Make sure the Security Parameters Index is the same as that on the NetScreen device for this user. It should match the Security Index marked Remote.

   Make sure the **ESP Encryption Key** and the **ESP Authentication Key** are the same as those on the NetScreen device for this user. Enter binary values—not ASCII values. The algorithm used to generate the keys in NetScreen device is not the same as the algorithm on the client. Therefore, you must copy the hexadecimal keys directly.

Checking the Surrounding Network

Check that the client has a secondary adapter with an address that resides on the destination LAN. In other words, if you set up the Remote Party Identity and Addressing to be an IP Subnet of 192.168.10.0/255.255.255.0 and the address of your NIC card is 192.168.10.50, the VPN will not work because Windows will try to go out of the NIC card. This will remain true even if the NIC card is not connected and seems disabled.

A good way to test the client is by pinging the address of the Trusted port on the NetScreen device from the client. If you can do this, but cannot ping any addresses on your Trusted LAN, then those machines probably are not using the NetScreen device as their gateway. For information on using the ping utility, see “Cannot Ping Between Unsecure Hosts and Secure Hosts” on page 10-43.

Using the Debugger

Follow the guidelines described in “Using the Debugger” on page 10-15.

Troubleshooting Mapped IP or Virtual IP Addresses

Mapped IPs (MIP) and Virtual IPs (VIP) are used to allow incoming access through a NetScreen device that is in Network Address Translation (NAT) mode. The basic premise is that you assign an Untrusted, public address to an internal (Trust or DMZ) address.

MIPs are used for one-to-one mappings. In other words, one Outside IP is mapped to one Internal IP.

VIPs have a dual purpose. One purposes is to do port mapping which means you can map one Untrusted IP to a server on the inside. You can tell the NetScreen device that the port that the outside users will access is any port that you want (user specified) and have that request forwarded to the internal server on the correct port by picking the service from the list. This is also known as port mapping. You can configure up to eight servers (each having a service mapped to a different port) per VIP.

On the NetScreen-100, VIPs are also used for load balancing which allows you to define a VIP and add up to eight servers to that VIP, each for the same service. You then pick one of the load balancing algorithms from the pull down menu. The NetScreen device will then load balance the access to this VIP between the servers you specified. See Chapter 7, “Traffic Shaping” for more information on load balancing.

Checking the MIP Configuration

Check the configuration of the problematic MIP by clicking its **Edit** option. Check the parameters according to the following guidelines.

WebUI

1. Virtual IP >> Mapped IP >> Edit: Check the following settings:

   Untrusted IP: This address should be the public address that you wish to have the Outside (Internet) users use to access your server. This cannot be the same address as the Untrusted Interface.

   Network Netmask: This netmask should remain 32 bits (or 255.255.255.255). This value is not the netmask of the Untrusted, DMZ, or Trusted side. Rather, this is the netmask that defines the relationship between the Untrusted IP and the Map to Address. Since we are mapping one outside IP to one inside IP, this should remain 32 bits.

   Map to Address: This is the actual address of the internal server.

2. Policy >> Incoming >> Detail: Edit the Access Policy associated with the problematic MIP and verify the settings according to the following settings:

   Source Address: Should be **Outside Any**.

   Destination Address: Typically an address like MIP(x.x.x.x), where x.x.x.x is the MIP in question.

   Service: Should be whatever you plan to allow. If you are not sure if the MIP is working correctly, consider letting in **ANY** service, then test that you can pass data through the device.

   Action: Typically **PERMIT**. For more information, see the authentication option in the “Troubleshooting Access Policies” on page 10-41.

Everything else: The remainder of the fields are optional. For more information, see the authentication option in the "Troubleshooting Access Policies" on page 10-41.

Note: *Even if the MIP is defined to the DMZ, the Access Policy for the MIP will show up in the Incoming Access Policy tab.*

CLI

Use the **get mip** and **get policy incoming** commands to check the MIP configuration parameters (and refer to the information in the preceding WebUI section).

Checking the VIP (Port Mapping) Configuration

Use these guidelines to check the configuration of the VIP when it is used for port mapping.

WebUI

Virtual IP >> Virtual IP*n* >> Edit: Verify the configuration of the VIP that is causing difficulty by using the following guidelines:

Virtual Server IP: This value should be an available Untrusted IP. In other words, this address should be an address that is not being used by anything on the Untrusted side, including the Untrusted port itself.

Service: The service that you want to be returned to the user when they access the VIP. For example, HTTP, FTP, etc.

Virtual Port: This is the port that the user will use to access the server. For example, suppose you used port 1000 and the service is HTTP. Also, assume the Virtual Server IP address is 100.100.100.100. The user will need to enter the following in their browser to access the service: http://100.100.100.100:1000

Status: This is an important field. This tells the administrator if the server is available. Basically, if the server is not responding to the NetScreen device ICMP checks, it will be deemed as down and “Not Available” will show up in this field.

Actual IP: The actual address of the server.

Actual Port: The actual port that the server is listening on for the service specified. If this is the incorrect port, then you will need to create a customer service to allow the access to the port you need.

CLI

Use the **get vip [server | session]** command to check the VIP configuration parameters (and refer to the information in the preceding WebUI section).

Checking the Load Balancing VIP Configuration

> ***Note:*** *This section only applies to the NetScreen-100.*

The NetScreen-100 can support up to four Virtual IP (VIP) addresses. In other words, up to four groups of load balanced servers can be utilized. For each of these VIPs there can be up to eight servers whose processing loads are being balanced. Also for each VIP, up to six services (HTTP, HTTPS, TELNET, FTP, MAIL, and POP3) can be utilized.

