AVAYA

Avaya VPNmanager Configuration Guide

Release 3.4

Notice

Every effort was made to ensure that the information in this document was complete and accurate at the time of release. However, information is subject to change.

Warranty

Avaya Inc. provides a limited warranty on this product. Refer to your sales agreement to establish the terms of the limited warranty. In addition, Avaya's standard warranty language as well as information regarding support for this product, while under warranty, is available through the following website: http://www.avaya.com/support

Preventing Toll Fraud

"Toll fraud" is the unauthorized use of your telecommunications system by an unauthorized party (for example, a person who is not a corporate employee, agent, subcontractor, or is not working on your company's behalf). Be aware that there may be a risk of toll fraud associated with your system and that, if toll fraud occurs, it can result in substantial additional charges for your telecommunications services.

Avaya Fraud Intervention

If you suspect that you are being victimized by toll fraud and you need technical assistance or support, in the United States and Canada, call the Technical Service Center's Toll Fraud Intervention Hotline at 1-800-643-2353.

How to Get Help

For additional support telephone numbers, go to the Avaya Web site:
http:/www.avaya.com/support/

If you are:
- Within the United States, click *Escalation Lists*, which includes escalation phone numbers within the USA.
- Outside the United States, click *Escalation Lists* then click *Global Escalation List*, which includes phone numbers for the regional Centers of Excellence.

Providing Telecommunications Security

Telecommunications security (of voice, data, and/or video communications) is the prevention of any type of intrusion to (that is, either unauthorized or malicious access to or use of) your company's telecommunications equipment by some party.

Your company's "telecommunications equipment" includes both this Avaya product and any other voice/data/video equipment that could be accessed via this Avaya product (that is, "networked equipment").

An "outside party" is anyone who is not a corporate employee, agent, subcontractor, or is not working on your company's behalf. Whereas, a "malicious party" is anyone (including someone who may be otherwise authorized) who accesses your telecommunications equipment with either malicious or mischievous intent.

Such intrusions may be either to/through synchronous (time-multiplexed and/or circuit-based) or asynchronous (character-, message-, or packet-based) equipment or interfaces for reasons of:
- Utilization (of capabilities special to the accessed equipment)
- Theft (such as, of intellectual property, financial assets, or toll-facility access)
- Eavesdropping (privacy invasions to humans)
- Mischief (troubling, but apparently innocuous, tampering)
- Harm (such as harmful tampering, data loss or alteration, regardless of motive or intent)

Be aware that there may be a risk of unauthorized intrusions associated with your system and/or its networked equipment. Also realize that, if such an intrusion should occur, it could result in a variety of losses to your company (including but not limited to, human/data privacy, intellectual property, material assets, financial resources, labor costs, and/or legal costs).

Responsibility for Your Company's Telecommunications Security

The final responsibility for securing both this system and its networked equipment rests with you - Avaya's customer system administrator, your telecommunications peers, and your managers. Base the fulfillment of your responsibility on acquired knowledge and resources from a variety of sources including but not limited to:
- Installation documents
- System administration documents
- Security documents
- Hardware-/software-based security tools
- Shared information between you and your peers
- Telecommunications security experts

To prevent intrusions to your telecommunications equipment, you and your peers should carefully program and configure:
- Your Avaya-provided telecommunications systems and their interfaces
- Your Avaya-provided software applications, as well as their underlying hardware/software platforms and interfaces
- Any other equipment networked to your Avaya products.

All Avaya media servers and media gateways are compliant with FCC Part 68, but many have been registered with the FCC before the SDoC process was available. A list of all Avaya registered products may be found at:
http://www.part68.org/
by conducting a search using "Avaya" as manufacturer.

European Union Declarations of Conformity

CE

Avaya Inc. declares that the equipment specified in this document bearing the "CE" (*Conformité Europeénne*) mark conforms to the European Union Radio and Telecommunications Terminal Equipment Directive (1999/5/EC), including the Electromagnetic Compatibility Directive (89/336/EEC) and Low Voltage Directive (73/23/EEC). This equipment has been certified to meet CTR3 Basic Rate Interface (BRI) and CTR4 Primary Rate Interface (PRI) and subsets thereof in CTR12 and CTR13, as applicable.

Copies of these Declarations of Conformity (DoCs) can be obtained by contacting your local sales representative and are available on the following Web site:
http://www.avaya.com/support

Japan

This is a Class A product based on the standard of the Voluntary Control Council for Interference by Information Technology Equipment (VCCI). If this equipment is used in a domestic environment, radio disturbance may occur, in which case, the user may be required to take corrective actions.

この装置は、情報処理装置等電波障害自主規制協議会（ＶＣＣＩ）の基準
に基づくクラスＡ情報技術装置です。この装置を家庭環境で使用すると電波
妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ず
るよう要求されることがあります。

China

BMSI (Chinese Warning Label)

警告使用者：
這是甲類的資訊產品，在居住的環境中使用時，
可能會造成射頻干擾，在這種情況下，使用者會
被要求採取某些適當的對策。

Hardware, including technical data, is subject to U.S. export control 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 licenses to export, re-export, or import hardware.

Acknowledgments:

This product includes software developed by the Apache Software Foundation (http://www.apache.org).

Documentation:

For the most current versions of documentation, go to http://www.avaya.com/support/

Table of Contents

Chapter 2   *Using VPNmanager*

Chapter 3   Setting up the network

Chapter 6   *Configuring user groups*

Chapter 7   *Configuring VPN objects*

Chapter 8   Establishing security

Chapter 9   *Using advanced features*

Appendix B   *Firewall rules template*

Glossary

Index

Preface

This Avaya™ VPNmanager Administrator's Guide is written for individuals who have an understanding of how computer networks are installed, configured, and managed. It provides detailed information about using the Avaya VPNmanager solution to build small, medium, or large scale Virtual Private Networks (VPNs).

This guide also provides information about the VPNmanager software. VPNmanager is a Java-based software application that brings convenience, ease of use, extended functionality, and platform independence to the management of VPNs.

What Products are Covered

Avaya's solution is a line of three products that are used for managing Virtual Private Networks. Each one, listed below, has been designed to meet the needs and requirements of either a small, medium, or large network.

- VPNmanager Service Provider
- VPNmanager Enterprise
- VPNmanager Small Office

VPNmanager Overview

The VPNmanager application lets network managers define, configure, and manage Virtual Private Networks (VPNs) from any location equipped with a computer running Windows 98, Window NT, Window 2000, or Solaris.

Network managers can configure and check the status of Avaya security gateways and VPN Service Units (VSU), add or remove remote sites and dial-in users to a VPN, configure user authentication servers using LDAP directory servers or RADIUS servers, and monitor the state of all security gateways, as well as the performance of private data transmissions using Java-interface technology.

Network-wide Visibility and Control

The logical VPNmanager representation of virtual private networks simplifies their installation and control. From a single workstation, network managers can assign users anywhere on the network to one or more logical Groups and integrate local and remote Groups into VPNs. The VPNmanager software provides global-level, VPN-level, group-level, client-level and equipment-level monitoring and control capabilities, and automates the task of managing configurations across multiple security gateways and Avaya VPNremote® Clients. Extensive alarm-reporting and statistics-gathering capabilities allow network managers to respond in real time to hardware, network, and security problems, and to plan the efficient growth and evolution of their networks.

Intranet and Extranet Support

The VPNmanager software makes it easy to extend intranet services to remote sites and users securely. In addition, the VPNmanager's sophisticated import and export capabilities enable network managers from different organizations to securely link with one another into private wide-area "extranets." Companies can quickly link and unlink to their suppliers, customers, consultants, and other business associates with flexibility and speed unmatched by traditional communications services.

Secure VPN Configuration

Several mechanisms are employed to insure security when managing VPNs. Industry-standard Secure Socket Layer (SSL) technology is used to keep configuration traffic between the VPNmanager and VSUs private. In addition, X.509 certificates are used by both VSUs and the VPNmanager console providing an authentication capability, thus allowing only authorized administrators to configure VSUs. Once authenticated, administrators can configure, modify, restart, or upgrade any security gateway in the corporate network. Finally, sensitive

cryptographic keying information stored in the VPNmanager database is encrypted using a password key to prevent compromising secure network traffic.