Use the following steps to check your VIP configuration from the NetScreen device web management tool. Note that the NetScreen device must be in NAT mode to use this feature. The VIP option is not present in Transparent Mode.

WebUI

Virtual IP >> Virtual IP(number) >> Edit: Check the parameters with the following guidelines:

Virtual Server IP: This value should be an available Untrusted IP. In other words, this address should be an address that is not being used by anything on the Untrusted side, including the Untrusted port itself.

Virtual Port: This is the port that the user will use to access the load balancing pool. For example, suppose you used port 1050 and the service is FTP. Also, assume the Virtual Server IP address is 100.100.100.100. The user must enter the following in their browser to access FTP: http://100.100.100.100:1050.

Service: This service defines the port that the Virtual Port is translated to on the DMZ or the Trusted side.

Load Balance: This is the method that will be used for load balancing (Round Robin, Weighted Round Robin, Least Connections, Weighted Least Connections). Refer to Chapter 7, "Traffic Shaping" for more information on these algorithms.

Server IP: The actual address of the server. For load balancing you will need to add at least two servers to the list.

CLI

Use the **get vip [server | session]** command to check the VIP configuration parameters (and refer to the information in the preceding WebUI section).

Checking the Incoming Access Policy for VIP Configuration

Check that the Incoming Access Policy is defined correctly. Go to the Access Policy page (Incoming tab) and click on the edit button for the Access Policy that is associated with the MIP having difficulty.

Note that even if the MIP is defined to the DMZ, the Access Policy for the MIP will show up in the Incoming Access Policy tab. Use the following discussion to troubleshoot the problematic MIP.

WebUI

Policy >> Incoming >> Detail: Compare the configuration of the Access Policy associated with the problematic MIP with the following guidelines:

Source Address: Outside Any

Destination Address: All Virtual IPs

Service: Select whatever you plan to allow. However, if your VIP is only allowing HTTP, an Access Policy with FTP as the service will be ignored.

Action: Permit.

All the other fields are optional.

CLI

Use the following **get policy** command to check the incoming Access Policy configuration parameters (and refer to the information in the preceding WebUI section):

get policy [all | incoming | outgoing| todmz| fromdmz | <number>]

Checking the Surrounding Network

With MIP/VIP problems, there are several things to check for regarding the surrounding network. The more common mistakes are:

- Servers not using the correct gateway—If the server that you are mapping to is not using the NetScreen device as a gateway, then the packets that get to the server will never be sent back. In other words, the server needs to know where to send the data when it receives it.
- Server Not Functioning Properly—If the MIP or VIP is not functioning properly, then the problem may be that the server itself is not functioning properly. If you suspect that this is the case, simply test the service on the server locally.

- Wrong IP Address—Sometimes the problem is simply that the server has the incorrect address. If this is the case, the status field in the VIP configuration shows **Not Available**.
- Check that the server can get out—One way to make sure that the server is correctly configured is to ping a server on the Internet from it. If it can get out, then chances are that it is configured correctly. However, for the MIP/VIP to work properly, the traffic must also be going through the NetScreen device out to the Internet.

Using the Debugger

Follow the steps below to use the debugger when a MIP or VIP is not working properly.

1. Turn on the debug buffer. This allows you to view the captured information at your leisure after it is captured. The command to do this is:

   set console dbuf

2. Set the value of the `ffilter`. This will narrow the results of the debug down to just the information that we are looking for.

   set ffilter src-ip <a.b.c.d>

   Where `<a.b.c.d>` is the address of the host on the Internet (or on the Untrusted side of the NetScreen device) that will be doing the test.

3. Turn on the debugger. To do this issue the following command:

   debug flow basic

4. Send some data (from the host specified in the **`set ffilter`** command above) to the MIP/VIP.
5. Turn off the debugger. The data has been captured to the debug buffer so we can stop the debugger by issuing the following command:.

   undebug flow basic

6. Examine the contents of the debug buffer by issuing the following command.

   get dbuf stream

7. Forward the results of this exercise along with the information listed in “Contacting Technical Support” on page 10-45.

Outbound Access Troubleshooting

The most common use of the NetScreen device is to allow outbound access to the Internet for the corporate LAN while providing protection to the LAN from outside, Internet users. The NetScreen device can do this in one of two modes: NAT or Transparent mode.

Checking NAT Mode Configuration

Use the following guidelines to check the interface addresses of the NetScreen device in NAT mode.

WebUI

Interface >> Trust: Verify your settings:

Trusted Interface: The inside IP should be the address that your inside (Trusted) hosts use as their gateway. The NetMask should be the mask of the internal LAN. The Traffic Bandwidth can remain 0. This means that all available bandwidth is used. When troubleshooting outbound access, this is the preferable setting. The Default Gateway is only used when you have users that are one or more hops away on the Trusted side. In other words, you have a router on your Trusted side and the hosts on the other side of that router use the NetScreen device for outbound access.

However, for the default gateway to work properly, IP Spoofing detection must be turned off.

Configure >> General Page: Verify your settings:

Detect IP Spoofing Attack: (cleared).

If you require that IP Spoofing remain enabled, then use the route table to add a static route instead.

Configure >> Route Table: Enter the following, and then click **OK**:

Network Address: The address of a specific subnet on the Trusted network.

Network NetMask: The subnet mask for that subnet.

Gateway IP Address: The internal router leading from the NetScreen device to the specified subnet.

Interface: Trusted.

Note: *When a NetScreen device is operating in NAT mode, the source IP address and source port number on outbound packets are changed as follows:*

- *The source IP address is changed to the IP address of the interface that the packet is leaving—either the Untrusted or DMZ interface.*
- *The NetScreen device changes the source port number to a unique number. The source port number is the determining factor by which the NetScreen device can route a returning packet to its proper destination.*

Interface >> Untrust: Verify your settings:

The Outside IP needs to be an address given to you by your ISP. It is typically a public IP address. The NetMask is the mask that your ISP gives you as well. The Default Gateway is the address of the router to the Internet. In some situations (DSL, cable modem) this router will not be at your location. Simply use whatever the ISP tells you to in this case. The Traffic Bandwidth can remain as 0. This means that all available bandwidth is used. When troubleshooting outbound access, this is the preferable setting.

Interface >> DMZ: Verify your settings

The DMZ IP should be the address that the servers on the DMZ will use for their gateway. The NetMask is the mask that is used on the DMZ LAN. The Traffic Bandwidth can remain as 0. This means that all available bandwidth is used. When troubleshooting outbound access, this is the preferable setting.

Note: *The DMZ Interface plays a special role when it comes to the NetScreen device data flow. Unlike data coming from the Trusted port, data coming from the DMZ is not NATed by default. By default, data coming from the DMZ LAN and going out the Untrusted (or Trusted for that matter) is routed without being changed. In order for NATing to happen, a Mapped IP needs to be created between an available Untrusted address and the DMZ server. The reasoning for this difference is that the DMZ is designed for servers. Since access to servers is given to the Outside Internet users, MIPs or public IP addresses will be necessary for their use.*

Configure >> Web: Verify your settings

This is the address of the NetScreen device's Management Tools. In other words, if you Telnet (or browse) to the NetScreen device, you will browse to this address. If this address is set to 0.0.0.0 then you access the management tool using the address of the Trusted (and/or Untrusted IPs). This address plays no role in the data path of the NetScreen device.

CLI

Use the following CLI commands to verify your settings (and refer to the preceding WebUI section):

get interface [dmz | trust | untrust]

get firewall

get route [ip <a.b.c.d>}

get admin

Checking the Policies

The two types of policies most likely to cause problems are Outgoing and From DMZ.

For outgoing policies, remember that the order of the policies is very important. Take the following two policies as an example.

set policy outgoing “inside any” “outside any” any deny

set policy outgoing “inside any” “outside any” any permit

This is a rather extreme example because the Access Policies are offsets of each other. The first Access Policy will block all data on the way out. The second Access Policy has no effect.

WebUI

Policy >> Outgoing >> Detail: Check the parameters of the Access Policies using the following guidelines:

Source Address: Because you are allowing access to the Internet for all your Internal users, this is typically set to Inside Any. However, you may wish to narrow this down by defining an address in the Address Book for the Local users to whom you wish to give access. These address book entries (when defined on the Trusted LAN) will show up in the Source Address pull down menu.

Destination Address: Outside Any.

Service: Typically Any. However, if you wish to limit your users to certain services, you can have several policies, each with a different service that you wish to let through. By default, the NetScreen device blocks everything. You only need to allow the services you want. You do not need to Deny anything.

Action: Permit.

The remainder of the parameters are optional. See the "Access Policy Troubleshooting" section for more information.

CLI

Use the **get policy outgoing** command to check the parameters of the Access Policies (and refer to the preceding WebUI section).

Check the Physical Connections

A common error is to use the wrong cable or incorrectly connect the NetScreen device to your Network. Here is a general overview of what to look for:

- **Trusted Port:** The Trust Port should be connected to the hub/switch of your Internal LAN. The green link light should be lit on the port. If this is not the case, try a different cable, a different port, or a differently wired cable (cross over or straight through).
- **Untrusted Port:** Typically, the Untrusted port is connected directly to the upstream router (or DSL/Cable modem). Again, the green link light should be lit. If this is not the case, try a different cable, a different port, or a differently wired cable (cross over or straight through).
- **DMZ Port:** The Trust Port should be connected to the hub/switch of your DMZ LAN. The green link light should be lit on the port. If this is not the case then try a different cable, a different port, or a differently wired cable (cross over or straight through).

Checking the Surrounding Network – NAT Mode

The health and configuration of the surrounding network components plays a vital role in the health of the Outbound connection. Check the following things to make sure everything is working correctly.

- Router/Internet Connection—One of the most common problems is that the Internet connection is not working properly. The easiest way to test this is to see is a host on the Trusted LAN can ping the Internet router. If it can, then the data is properly traversing the NetScreen device and the problem lies beyond the router. You can also put a PC on the Untrusted LAN, giving it the same address as the NetScreen device's Untrusted port and same gateway as the Untrusted port. If that PC also has trouble getting to the Internet, then the problem is not with the NetScreen device, but rather with the Internet connection. Contact your ISP for assistance.
- Network Settings Incorrect on the Clients—Another possibility is that the clients on the Trusted LAN are configured improperly. They should be on the same subnet as the NetScreen device's Trusted port and should be using the NetScreen device's Trusted port as their gateway. (This, of course, will not be true if the clients are one or more hops away on the Trusted side.) Another possibility is that the DNS server is incorrectly configured. This would have the effect of allowing access to IP addresses, but not to fully qualified names (for example, www.netscreen.com). You can test this by pinging the upstream router (Internet router) from the client.