No Special Consoles Required

The VPNmanager software runs on host environments that support the Java Virtual Machine (see the VPNmanager README file for a current list of supported platforms). Expensive management consoles and proprietary management interfaces are not needed. Regardless of the host platform, the VPNmanager software presents the same appearance and user controls.

Complementary to SNMP Management Tools

The VPNmanager software is designed specifically for securely defining, configuring, monitoring, and upgrading VPNs. The VPNmanager software is required to configure and modify VPNs. Secure traffic running between VSUs or between VSUs and VPNremote Clients does not require an active VPNmanager. After configuring the required VPNs, the VPNmanager can be shutdown if desired, or used to monitor security gateway activity. In addition, standard MIBs available with the VSUs enable monitoring from standard SNMP management stations.

Using VPNmanager Help

The VPNmanager comes with a context-sensitive-Help system. Use the Help system for getting information about a specific command in the VPNmanager graphical user interface (GUI).

Related Documentation

Be sure to read the VPNos Configuration Guide.It contains important information on the proper procedure for setting up your VSUs, which is a prerequisite to setting up a Virtual Private Network.

VPNremote Client software installation and usage information is found in the *VPNremote Client Administrator's Guide*. This software allows the network administrator to pre-configure the VPNremote client software for distribution to end users via the web, or on portable storage media such as a CD or floppy disk.

You can download these documents from www.avaya.com. Click on Product Documentation, select **VPN and Security**.

How This Book Is Organized

With this release of VPNmanager, the administrator's guide was redesigned to present information in the order that you use VPNmanager to configure a secure network.

Note: Depending on the VPNmanager version, some features described in this guide may not apply.

Chapter 1, Overview of implementation, provides an overview of how to use VPNmanger for centralized administration of your VPN and security gateway. It includes a checklist for implmeneting the network.

Chapter 2, Using VPNmanager, explains how to log in to VPNmanager. It also explains how to use the VPNmanager interface, inlcuding the VPNmanager main console and the configuration console. The VPNmanager Preferences are described here.

Chapter 3, Setting up the network explains how to create a domain and create and configure a security gateway. This chapter explains how to configure the Device object, including multiple zones, NAT services, DNS, and Static Route.

Chapter 4, Configuring IP Groups, describes how to configure IP Group Objects for Data Terminal Equipment (DTE) such as computers, printers, and network servers as members of your VPN.

Chapter 5, Configuring remote access users, describes how to setup and maintain individual remote access users in the VPN. This chapter includes Dyna-Policy configuration and information about the Policies tab including Client IP configuration RADIUS/ACE services, and client attributes

Chapter 6, Configuring user groups, describes how to setup and maintain logical groups that the individual VPN remote users reside.

Chapter 7, Configuring VPN objects, explains VPN Objects as the method for linking VSUs, remote terminals, and LAN terminals in a fully configured VPN.

Chapter 8, Establishing security, describes the levels of Firewall policy management and Denial of Service available, how to configure the

security gateway for Voice over IP and how to create and map Quality of Service (Qos) rules.

Chapter 9, Using advanced features, describes about using certificates, configuring the Directory Server, resilient tunnels, and high availability groups.

Chapter 10, Monitoring your network, describes the monitoring and reporting features of the VPNmanager software. This includes SNMP, Syslog Services, Reports, and Alarms. These features allow virtual real-time monitoring of the VPN performance and specific security gateways.

Chapter 11, Device management, describes how to optimize the VPNmanager, check connectivity, reset the device time, reboot, resetup a security gateway and how to import a VPN.

Chapter 12, Upgrading firmware and licenses, describes how to use the automatic upgrade feature to upgrade the firmware for a security gateway or for a group of secruity gateways, and how to add new licenses to your security gateway.

Appendix A, Using SSL with Directory Server, describes the benefit of using secure socket layer (SSL) with the Directory Server.

Appendix B, Firewall rules template, describes the predefined firewall templates that are included in the VPNmanager.

Contacting Technical Support

Technical Support is available to support contract holders of Avaya VPN products.

Domestic support

- Toll free telephone support: (866) 462-8292 (24x7)

- Email: vpnsupport@avaya.com

- Web: http://www.support.avaya.com

International Support

- For regional support telephone numbers, go to http:www.avayanetwork.com/site/GSO/default.htm

Chapter 1   Overview of implementation

Planning how your virtual private network should be configured is critical to the successful deployment of a secure virtual private network.This chapter provides an overview of the major features that you will configure.

Note: This chapter does not explain how to set up a VPN or how to determine what type of security policies are required. You should understand about networking, establishing firewall policies, and VPNs before implementing a VPN using VPNmanager.

Components of the Avaya security solution

The Avaya security solution consists of the following:

- Avaya VPNmanager™

- Avaya™ SG security gateways and VPN Service Units (VSU)

   Note: Beginning with VPNmanager 3.4, this configuration guide uses "security gateway" to refer to both the security gateway and the VSU. The VPNmanager application uses the word "Device" to refer to both of these components.

- Avaya VPNremote™ Client

Security gateways

The security gateways are designed to provide firewall coverage and VPN gateway functionality for enterprises migrating towards converged network environments. The security gateway performs cryptography, authentication, and filtering tasks at the boundary of the VPN.

After the security gateway is installed and configured, the security gateway is transparent to users who are logged into the VPN.

VPNremote Client software

VPNremote Client software is a communications application that runs on remote computers that use dialup, DSL and cable connection supplied by Internet Service Providers (ISP), to connect to the corporate VPN. When communicating with a VPN, the software seamlessly performs authentication and cryptography tasks. To install and use the software, an account with an ISP must first be created.

The software is installed on the remote user's computer and then Client Configuration Download (CCD) can be used to configure the remote user's Dyna-Policy™ for authentication to a specific VPN.

When remote users log in, they connect to the ISP and type in their user authentication information, if asked. Upon authentication, any traffic that uses the VPN is safely encrypted as it is transported through the public networks.

VPNmanager software

VPNmanager software lets network managers define, configure, manage VPN and firewall policies, upgrade firmware, and manage remote user access policies from a central location.

The VPNmanager software combines three components, the VPNmanger Console, the policy server and the directory server.

- The VPNmanager console is a client that is used for configuring, managing, and monitoring one or more VPNs. The console is a Java application that can be run anywhere and is used as a front-end to the policy server and the directory server.

- The directory server is a database used to store configuration data. The configuration includes VPNs, security gateways, and the security gateway's security policy.

- The policy server distributes configuration and security policies. The VPNmanager console is a client that communicates with the policy server to retrieve security policies. The policy server then communicates with the directory server.

The VPNmanager Console and the directory server can reside on separate, dedicated servers within the network to provide better performance for updating and configuring large numbers of security gateways. You can use either an existing Sun One Server or Microsoft Active Directory Server to store the policies that are created.

VPNmanager software consists of different versions to meet the needs of various networks.

- **VPNmanager Small Office.** Use the small office version for managing up to five security gateways and unlimited VPNremote Clients.

- **VPNmanager Enterprise.** Use the VPNmanger Enterprise version for managing an unlimited number of devices and VPNremote Clients. The Sun One™ server directory server installation is included in this version.

- **VPNmanager Service Provider.** Use this version to manage an unlimited number of devices and VPNremote Clients. The Service Provider also supports multiple VPN domains, which meets the needs of ISPs. The Sun One™ server directory server installation is included in this version.

- **VPNmanager Enterprise Client.** Use the Enterprise Client version for managing an unlimited number of security gateways and VPNremote Clients. A directory server installation is not included in this version.

- **VPNmanager Service Provider Client.** Use the Service Provider Client version to manage an unlimited number of security gateways and VPNremote Clients. The Service Provider also supports multiple VPN domains. A directory server installation is not included in this version.

Overview of the VPN management hierarchy

With the VPNmanager software, you can configure and manage VPNs and firewalls from a central location. By focusing on security policy instead of individual device management, administration of large-scale networks is simplified. Central management allows you to make configuration updates automatically to all affected security gateways. This distributed approach also applies to firewall management.

The VPNmanager software is built on a policy-based architecture that allows the administrator to start at a high-level with a VPN domain, then move down the hierarchy to create user groups, IP groups for protected resources, and security groups that define membership and policies of the VPN.

Figure 1. Domain hierarchy
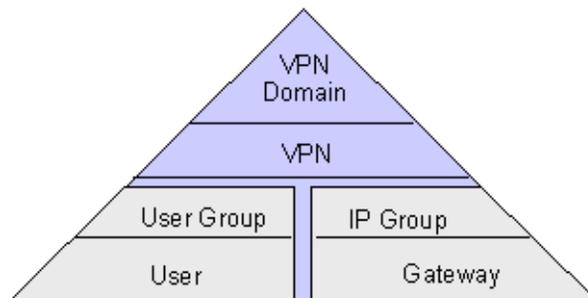
At the peak of the hierarchy is the VPN domain. A domain is assigned a name to identify it from other domains. Usually one domain is configured for an entire organization. A domain is built of one or more VPNs.

Each VPN is built of users, user groups and IP groups. VPNs are assigned names. These names can associate the VPN to a regional location or purpose.

Users are the individual remote access users who log in to the VPN through a security gateway. The VPNremote Client software is used to connect to the VPN services.

A User Group contains or organizes user accounts. These accounts are assigned to remote VPN members who dial in to the network and run VPNremote Client software to access the VPN.

An IP Group contains the IP addresses that belong to a specific LAN. Any device connected to the LAN can use these addresses. A VPN can have many IP Groups so addresses can be consolidated to meet the needs of an organization.

The security gateway is configured to provide VPN gateway functionally and firewall coverage.

VPNmanager security management includes creating domain-level firewall rules and device-level firewall rules. VPNmanager provides multiple firewall templates that can be used as a general rule set or as a starting point for creating a customized firewall template. You can apply these templates at the domain level for all security gateways, for a specific security gateway (device-level), or for a defined device group.

Preparing to configure your network

Before you use VPNmanager to build your VPN and establish your VPN security policies, you need to know how the VPN should be implemented. This section gives a overview of what information you should know before you begin.

The following are functions or tasks that need to be addressed:

- How the security gateway will be configured for your network

- Which remote users will be configured on a security gateway

- What IP addresses to configure and group

- What type of security policies you want to implement

- What VPN services to use

- What advanced features, such as VoIP, Failover, or SNMP will be implemented

Security gateway

The security gateway is preconfigured with default settings for the media interface zones and Network Address Translation (NAT) . You may need to change default configuration for your specific network environment.

Up to six media interfaces can be configured with different *zone* interfaces. The number of zones that can be configured depends on the security gateway model (Table 1). Ethernet0 and Ethernet1 are present in all models and are assigned to the public and the private zones. The media interfaces that remain are unused and can be configured as required.

- **Public zone**. Public zone provides connection to the Internet, usually by way of a wide area network (WAN).

- **Private zone**. Private zone is used to provide connection to your private local area network (LAN) or to your corporate LAN.

- **Public-backup zone**. Public-backup zone is the backup interface to the primary public interface for use when Failover is configured.

- **Semiprivate zone**. Semiprivate zone is used for media such as wireless LAN, where the network is considered part of the protected network, but the media may be vulnerable to attack. The semi-private zone provides the additional security measure of IPSec encryption to prevent compromise to the network, for example, VPN over wireless protection.

- **DMZ zone**. DMZ (Demilitarized zone) is used for an area in the company network that needs to be accessible from the public networks, for example, email, FTP, and Web servers, but the area is not considered part of the internal private network. Servers in the DMZ typically have publicly routable IP addresses or should use advanced NAT within the security gateway.

- **Management zone**. Management zone is used to simplify network deployments, to eliminate enterprise network dependencies on switches or routers. The management network interface is usually used as an access point for a dedicated VPNmanager management station or as a dedicated interface for dumping log messages to a syslog server.

Table 1. Network zones

Media type	SG5 and SG5X	SG200	SG203	SG208
Ethernet0	Public	Public	Private	Private
Ethernet1	Private	Private	Public	Public
Ethernet2	NA		• Unused • Public backup • Private • Semiprivate • DMZ • Management	• Unused • Public backup • Private • Semiprivate • DMZ • Management
Ethernet3 to Ethernet5	NA	NA	• Unused • Public backup • Private • Semiprivate • DMZ • Management	• Unused • Public backup • Private • Semiprivate • DMZ • Management

Static Routes

Static routes are specified when more than one router exists on a network to which the security gateway must forward either VPN traffic or non-VPN traffic. You can build a static route table with up to 32 network address/mask pairs.

IP groups

Data Terminal Equipment (DTE); such as computers, printers, and network servers, are devices that can be members of a VPN. To make these devices members, you create IP Groups. An IP Group is composed of a set of hosts (workstations and servers) that are located behind a common security gateway. The hosts are defined by their IP address and mask. VPNs are made up of IP groups at multiple locations linked across

a public IP network (Internet). Assigning workstations and servers to different IP groups offers a powerful way to limit VPN traffic to specifically designated users.

Remote users and user groups

VPNremote Client users who log in to the VPN through the security gateway must have their user authentication configured on that security gateway.

If RADIUS is not used, you must configure the user name and the password for each remote user. With RADIUS, you can configure a remote user as a default user. When a remote user is configured as a default user, the user password is not required to log in. The user is authenticated by a third-party authentication server, such as RADIUS.

You can also change the default Internet Key Exchange (IKE) identity, the split tunneling option and the security option.

You can configure User Groups to setup and maintain logical groups of users.

VPN

A VPN object is the method used to link security gateways, remote terminals, and LAN terminals in a fully configured virtual private network. Creating a VPN involves naming each VPN, adding users and user groups, and adjusting the IKE and IPSec security protocols for VPN traffic.

Security policies

VPNmanager security policy management provides the following security features that can be configured:

- Firewall rules

- Denial of Service (DoS) categories

- Quality of Service (QoS) rules

- Bandwidth management

In addition, encrypton security options include Internet Key Exchange (IKE) with IPSecuirty protocol (IPSec). It applies globally to the VPN.

Firewall policies

VPNmanager firewall policy management includes domain firewall rules, device firewall rules, and firewall templates. The VPNmanager software provides multiple firewall templates that can be used as a general rule set or as a starting point for creating a customized firewall template. You can apply these templates at the domain level for all security gateways, for a specific gateway, or for a defined group. The integrated SMLI (Stateful Multi-Layer Inspection) Firewall supports firewall rules criteria based on the following:

- Source/Destination IP address or range

- TCP/UDP/ICMP protocol

- Port or port ranges

- IP protocol

- Interface

- Direction

A set of common network services is provided, and custom network services or objects can be easily defined for use in both firewall and QoS policies. Firewall rules can be individually enabled to track state information on TCP/UDP/ICMP packet flows and can be user-configured with advanced state timers. Login can also be enabled for each rule.

Note: Domain level rules and firewall templates are available for VPNos release 4.2 and later.

Denial of Service

The following Denial of Service (DOS) categories are enabled to protect the security gateway from attack by hackers.

Ping of Death. The ping of death sends packets with invalid lengths. When the receiving system attempts to rebuild the packets, the system crashes because the packet length exhausts the available memory.

IP Spoofing. This attack sends an IP packet with an invalid IP address. If the system accepts this IP address, the attacker appears to reside on the private side of the security gateway. The attacker is actually on the public side, and bypasses the firewall rules of the private side.

Smurf Attack. This attack floods the system with broadcast IP packet pings. If the flood is large enough and long enough, the attacked host is unable to receive or distinguish real traffic.

Tear Drop. This attack sends IP fragments to the system that the receiving system cannot reassemble and the system can crash.

Flood Attack. This attack floods the system with TCP connection requests, which exhausts the memory and the processing resources of the firewall. Flood attacks also attack the UDP ports. This attack attempts to flood the network by exhausting the available network bandwidth.

WinNuke Attack. This attack attempts to completely disable networking on computers that are running Windows 95 or Windows NT. This attack can be swift and crippling because it uses common Microsoft NetBIOS services.