Checking the Configuration—Transparent Mode

The only difference between NAT mode and Transparent mode is that there are no addresses on the Interfaces and the policies need (or at least should be) to be more specific in Transparent mode. Also, the DMZ port is not used in Transparent Mode. Check the Outgoing policies according to the following rules:

For outgoing policies, remember that the order of the policies is very important. Take the following two policies as an example.

set policy outgoing “inside any” “outside any” any deny

set policy outgoing “inside any” “outside any” any permit

This is a rather extreme example because the policies are offsets of each other. The first Access Policy will block all data on the way out. The second Access Policy has no effect.

WebUI

Policy >> Outgoing >> Detail: Check the parameters of the policies using the following guidelines:

Source Address: Since you are allowing access to the Internet for all your Internal users, this is typically set to Inside Any. However, you should narrow this down by defining an address in the Address Book for the Local LAN. For example, you can define an address book entry for 192.168.10/255.255.255.0. These address book entries (when defined on the Trusted LAN) will show up in the Source Address pull down menu.

Destination Address: Outside Any.

Service: Typically Any. However, if you wish to limit your users to certain services, you can have several policies, each with a different service that you wish to let through. By default, the NetScreen device blocks everything. You only need to allow the services you want. You will not need to Deny anything.

Action: Permit.

The rest of the parameters are optional.

CLI

Use the **get policy outgoing** command to check the parameters of the Access Policies (and refer to the preceding WebUI section).

Checking the Surrounding Network – NAT Mode

The health and configuration of the surrounding network components plays a vital role in the health of the Outbound connection. Check the following things to make sure everything is working correctly.

- Router/Internet Connection—One of the most common problems is that the Internet connection is not working properly. The easiest way to test this is to see is a host on the Trusted LAN can ping the Internet router. If it can, then the data is properly traversing the NetScreen device and the problem lies beyond the router. You can also put a PC on the Untrusted LAN, giving it the same address as the NetScreen device's Untrusted port and same gateway as the Untrusted port. If that PC also has trouble getting to the Internet, then the problem is not with the NetScreen device, but rather with the Internet connection. Contact your ISP for assistance.
- Network Settings Incorrect on the Clients—Another possibility is that the clients on the Trusted LAN are configured improperly. They should be on the same subnet as the Untrusted Side router. (This, of course, will not be true if the clients are one or more hops away on the Trusted side.) Another possibility is that the DNS server is incorrectly configured. This would have the effect of allowing access to IP addresses, but not to fully qualified names (like www.netscreen.com). You can test this by pinging the upstream router (Internet router) from the client. The clients should have a gateway as the Untrusted side (Internet) router.

Using the Debugger

Follow the steps below to use the debugger when outbound access is not working properly.

1. Turn on the debug buffer. This allows you to view the captured information at your leisure after it is captured. The command to do this is:

   set console dbuf

2. Set the value of the `ffilter`. This will narrow the results of the debug down to just the information that we are looking for.

   set ffilter src-ip <a.b.c.d>

   Where <a.b.c.d> is the address of the host on the Internet (or on the Untrusted side of the NetScreen device) that will be doing the test.

3. Turn on the debugger. To do this issue the following command:

   debug flow basic

4. Send some data (from the host specified in the **`set ffilter`** command above) to the MIP/VIP.

5. Turn off the debugger. The data has been captured to the debug buffer so you can stop the debugger.

   undebug flow basic

6. Examine the contents of the debug buffer.

   get dbuf stream

   Forward the results of this exercise along with the information listed in“Contacting Technical Support” on page 10-45 or go to support@NetScreen.com.

Troubleshooting Access Policies

This section focusses on creating a proper and efficient access control list (ACL).

Access Policy Order Matters

The order of the policies is very important. The general rule is that the first Access Policy that fits will be used. The NetScreen device will not look further down the list. What does “fits” mean? Consider the following:

- Source IP: If the Access Policy says **Inside Any**, then any packet coming from the Trusted side will fit this parameter. Of course, if this parameter is more specific (through an address book definition), say a particular host, then only that host will fit the parameter.
- Destination IP: This is the same situation as that of the Source IP. **Outside Any** will fit all clients, while a more specific definition will not.
- Service: The last parameter that matters in the “Does this Access Policy fit this packet” question is Service. If the packet matches the Source IP and Destination IP, it still needs to match the service (based on source and destination port) before the Access Policy will be applied to the packet.

Example 1

Assume that an address book entry called “Bob’s Computer” has been defined for host 192.168.10.55.

Lets assume that we have the following Outgoing Policies:

1. set policy outgoing “inside any” “outside any” any permit
2. set policy outgoing “Bob’s Computer” “outside any” any deny

Bob’s Computer is permitted Internet access because the first Access Policy that applies to it is Access Policy 1. This Access Policy allows Bob’s packets through. The correct order of the policies is:

1. set policy outgoing “Bob’s Computer” “outside any” any deny
2. set policy outgoing “inside any” “outside any” any permit

This brings us to our second rule. The more specific the Access Policy, the higher on the ACL it should be. Only packets with a source IP of 192.168.10.55 fit the first Access Policy. The second Access Policy, allowing Internet access, applies to all other addresses.

Example 2

Still using “Bob’s Computer” as an entry in the Address book, assume we have the following ACL:

1. set policy outgoing “Bob’s Computer” “outside any” FTP deny
2. set policy outgoing “inside Any” “outside any” any permit

Is Bob allowed to get to http://www.netscreen.com?

The answer is yes. We are only denying FTP traffic for Bob. All others services are permitted able to traverse the NetScreen device.

Example 3

Assume that we have a dial-up VPN user named “Bob” and that we have the following Access Policies:

1. set policy outgoing “inside Any” “Outside Any” “Any” Permit
2. set policy outgoing “inside Any” “Dial-Up User” “Any” Encrypt “VPN User – Bob”

Can the dial-up user access the resources on the LAN?

The answer is no. The packets that are returned to the dial-up user are sent out in the clear because Access Policy 1 applies to all packets. The order of these policies should be reversed.

Logging

When you check the logging option, the NetScreen device keeps track of all uses of the Access Policy. The logging feature can be used to check if a Access Policy is being used when you think it is. You can check the log by clicking on the log icon in the ACL summary.

Counting

You can also utilize counting to find out how much bandwidth a particular Access Policy is taking up. Once you click on the hour glass icon in the Access Policy Summary Screen, you can look at the amount of bandwidth used in the last second, minute, hour, day, or month. The results of the counting is shown in a graph.

Alarm Threshold

In addition to counting, the alarm threshold can be used to determine the amount of bandwidth being used by a particular Access Policy. You can either set a bytes/sec threshold or bytes/minute threshold. If this threshold is exceeded, an alarm is generated.

Schedule

An Access Policy can be set to work on a schedule. The advantage of this is that it allows bandwidth (especially when it is scarce) to be allocated for “special events” at certain times. For example, suppose every evening at 9PM to 10PM an online backup takes place through the VPN. The Access Policy for all other traffic can be put on a schedule that states it is active from 10PM to 9PM everyday. This will have the effect of turning off the other policies during the 9PM – 10PM period, giving more bandwidth to the Encryption Access Policy. Before a schedule can be picked from the pull down list in the Access Policy, a schedule needs to be created. This can be done by clicking on the Schedule button on the left and adding a schedule.

Cannot Ping Between Unsecure Hosts and Secure Hosts

Be sure to enable pings on the Trusted interface and make sure the Untrusted configuration is enabled.

Each router adjacent to the firewall must contain a static route specifying the firewall as the gateway for destination networks beyond the firewall.

> ***Note:*** If your Trusted network uses addresses that are not registered and routable on the Untrusted network, including private addresses as specified in RFC 1597, packets will not be routed back to the sender. Use a client with a registered address. The firewall's NAT feature may be used for TCP and UDP traffic, but NAT will not translate addresses in ICMP packets like ping.

To enable the hosts on the Trusted and Untrusted sides to send and respond to pings between each other, do the following:

WebUI

1. Policy >> Outgoing >> New Policy: Enter the following information, and click **OK**.

   Source Address: Inside Any

   Destination Address: Outside Any

   Service: PING

   Action: Permit

2. Policy >> Incoming >> New Policy: Enter the following information, and click **OK**.

   Source Address: Outside Any

   Destination Address: (the IP address of your workstation)

   Service: PING

   Action: Permit

CLI

1. set policy outgoing “inside any” “outside any” ping permit
2. set policy incoming “outside any” “inside any” ping permit

CONTACTING TECHNICAL SUPPORT

NetScreen offers technical support via phone or through the web, for users with a service level agreement. All NetScreen products must be registered, either via the web site, or by faxing in the registration information. Once you are registered, you are entitled to 30 days of free support, and are assigned a service level ID and password.

- Voice: 1-408-330-7800
- FAX: 408-330-7850
- E-mail: support@netscreen.com
- World Wide Web: http://www.netscreen.com (where you can create a trouble ticket online, query the online knowledge base, and view a list of the top 10 support items).
- Please be prepared to provide the following information:
  - Your name
  - Company name
  - Telephone number
  - Fax number
  - E-mail address
  - Model number
  - 8-digit Serial number
  - Description of the problem
  - Service ID number
  - Output of **get tech-support** command which includes a copy of the configuration, session information, etc.
  - Any error messages
  - Operating system level and application types and levels
  - Network diagram

Glossary

A

10BaseT. The most common form of Ethernet is called 10BaseT, which denotes a peak transmission speed of 10 Mbps using copper twisted-pair cable. Ethernet is a standard for connecting computers into a local area network (LAN). The maximum cable distance is 100 meters (325 feet), the maximum devices per segment is 1, and the maximum devices per network are 1024. See also *100BaseT*.

100BaseT. Another term for fast Ethernet, an upgraded standard for connecting computers into a local area network (LAN). 100BaseT Ethernet works just like regular Ethernet except that it can transfer data at a peak rate of 100 Mbps. It is also more expensive and less common than its slower 10BaseT sibling.

Access Policies. Access Policies provide the initial protection mechanism for the firewall, allowing you to determine what traffic passes across it based on IP session details. They protect the Trusted network from outsider attacks, such as the scanning of Trusted servers. Access Policies create an environment in which you set up security Policies to monitor traffic attempting to cross your firewall.

Authentication Header (AH). See *ESP/AH*.

Authentication. Authentication ensures that digital data transmissions are delivered to the intended receiver. Authentication also assures the receiver of the integrity of the message and its source (where or whom it came from). The simplest form of authentication requires a user name and password to gain access to a particular account. Authentication protocols can also be based on secret-key encryption, such as DES, or on public-key systems using digital signatures.

Bridge. A device that forwards traffic between network segments based on data link layer information. These segments would have a common network layer address.