Buffer Overflow. This attack overflows the internal buffers of the application by sending more traffic than the buffers can process.

QoS

Quality of Service (QoS) allows you to classify and prioritize traffic based on DHCP values and TCP/IP services and networks. The bandwidth available to a class of traffic can be allotted to a specific percentage of the total upstream bandwidth. Configuring QoS allows VoIP traffic to receive a higher priority. If QoS is disabled, all traffic receives the same priority.

VoIP

The security gateway can be configured to protect and enable the communication of VoIP telephones either within a VPN or firewall. The security gateway can be configured to secure Avaya Multivantage™ and IP Office™ VoIP solutions as follows:

- Secure site-to-site voice trunks such as between headquarters and branch offices or between main offices and home offices using VPNs.

- Secure VoIP servers or endpoints (IP telephones) by providing perimeter security using the VoIP aware firewall filtering that is able to dynamically open and close all ports required to pass VoIP communication between servers and endpoints

- Allow voice-secure communication with Avaya's IP Softphone and IP Office Phone Manager Pro using VPNremote Client

- Enable NAT traversal of H.323 VoIP traffic

- Optimize bandwidth for VoIP traffic using the security gateway's Quality of Service (QoS) policies

In order to successfully use VoIP it is important to thoroughly plan the implementation of the feature. Avaya suggests that you read the *Avaya IP Telephony Implementation Guide* before implementing VoIP.

Additional features

The following is a list of some of features that can be configured depending on your VPN networking requirements.

NAT

Network Address Translation (NAT) is an Internet standard that allows private (nonroutable) networks to connect to public (routable) networks. To connect private networks and public networks, address mapping is performed on a security gateway that is located between the private network and the public network.

You can set up three types of NAT mapping on the security gateway:

- **Static NAT**. With static NAT, addresses from one network are permanently mapped to addresses on another network.

- **Port NAT**. With port NAT, addresses from internal, nonroutable networks are translated to one routable address in Port NAT.

- **Port Redirection**. With port redirection, addresses from a specific IP address and a specific port are redirected to another IP address and port.

By default, NAT is enabled and the *Share public address to reach the Internet* feature is selected. NAT affects only clear traffic.

SNMP

The VPNmanager uses the SNMP protocol to monitor the security gateway. The security gateway includes a SNMP agent that supports MIB-II and a proprietary MIB. This agent is read-only and cannot be used to configure the security gateway. The agent can send traps to a list of trap agents that you configure. SNMPv1, SNMPv2c, or VNMPv3 can be selected.

Syslog

The security gateway has a syslog messaging facility for logging system error messages. The message can be automatically sent to a destination running a Syslog server.

Client IP address pooling

Access control devices (ACD), such as firewalls, guard networks from unauthorized users. Analyzing source addresses is one method ACDs use to decide which packets can enter a network. The addresses that ISPs dynamically assign to VPNremote Client users is naturally blocked because it is impossible to know ahead of time which address is assigned. You need to configure the VPNremote Client IP address pools feature with the source IP addresses that can be recognized by an ACD so that user access is not blocked.

SSL for Directory Server

As an added benefit, all communications with the directory server can be secured by SSL (Secure Sockets Layer).

You can configure your VPN to run SSL at any time. However it is recommended that you configure SSL before you put the VPN into service, so that the VPN services do not have to be stopped.

Sequence to configure your VPN

The suggested order to set up your VPN is as follows. Refer to the chapters in this *VPNmanager Administrator's Guide* for details about how to create and configure these features.

1. Create a VPN domain

2. Create the VPN

3. Create a security gateway

4. Configure needed static routes on the gateway

5. Create IP groups

6. Associate IP groups with the security gateway

7. Associate IP groups with the VPN

8. Create new users

9. Associate users with VPNs

10. Create a VPNremote Client address pool on the gateway

11. Configure firewall rules

12. Associate firewall rules with the correct gateway and security zone

13. Configure other features such as QoS, VoIP gateway, DHCP, NAT, routing, etc.

Chapter 2   Using VPNmanager

With Avaya VPNmanager you can define, configure, and manage VPNs

and firewall policies, upgrade firmware, and manage remote user access policies. The VPNmanager graphical interface is modularized by functions and tasks to make configuring a VPN fast and easy.

This chapter describes how to:

- Log in

- Navigate the VPNmanager Console interface

- Configure Preferences for the VPNmanager Console

- How to communicate with the security gateway

About VPNmanager administrators

When the VPNmanager software was installed, during the policy server login configuration, you configured the centralized management VPNmanager login ID and password.

The VPNmanager administrator is called the *super user*. The super user is the administrator that manages and administers the VPNmanager console. The super user can add additional VPNmanager administrators, delete administrator, and change their admin directory passwords.

A VPNmanager administrator can also be set up as a SNMPv3 administrator.

To add an administrator

The Admin object is used to change the super user password and to create administrators.

1. Select **Admin** from the New Objects list. The **New Admin** dialog opens

2. Enter the administrator's name and the admin directory password.

3. Click **Apply** and then click **Close.**

To configure an administrator to be an SNMPv3 admin

1. From the **Configuration Console>Admin Contents** column, select the admin to be configured as an SNMPv3 admin. Select the SNMP tab to bring it to the front.

2. Check **Enable**.

3. For the **Security Level**, select either

   - Authentication and Privacy

   - Authentication and No Privacy

   Based on the selection, the privacy settings are enable or disabled.

4. In the **Authentication Protocol** field, select either the default, HMAC_SHA1or HMAC_MD5 and enter a password.

5. For the privacy settings, the only available value is DES_CBC. Enter the privacy password.

6. When finished, click **Save**. When you configure SNMPv3 for a device, the admin name is listed.

Log into the VPNmanager console

You log in to the VPNmanager from your computer's Start menu, *Programs>Avaya> VPNmanager>Console*. You use the super user name and password that were configured when the VPNmanager software was installed.

Figure 2. VPNmanager login screen
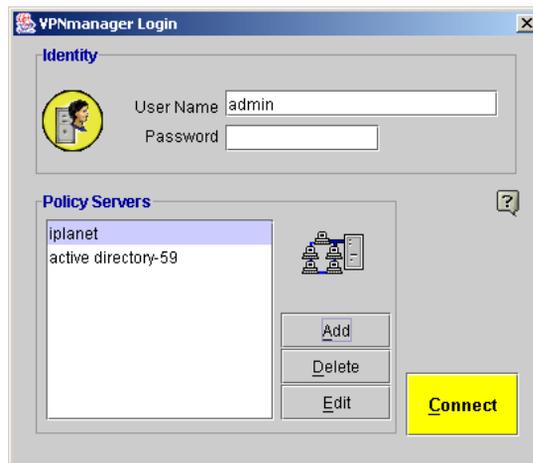
The first time you log in to the VPNmanager Console, you log in as the super user and add the policy server address or the name associated with the address. See "Add a policy server" on page 39.

Administrators that the super user creates can log in.

To log in:

1. In the User Name field, type the administrator name, if it is not displayed.

2. Type the password that was configured when the VPNmanager software was installed.

3. The IP address or name of the policy server is listed in the Policy Servers list. Select the Policy Server, if it is not highlighted and click **Connect** to log into the server.

Add a policy server

The policy server is installed during the installation of the VPNmanager Console. The policy server distributes configuration and security policies. The VPNmanager console is a client that communicates with the policy server to retrieve security policies. The policy server then communicates with the directory server.

You add the policy server address the first time you login into the VPNmanager Console.

1. From the VPNmanager Login dialog, click **Add**

2. Enter the name that identifies the Policy Server, if available. This is the "user friendly name"

3. Enter the IP address of the Policy Server.

4. Enter the port. The default is 443.

5. Click **OK**. The name or address is displayed on the login screen

You can edit or delete the policy server information.

Open Domain

When you connect to the directory server, an Open Domain screen appears. A list of all domains is displayed, with the last-selected domain highlighted.

Note: The Open Domain screen does not appear if you add a context and then click Connect on the first logon dialog.

At this point, the main console display screen appears and the selected VPN appears in the View VPN window.

Navigating the main window

The VPNmanager Console consists of the console main window, the Configuration Console window and dialogs to configure and monitor domains, VPNs, and the security gateway and network configurations related to them.