Circuit-level Proxy. Proxy or Proxy Server is a technique used to cache information on a Web server and acts as an intermediary between a Web client and that Web server. It basically holds the most commonly and recently used content from the World Wide Web for users in order to provide quicker access and to increase server security. This is common for an ISP especially if they have a slow link to the Internet. On the Web, a proxy first attempts to find data locally, and if it's not there, fetches it from the remote server where the data resides permanently. Proxy servers are also constructs that allow direct Internet access from behind a firewall. They open a socket on the server, and allow communication via that socket to the Internet. For example, if your computer is

inside a protected network, and you want to browse the Web using Netscape, you would set up a proxy server on a firewall. The proxy server would be configured to allow requests from your computer, trying for port 80, to connect to its port 1080, and it would then redirect all requests to the proper places.

Data Encryption Standard (DES). A 40- and 56-bit encryption algorithm that was developed by the National Institute of Standards and Technology (NIST). DES is a block encryption method originally developed by IBM. It has since been certified by the U.S. government for transmission of any data that is not classified top secret. DES uses an algorithm for private-key encryption. The key consists of 64 bits of data, which are transformed and combined with the first 64 bits of the message to be sent. To apply the encryption, the message is broken up into 64-bit blocks so that each can be combined with the key using a complex 16-step process. Although DES is fairly weak, with only one iteration, repeating it using slightly different keys can provide excellent security.

Data Encryption Standard-Cipher Block Chaining (DES-CBC). Until recently, the most significant use of triple-DES (3DES) was for the encryption of single DES keys, and there was really no need to consider how one might implement various block cipher modes when the block cipher in question is actually one derived from multiple encryption. However, as DES nears the end of its useful lifetime, more thought is being given to an increasingly widespread use of triple-DES. In particular, there are two obvious ways to implement the CBC mode for triple-DES. With single-DES in CBC mode, the ciphertext is exclusive-ored with the plaintext before encryption. With triple-DES however, we might use feedback around all three DES operations from the ciphertext to the plaintext, something which is called outer-CBC. Alternatively, we might run the feedback around each individual encryption component, thereby making, in effect, triple-(DES-CBC). This is referred to as inner-CBC, since there are internal feedbacks that are never seen by the crypto-analyst. Performance-wise, there can be some advantages to use the inner-CBC option, but research has established that outer-CBC is in fact more secure. Outer-CBC is the recommended way for using triple-DES in the CBC mode.

De-Militarized Zone (DMZ). From the military term for an area between two opponents where fighting is prevented. DMZ Ethernets connect networks and computers controlled by different bodies. They may be external or internal. External DMZ Ethernets link regional networks with routers.

Encryption. Encryption is the process of changing data into a form that can be read only by the intended receiver. To decipher the message, the receiver of the encrypted data must have the proper decryption key. In traditional encryption schemes, the sender and the receiver use the same key to encrypt and decrypt data. Public-key encryption schemes use two keys: a public key, which anyone may use, and a corresponding private key, which is possessed only by the person

who created it. With this method, anyone may send a message encrypted with the owner's public key, but only the owner has the private key necessary to decrypt it. PGP (Pretty Good Privacy) and DES (Data Encryption Standard) are two of the most popular public-key encryption schemes.

ESP/AH. The IP level security headers, AH and ESP, were originally proposed by the Network Working Group focused on IP security mechanisms, IPSec. The term IPSec is used loosely here to refer to packets, keys, and routes that are associated with these headers. The IP Authentication Header (AH) is used to provide authentication. The IP Encapsulating Security Header (ESP) is used to provide confidentiality to IP datagrams.

Ethernet. A local area network technology invented at the Xerox Corporation, Palo Alto Research Center. Ethernet is a best-effort delivery system that uses CSMA/CD technology. Ethernet can be run over a variety of cable schemes, including thick coaxial, thin coaxial, twisted pair, and fiber optic cable. Ethernet is a standard for connecting computers into a local area network (LAN). The most common form of Ethernet is called 10BaseT, which denotes a peak transmission speed of 10 Mbps using copper twisted-pair cable.

Extranet. The connecting of two or more intranets. If an intranet as a company's internal Web site which allows users inside the company to communicate and exchange information, an extranet connects that virtual space with another company's intranet, thus allowing these two (or more) companies to share resources and communicate over the Internet in their own virtual space. This technology greatly enhances business to business communications.

Filtering, dynamic. IP service that can be used within VPN tunnels. Filters are one way the NetScreen-10/100 controls traffic from one network to another. When TCP/IP sends data packets to the firewall, the filtering function in the firewall looks at the header information in the packets and directs them accordingly. The filters operate on criteria such as IP source or destination address range, TCP ports, UDP, Internet Control Message Protocol (ICMP), or TCP responses. See also *Tunneling and Virtual Private Network (VPN)*.

Firewall. A device that protects and controls the connection of one network to another, for traffic both entering and leaving. Firewalls are used by companies that want to protect any network-connected server from damage (intentional or otherwise) by those who log in to it. This could be a dedicated computer equipped with security measures or it could be a software-based protection.

Hub. This hardware is used to network computers together (usually over an Ethernet connection). It serves as a common wiring point so that information can flow through one central location to any other computer on the network thus enabling centralized management. A hub is a hardware device that repeats signals at the physical Ethernet layer. A hub retains the behavior of a standard bus type network (such as Thinnet), but produces a star topology with the hub at the center of the star. This configuration enables centralized management.