When you log in to VPNmanager for the first time, the main window is blank.The title bar shows *No Domain Open.* When you open a domain, the title bar shows the name of the domain that is opened.

The main window includes a menu bar, a toolbar, the view VPN pane, and the alarms monitoring pane.

Figure 3. VPNmanager console main window
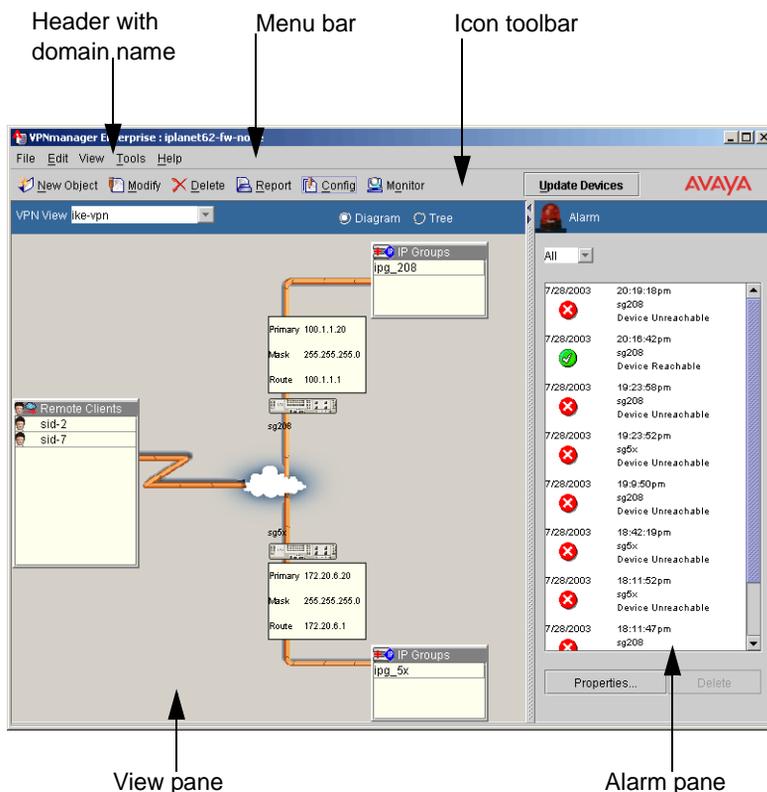
Menu bar

The menu bar on the main VPNmanager screen includes the following commands File, Edit, View, Tools, and Help.

File menu

The File menu includes the following commands:

- **Domain**. You can create a new domain, open, close, or delete an existing domain, and select from a list of recent domains that were accessed.

When you select to create **New**, a dialog to create a new domain name is displayed. This name is the unique name assigned to an overall virtual private network. A VPN domain is a collection of VPN devices that compose a VPN network. See "New VPN Domain" on page 61.

When you select **Delete** a list of all available domains is displayed. You can delete just the users within the domain, just the user groups within the domain, or all objects with the domain.

Note: When you delete VPNs that include groups associated with RADIUS-enabled security gateways, the VPNremote Client configuration records should be removed from the RADIUS database. See "RADIUS/ACE Services" on page 141.

- **New Object**. When New Object is selected, a list of objects that can be created are displayed. When you select one of these commands, either a dialog or a wizard is opened to configure the information. Table 2 describes the new objects that can be configured.

- **Logoff**. Logoff closes the current directory server without exiting VPNmanager. The Login screen appears immediately after you log off.

- **Exit**. Exit closes the VPNmanager console.

Figure 4. File Menu>New Object list
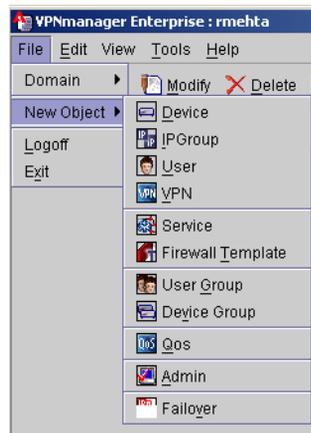
Table 2. New object

Objects	Description
Device	You create a new security gateway within a domain and configure the port interfaces
IPGroup	You configure new IP groups to assign workstations and servers.
User	For each remote user, you configure the name and password for authentication
VPN	To create a virtual private network, you give it a name and select a key management method.
Service	You create services to specify different traffic types.
User Group	You can set up logical groups in which the individual VPN users reside.
Device Group	
QoS	You create a quality of service (QOS) policy to classify and prioritize traffic based on a DSCP value and TCP/IP services and networks.
Admin	
Firewall Templates	You can create new firewall rules to apply to either the domain or the device and set the priority for these rules.

Edit menu

From Edit, you can chose one of the following commands:

- **Delete Object**. Select an object from the VPN diagram and then select **Edit>Delete Object**.

- **Modify Object**. Select an object form the VPN diagram and then select **Edit>Modify Object**.

- **Preferences**. **Edit>Preferences** brings up a window with tabs to select from. for a description of the tabs and how to configure VPNmanager preferences.

View menu

From View, you can select to view the Configuration, the Monitoring Screen, or the Report Wizard.

- **Configuration**. Select **View>Configuration** to open the Configuration Console, or you can click the Config icon on the toolbar. From the Configuration console you can configure and modify the VPN network. See "Configuration Console window" on page 49.

- **Monitoring Screen**. Select **View>Monitoring Screen** to open the Monitoring wizard for the domain that is opened, or you can click the Monitor icon on the toolbar. The Monitor wizard assists you in selecting the various VPN objects you wish to monitor. A number of prebuilt MIB-II and VPNet Enterprise MIB parameter groups can be selected to monitor desired VPN functions, or you can build a custom monitoring group from a comprehensive list of enterprise MIB objects. Examples of ready-to-use groups include an Attack log, Traffic log, security gateway CPU usage, and throughput.You select a type of group to monitor, or you can define a customer group to monitor. See "Using Monitor" on page 283.

- **Report Wizard**. Select **View>Report Wizard** to open Reports, or you can click the Reports icon on the toolbar. The wizard guides you through creating various reports showing details of your network or an object in the network. See "Report Wizard" on page 303.

Tools menu

From Tools, you can access the following commands.

- **Update Devices**. Update Devices is used to update the security gateway configuration with the configuration currently in the Directory Server database.

- **Trace Console**. Trace Console is used to log some debugging information. This information is used by Avaya support to diagnose and troubleshoot any problems that may occur.

Help menu

From Help, you can access the VPNmanager Help, and About VPNmanager.

> *Note: Many of the VPNmanger screens display a "?" icon that, when selected, opens a Help topic relevant to the screen.*

Toolbar

The toolbar on the main VPNmanager screen contains buttons that are shortcuts for the tasks on the Menu bar and the Device Update button.

Figure 5. Icons on toolbar
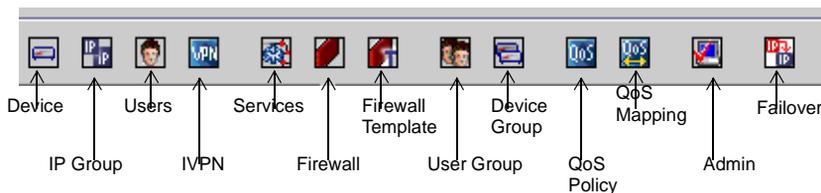
Table 3. Toolbar commands

Toolbar commands	Description
New Object	The New Object button is a shortcut to the *File>New Object* command to create new objects within any of the categories listed in Table 2. When you select one of these commands, either a dialog or a wizard is opened to configure the information.
Modify	The Modify command is used to modify objects from the network diagram view. To use Modify, first select the object to be modified from the network diagram view in the monitor pane and then click Modify.
Delete	Delete is used to delete objects from the network diagram view. To use this Delete, first select the object to be deleted from the network diagram view in the monitor pane and then click Delete.

Table 3. Toolbar commands (Continued)

Toolbar commands	Description
Report	The Report button is a shortcut to the *View>Report Wizard* command that guides you through the steps to create a report about your network.
Config	The Config button is a shortcut to the *View>Configuration* command that opens the Configuration Console dialog. From this dialog you can configure new objects, modify, and view existing content and details about the domain.
Monitor	The Monitor button is a shortcut to the *View>Monitor Screen* command*,* to open the monitoring wizard for the domain that is open.
Update Devices	Update Device is a shortcut to *Tools>Update Devices* used to update the security gateway configuration with the configuration currently in the Directory Server database.

VPN view pane

The VPN view pane is empty until you define your VPN. As devices are configured and added to the VPN, they are displayed in the view pane. The VPN view pane automatically selects one of three presentation types: network diagram view, tiled view, or tree view. The VPN view is determined by the complexity of the VPN. When the VPN contains fewer than six security gateways, a familiar network diagram view is presented. When more than five security gateways exist, the view switches to a tiled display in a vertically scrolling window. Alternately, a third presentation style, the tree view, can also be selected to deal with complex VPNs.

In addition to displaying the individual security gateways in the VPN, a list of Remote Access Users associated with each security gateway is also displayed providing a comprehensive VPN overview at a glance. Double-clicking on an object automatically opens the configuration console details window.

At the top of the VPN View pane is the VPN View selection bar.

VPN view selection toolbar. The VPN View selection bar contains two elements, a list from which the desired VPN is selected, and two radio buttons to select the view styles (Diagram or Tree).

Note: If more than five security gateways are present in the VPN, only the tiled or tree views are available.

All security gateways in the selected VPN selected are displayed, however, only one security gateway can be in focus at any time. The security gateway in focus is indicated by a dashed line around the box and a yellow background.

Remote clients associated with the VPN currently in focus are displayed in a two-column scrolling list box. This list always appears at the top of the Tiled View pane. Clients are listed alphabetically.

Status Icons. The functional status of each security gateway in the VPN is indicated with an icon on the security gateway graphic. A green dot with a checkmark in it means full functionality, while a red dot with an "x" indicates an alarm.

Network Diagram View

In this view, all security gateways, their IP address, associated IP Groups, and a list of all remote client users in the currently selected VPN are displayed in a circular pattern around the Internet cloud which appears in the center. The security gateways are displayed graphically along with a device status icon directly over the security gateway graphic.

Figure 6. VPNmanager Network Diagram View
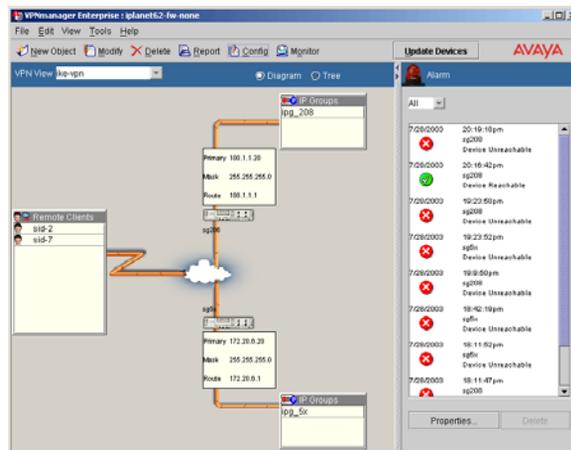
Tiled View

When six or more security gateways are present in the selected VPN, the presentation automatically switches from the diagram view to the tiled view.

Figure 7. VPNmanager, Tiled View
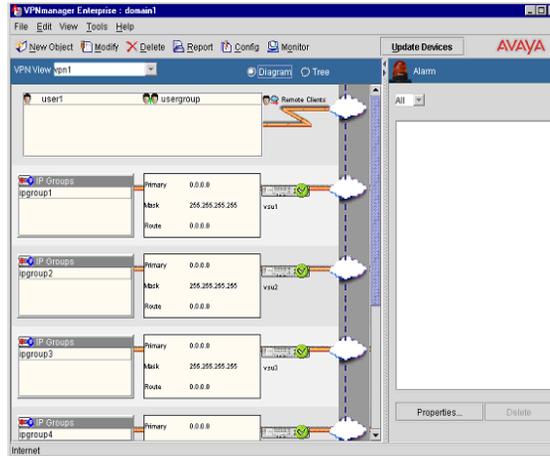
Tree View

An alternative presentation style to the diagram and tiled views, the tree view mimics the Windows-style vertical directory presentation. Its main benefit is that in large or complex VPNs, sections can be collapsed to simplify the view. A [+] or [-] box is displayed to the left of an entry indicating that the entry is collapsed or expanded.

Figure 8. VPNmanager, Tree View
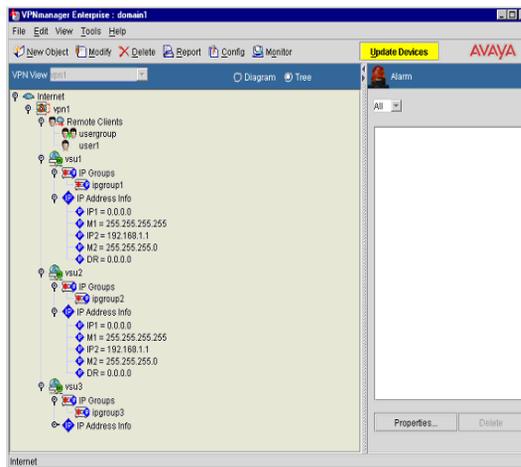
Alarm monitoring pane

To the right of the VPN view pane is the alarm monitor pane. The alarm monitor pane contains summary alarm information, including a time stamp, security gateway name, and alarm type. Alarm information is presented in a vertically scrolling list. A rotating red beacon appears at the top of this screen when a critical alarm is received. See "Monitoring alarms" on page 300.

Configuration Console window

You select *View>Configuration* or click the *Config* icon on the tool bar to open the Configuration Console window. From this window you configure and modify the VPN network configuration. The Configuration Console window includes a menu bar, toolbars, contents pane and a details pane.

Figure 9. Configuration console window
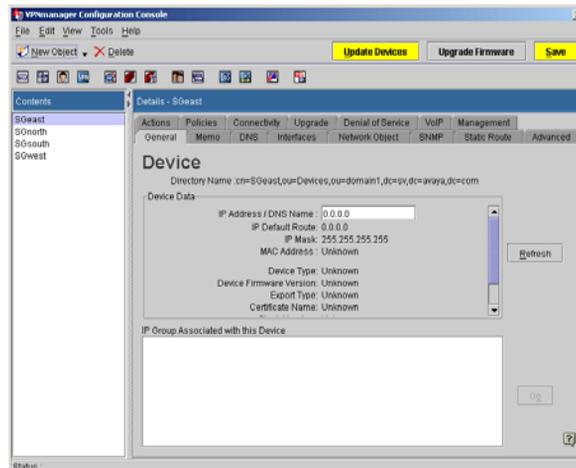
Configuration Console Menu bar

The menu bar on the Configuration Console window includes the following commands File, Edit, View, Tools, and Help.

File menu

The File menu includes the following commands:

- **New Object**. You can create new objects within any of the categories listed in Table 2, "New object," on page 43.

- **Save Changes**. This command saves any changes made through the Configuration Console.

- **Discard changes**. This command clears any changes you have made and reverts the configuration to the last saved version.

- **Close**. This command closes the Configuration Console window.

Edit menu

The Edit menu includes the following commands:

- **Delete Object**. This command deletes the currently selected object.

- **Preferences**. Preferences provides access to global settings for both the machine on which the VPNmanager resides and the domain currently in focus. See "Preferences" on page 53.

View menu

From the View menu, you can view the configured objects, and you can refresh the screen.

Tools menu

The Tools menu consists of functions used for normal VPN maintenance. These functions include the following.

- **Update Devices**. To update the selected security gateway configuration, click **Update Devices**. You select the security gateway to update. This reconfigures all security gateway parameters for the selected gateway and can take several minutes to complete. This function is the same as the *Update Devices* button on the far right side of the toolbar on the VPNmanager main screen. See "Update Devices" on page 53.

- **Upgrade Devices Firmware**. This function is used to download new firmware to selected devices. See "Upgrading a security gateway's firmware" on page 326.