Internet Control Message Protocol (ICMP). Occasionally a gateway or destination host will communicate with a source host, for example, to report an error in datagram processing. For such purposes the protocol, the Internet Control Message Protocol (ICMP), is used. ICMP uses the basic support of IP as if it were a higher level protocol, however, ICMP is actually an integral part of IP, and must be implemented by every IP module. ICMP messages are sent in several situations: for example, when a datagram cannot reach its destination, when the gateway does not have the buffering capacity to forward a datagram, and when the gateway can direct the host to send traffic on a shorter route. The Internet Protocol is not designed to be absolutely reliable. The purpose of these control messages is to provide feedback about problems in the communication environment, not to make IP reliable.

Internet. Also known as “the Net.” Originally designed by the U.S. Defense Department so that a communication signal could withstand a nuclear war and serve military institutions worldwide. The Internet was first known as the ARPAnet. A system of linked computer networks, international in scope, that facilitates data communication services such as remote login, file transfer, electronic mail, and newsgroups. The Internet is a way of connecting existing computer networks that greatly extends the reach of each participating system.

Internet Key Exchange (IKE). The method for exchanging keys for encryption and authentication over an unsecured medium, such as the Internet.

Internet Protocol (IP). An Internet standard protocol that defines a basic unit of data called a datagram. A datagram is used in a connectionless, best-effort, delivery system. The Internet protocol defines how information gets passed between systems across the Internet.

IP Address. Each node on a TCP/IP network usually has an IP address. The IP address has a network number portion and a host number portion, as shown in the following table:

Table C-1 IP Address Classes and Formats

Class	Number of Nodes	Address Format
A	> 32,768	nnn.hhh.hhh.hhh
B	256–32,768	nnn.nnn.hhh.hhh
C	<256	nnn.nnn.nnn.hhh

This format is called decimal dot format. The “n” represents a digit of a network number and “h” represents a digit of a host number; for example, 128.11.2.30. If you are sending data outside of your network, such as to the Internet, you need to obtain the network number from a central authority, currently the Network Information Center. See also *Subnet Mask*.

IP Gateway. Also called a router, a gateway is a program or a special-purpose device that transfers IP datagrams from one network to another until the final destination is reached.

IP Security (IPSec). Security standard produced by the Internet Engineering Task Force (IETF). It is a protocol suite that provides everything you need for secure communications—authentication, integrity, and confidentiality—and makes key exchange practical even in larger networks. See also *DES-CBC, ESP/AH*.

ISAKMP. The Internet Security Association and Key Management Protocol (ISAKMP) provides a framework for Internet key management and provides the specific protocol support for negotiation of security attributes. By itself, it does not establish session keys, however it can be used with various session key establishment protocols to provide a complete solution to Internet key management.

Intranet. A play on the word Internet, an intranet is a restricted-access network that works like the Web, but isn't on it. Usually owned and managed by a corporation, an intranet enables a company to share its resources with its employees without confidential information being made available to everyone with Internet access.

Key Management, Manual. The only reasonable way to protect the integrity and privacy of information is to rely upon the use of secret information in the form of private keys for signing and/or encryption. The management and handling of these pieces of secret information is generally referred to as “key management.” This includes the activities of selection, exchange, storage, certification, expiration, revocation, changing, and transmission of keys. Most of the work in managing information security systems lies in the key management.

Load balancing. Load balancing is the mapping (or re-mapping) of work to processors, with the intent of improving the efficiency of a concurrent computation.

Local Area Network (LAN). Any network technology that interconnects resources within an office environment, usually at high speeds, such as Ethernet. A local area network is a short-distance network used to link a group of computers together within a building. 10BaseT Ethernet is the most commonly used form of LAN. A hardware device called a hub serves as the common wiring point, enabling data to be sent from one machine to another over the network. LANs are typically limited to distances of less than 500 meters and provide low-cost, high-bandwidth networking capabilities within a small geographical area.

MD5. Message Digest (version) 5, an algorithm that produces a 128-bit message digest (or hash) from a message of arbitrary length. The resulting hash is used, like a “fingerprint” of the input, to verify authenticity.

Media Access Control (MAC) Address. An address that uniquely identifies the network interface card, such as an Ethernet adapter. For Ethernet, the MAC address is a 6 octet address assigned by IEEE. On a LAN or other network, the MAC address is a computer's unique hardware number. (On an Ethernet LAN, it's the same as the Ethernet address.) When you're connected to the Internet from your computer (or host as the Internet protocol thinks of it), a correspondence table relates your IP address to your computer's physical (MAC) address on the LAN. The MAC address is used by the Media Access Control sublayer of the Data-Link Control (DLC) layer of telecommunication protocols. There is a different MAC sublayer for each physical device type.

Network Address Translation (NAT). A standard for translating secure IP addresses to temporary, external, registered IP address from the address pool. This allows Trusted networks with privately assigned IP addresses to have access to the Internet. This also means that you don’t have to get a registered IP address for every machine in your network.

RJ-45. Resembling a standard phone connector, an RJ-45 connector is twice as wide (with eight wires) and is used for hooking up computers to local area networks (LANs) or phones with multiple lines.

Router. This hardware device routes data from a local area network (LAN) to a phone line's long distance line. Routers also act as traffic cops, allowing only authorized machines to transmit data into the local network so that private information can remain secure. In addition to supporting these dial-in and leased connections, routers also handle errors, keep network usage statistics, and handle security issues.

Security Association. The combination of a Security Parameters Index and a destination address. Required for both Authentication Header and Encapsulating Security Payload protocols. See also Security Parameters Index.

Security Parameters Index. (SPI) is a hexadecimal value which uniquely identifies each tunnel. It also tells the NetScreen device which key to use to decrypt packets.