- **Import VPN**. A secure, inter-company extranet can be created by exporting a VPN configuration to a file that is then imported by other VPNmanager installations. See the "Importing and exporting VPN configurations to a device" on page 318

- **Export VPN**. Export VPN can be used to export the VPN configuration which in turn can be imported into other VPNmanger installations.

- **Export RADIUS**. This function is used to export VPN information to an existing RADIUS database. This is primarily for backwards compatibility, but also useful if you wish to convert your existing VPN (using local security gateway-based user authentication) into a dynamic VPN for future scalability. It is, however, expected that LDAP will be the preferred method of building dynamic VPNs.

- **Policies Manager** The *Policies Manager* displays a list from which specific policy services can be selected. Select a service and click **GO** to start the *Policy Manager* for the selected service.

  The types of policies that can be configured depend on the firmware version of the security gateway. Only policies that can be configured are displayed.

  Table 4. lists the policies that could be configured.

Table 4. Policy Services

• Client IP Configuration*	• Syslog*
• My Certificates	• NAT*
• Issuer Certificates	• Packet Filtering
• IKE Certificate Usage	• Firewall*
• RADIUS/ACE	• Client Attributes*

*Policies that can be configured for security gateways with VPNos 4.x. Beginning with VPNos 4.31, the Firewall configuration is not part of Policy Manager.

Toolbar

The toolbar includes the following shortcut buttons.

- **New Object.** You can select one of the icons in the toolbar below New Object and then click New Object to launch the appropriate configuration dialog, or you can click the arrow tip next to New Object and select one of the object types to launch the appropriate configuration dialog.

- **Delete** deletes the selected object.

Contents pane

The Contents pane displays a list of all available members of the object type currently selected.

Details pane

The Details pane displays specific information about the selected object. Details are organized into categories presented as tabs across the top of the screen.

Update Devices

Located in the upper right-hand corner of the VPNmanager Console window is the *Update Devices* button. Use it whenever you make changes to your VPN.

To update the security gateway devices:

1. Make your changes to the VPN.

2. Click **Update Devices** to open the *Update Devices* dialog.

3. Select the security gateways to be updated.

4. Click **OK** to view the status of the update.

   If the *Update Configuration* dialog appears, do the following.

   - In the **User Name** text box, type in the superuser name you configured through the *Console Quick Setup Menu* when the device was being installed. If the device had a firmware upgrade from 3.x, type in **root**.

   - In the **Password** text box, type in the Superuser password configured at the *Console Quick Setup Menu* when the device was being installed. If the device had a firmware upgrade from 3.x, and had an existing security gateway Console password, type in that password. If the security gateway did not have an existing security gateway Console password, type in **password**.

   - Click **OK**.

5. The *Update Devices* dialog will tell you when the update is completed.

Preferences

Preferences provides access to global settings for both the machine on which the VPNmanager resides and the domain currently in focus. Preferences is located in the Edit menu in the VPNmanager Main Console.

When you select Preferences, a series of tabs are displayed. A short description of the tabs follows:

General tab

The Preferences General tab is used to set how you want to save changes on the VPNmanager. You can choose either "Save configuration changes automatically", or "Alert me before saving configuration changes".

Figure 10. Preferences, General Tab
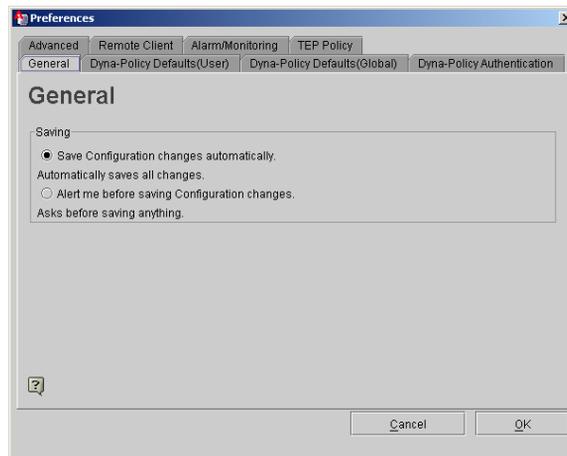
Save Configuration changes automatically When this radio button is active, any changes made to an object are automatically saved upon moving to another object.

Alert me before saving configuration changes When this radio button is active, any changes made to an objects triggers a Save prompt upon attempting to move to another object.

Dyna Policy Defaults (User)

The Dyna Policy Defaults (User) tab is used to define how the Dyna Policy configuration data (VPN session parameters) are handled on the remote user's computer. See

Dyna Policy Defaults (Global)

The Dyna Policy Defaults (Global) tab is used to define the Dyna Policy defaults for the maximum number of login attempts a remote client can make before being locked out for a predetermined time, in minutes. See "Dyna-Policy Defaults (Global) tab" on page 123.

Figure 11. Preferences, Dyna-Policy (Global) Tab
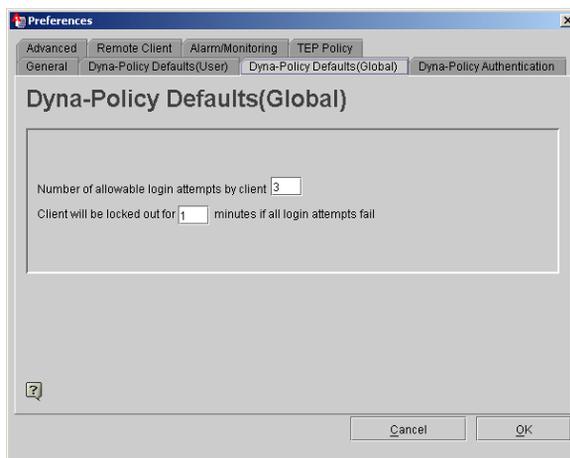
Dyna Policy Authentication

The Dyna Policy Authenticating tab offers a selection of how user authentication and Client Configuration Download (CCD) are performed. Choices are Local (security gateway-based), RADIUS, or LDAP. Whichever method selected is global (across the entire VPN). Selection is made by clicking on the desired radio button.

See Configuring a remote user object for details about configuring Dyna Policy.

Figure 12. Preferences, Dyna-Policy Authentication Tab

Advanced

The Advanced tab is used to either hide or display the LDAP directory context field that appears in a number of places throughout the VPNmanager Console. Users familiar with the LDAP directory structure may prefer having this field displayed.

Figure 13. Preferences, Advanced Tab
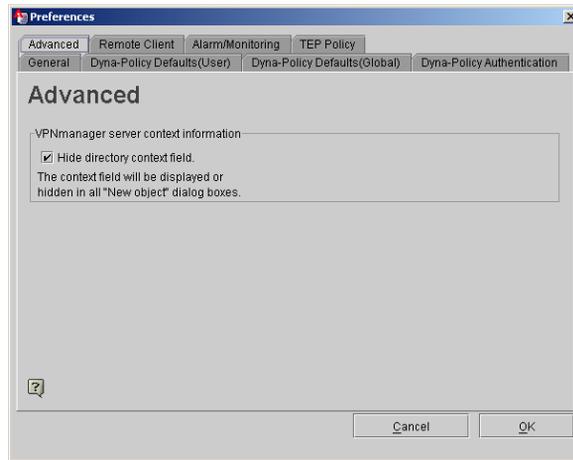
Remote Client

The Remote Client tab is used to establish a path (tunnel) to a secure DNS server to resolve client DNS names (as opposed to using a public DNS server), and to set the remote client user idle time-out period. See "Remote Client tab" on page 125.