Server Farm. A server farm is a network where clients install their own computers to run Web servers, email, or any other TCP/IP based services they require, making use of leased permanent Internet connections with 24-hour worldwide access. Instead of expensive dedicated-line connections to various offices, servers can be placed on server farm networks to have them connected to the Internet at high-speed for a fraction of the cost of a leased line.

SHA-1. Secure Hash Algorithm-1, an algorithm that produces a 160-bit hash from a message of arbitrary length. (It is generally regarded as more secure than MD5 because of the larger hashes it produces.)

Subnet Mask. In larger networks, the subnet mask lets you define subnetworks. For example, if you have a class B network, a subnet mask of 255.255.255.0 specifies that the first two portions of the decimal dot format are the network number, while the third portion is a subnet number. The fourth portion is the host number. If you do not want to have a subnet on a class B network, you would use a subnet mask of 255.255.0.0. A network can be subnetted into one or more physical networks which form a subset of the main network. The Subnet Mask is the part of the IP address which is used to represent a subnetwork within a network. Using Subnet Masks allows you to use network address space which is normally unavailable and ensures that network traffic does not get sent to the whole network unless intended. Subnet Masks are a complex feature, so great care should be taken when using them. See also *IP address*.

Three-Way Handshake. A TCP connection is established with a triple exchange of packets known as a three-way handshake. The procedure transpires as follows:

1. The initiator sends a SYN (synchronize/start) packet.
2. The recipient replies with a SYN/ACK (synchronize/acknowledge) packet.
3. The initiator responds with an ACK (acknowledge) packet.

4. At this point, the two endpoints of the connection have been established and data transmission can commence.

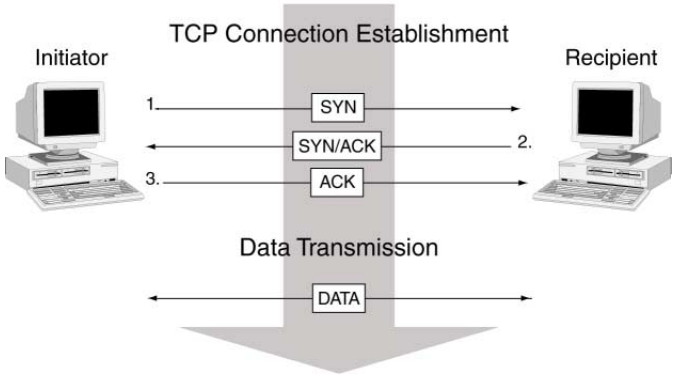

Transmission Control Protocol/Internet Protocol (TCP/IP). A set of communications protocols that support peer-to-peer connectivity functions for both local and wide area networks. A communications protocol which allows computers with different operating systems to communicate with each other. Controls how data is transferred between computers on the Internet.

Trunk Port. A trunk port allows a switch to bundle traffic from several VLANs through a single physical port, sorting the various packets by the VLAN identifier (VID) in their packet headers.

Tunneling. A method of data encapsulation.With VPN tunneling, a mobile professional dials into a local Internet Service Provider's Point of Presence (POP) instead of dialing directly into their corporate network. This means that no matter where mobile professionals are located, they can dial a local Internet Service Provider that supports VPN tunneling technology and gain access to their corporate network, incurring only the cost of a local telephone call. When remote users dial into their corporate network using an Internet Service Provider that supports VPN tunneling, the remote user as well as the organization knows that it is a secure connection. All remote dial-in users are authenticated by an authenticating server at the Internet Service Provider's site and then again by another authenticating server on the corporate network. This means that only authorized remote users can access their corporate network, and can access only the hosts that they are authorized to use.

User Datagram Protocol (UDP). A protocol in the TCP/IP protocol suite, the User Datagram Protocol or UDP allows an application program to send datagrams to other application programs on a remote machine. Basically UDP is a protocol that provides an unreliable and connectionless datagram service where delivery and duplicate detection are not guaranteed. It does not use acknowledgments, or control the order of arrival.

Universal Resource Locator (URL). A standard way developed to specify the location of a resource available electronically.Also referred to as a location or address, URLs specify the location of files on servers. A general URL has the syntax protocol://address. For example, http://www.srl.rmit.edu.au/pd/index.html specifies that the protocol is http and the address is www.srl.rmit.edu.au/pd/index.html.

Unshielded Twisted Pair (UTP). Also known as 10BaseT. This is the standard cabling used for telephone lines. It is also used for Ethernet connections.

Virtual Local Area Network (VLAN). A logical rather than physical grouping of devices that constitute a single broadcast domain. VLAN members are not identified by their location on a physical subnetwork but through the use of tags in the frame headers of their transmitted data. VLANs are specified in the IEEE 802.1Q standard.

Virtual Private Network (VPN). A VPN is an easy, cost-effective and secure way for corporations to provide telecommuters and mobile professionals local dial-up access to their corporate network or to another Internet Service Provider (ISP). Secure private connections over the Internet are more cost-effective than dedicated private lines. VPNs are possible because of technologies and standards such as tunneling, screening, encryption, and IPSec.

Virtual System. A feature unique to the NetScreen-1000, a Virtual System is a subdivision of the main system that appears to the user to be a stand-alone entity. Virtual Systems reside separately from each other in the same NetScreen-1000 device. Each one can be managed by its own Virtual System Administrator.

Windows Internet Naming Service (WINS). WINS is a service for mapping IP addresses to NetBIOS computer names on Windows NT server-based networks. A WINS server maps a NetBIOS name used in a Windows network environment to an IP address used on an IP-based network.

Index

Numerics

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

Q

R

S

T

W