Figure 14. Preferences, Remote Client Tab
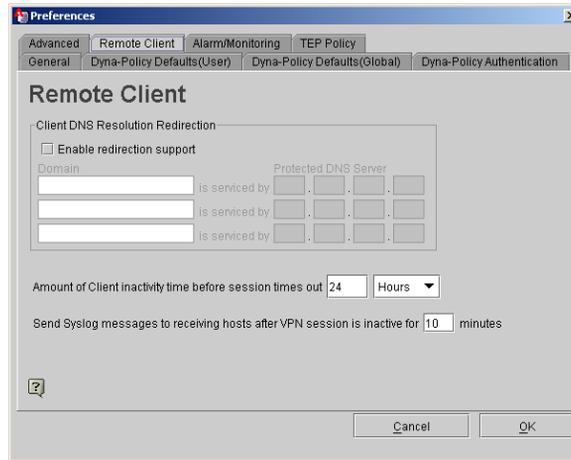
Alarm/Monitoring

The Alarm/Monitoring tab is used to define high-level functions of the alarm console. See "Monitoring alarms" on page 300

Figure 15. Preference, Alarm/Monitoring
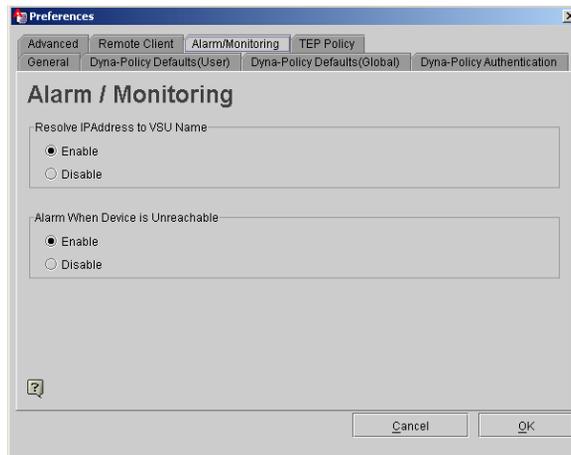
Resolve IP Address to Device Name Enable/Disable. When enabled, the IP address of the alarming security gateway is translated into the security gateway name for display in the Alarm Console. When disabled, only the alarming IP address is shown.

Functionally, all security gateways in all domains are scanned and a map file is created to cross reference the security gateway IP addresses to their respective security gateway names. Default is enabled.

Alarm When Device is Unreachable Enable/Disable. When enabled, this function causes the VPNmanager to continuously send SNMP packets to all security gateways to verify that they are running. If a security gateway fails to respond, an error message is displayed in the Alarm Console.

TEP Policy

The Tunnel End Point (TEP) Policy tab lets you control the security policy applied to the traffic that flows between the end points of a tunnel. The default is off, that is, do not apply configured VPN policies to TEP traffic. See "TEP Policy" on page 242.

Figure 16. Tunnel End Point Policy
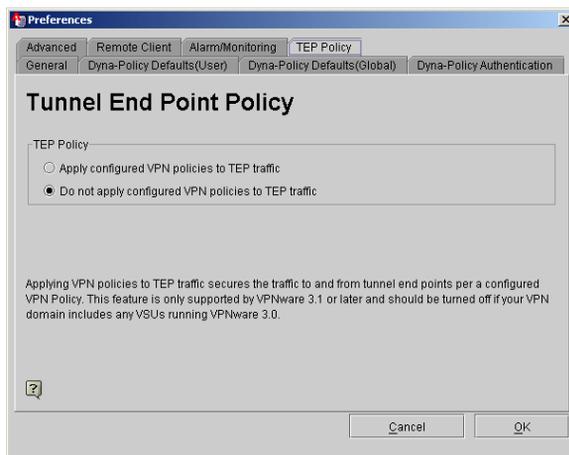
Chapter 3   Setting up the network

This chapter describes the following features that are configured for the domain and the security gateway

- New VPN domain
- Security gateway including:
    - Domain name system resolution
    - Zone interfaces
    - NAT policies
    - Static route table
    - Routing/RIP

New VPN Domain

A domain can be created to meet the networking needs of an entire

organization, or a domain can be created to meet the needs of specific departments of an organization. Existing VPN configurations can be imported into other domains creating interconnected domains.

When you log in to the VPNmanager Console the first time, you must create a domain. You create a domain name and select firewall rules to be applied to the domain, see the "Establishing security" chapter, Firewall rules set up . After the domain is created, you can configure all the objects that are contained in the domain.

To create a new domain:

1. From the *VPNmanager Console main* window menu, select **File>Domain>New**. The **New Domain** dialog is displayed.

2. In the **Name** text box, type in a name for the domain.

   Names can be up to 255 characters and can use any characters, except a comma (,).

   All VPN components must have unique names. To prevent naming conflicts:

   • Check the names of existing VPNs to avoid duplication.

   • Use organization names (for example, "WorldWideSales_VPN" or "ApplicationsEngineering_VPN") since VPNs usually represent functional organizations within a corporation.

   Note: *Once the domain name is created, you cannot change it.*

3. In the **Security** text box, select the firewall template to be applied to this domain. For detailed information regarding the security policies included in this template, see Establishing security.

Select	Level of security
High	The high security template enforces very strict security policies on the traffic going to and from the security gateway.
Medium	The medium security template enforces strict security policies on the traffic going to and from the security gateway.
Low	The low security template enforces security policies on the traffic going to and from the security gateway.
None	Firewall rules are not enforced. All traffic is permitted into and out of the network.

4. Click **Apply** to create the domain

The name of your new VPN domain appears in the title bar of the VPNmanager Console main window. The domain is open and ready to be configured.

Configuring a security gateway

The New Object>Device function is used to create security gateways and VPN Service Units (VSU) in a VPN environment. The security gateway acts as the end-points of VPN tunnels.

Note: Beginning with VPNmanager 3.4, this configuration guide uses the term"security gateway" to refer to both the security gateway and the VSU. The VPNmanager application uses the term "Device" to refer to both of these components.

In order to configure a security gateway, the security gateway must have an IP address and can be reached over the network. When you select New Object for the device, a setup wizard is launched that allows you to configure the following security gateway functions:

- Name for the security gateway.

- IP address that is used to identify the security gateway to the VPNmanager console.

- SNMP community string. VPNmanager uses the SNMP protocol to monitor the security gateways. See "Using SNMP to monitor the device" on page 277

- Whether the security gateway dynamically builds a routing table using RIP updates. See "Routing/RIP tab" on page 91

- Static routes, if more than one router exists on a network to which the security gateway forwards traffic.

Creating a new security gateway

Before you create and configure the security gateway, make sure that you understand how the features work. Review the information in this chapter and in the Establishing security and the Using advanced features chapters.

To create a new security gateway:

1. From the *VPNmanager Console main window* menu, select **New Object>Device**. The *Setup Wizard* dialog is displayed.

2. Type a name for your new security gateway. Click **Next**.

3. Click **Next**. The **Device Configuration** dialog is displayed. Select one of the following options.

   • Select **New Device Setup** to create a security gateway with a unique configuration. Continue with step 5.

   • Select **Clone setup from another device** to create a security gateway with an existing security gateway configuration. Click **Browse** to select the security gateway to clone. Complete steps 4 and 5 and then go to step 8.

4. Click **Next**, to select which of the following parameters from the selected security gateway to inherit.

   • **Syslog**. Inherits the syslog server addressing

   • **SNMP Community String**. Inherits the SNMP community strings

   • **SNMP Trap Targets**. Inherits the SNMP trap targets

   • **LDAP Servers**. Inherits the LDAP servers addresses

5. Click **Next**, to enter the public IP address information for the security gateway.

   a. Select one of the following:

   • **Unknown**, if the address is not known. The General tab can be used to configure this address at a later time.

   • **IP Address**, to enter the primary IP address of the new security gateway. Optional, add a secondary address if VPNmanager is located on the public network. If VPNmanager is located on the private network the secondary address is required.

   • **DNS Name,** to enter the name of the Domain Name Service of the new security gateway. See "DNS tab" on page 69.

   b. Enter the private IP address and private mask of the private ethernet port

   c. Select **Use this address when directly communicating with this device**, if the VPNmanger is on the private side of the security gateway and needs to communicate using the security gateway's private IP address.

   Note: Entering a security gateway IP address from the VPNmanager Console does not change the security gateway's address. The address and subnet mask of a security gateway can only be changed

with a computer connected directly to the security gateway's console interface. The address entered here is used to identify the security gateway so VPNmanager Console can communicate with it.

6. Click **Next**, the Primary Directory Server address is displayed. Modify the address if it is incorrect.

   • Select **Allow the device to use Server to authenticate administrators**, if the server should be used to authenticate.

   • Select **SSL** to use SSL for auto detection for all communications between the VPNmanager and the security gateway. If SSL is not selected, the SNMP traffic occurs in the clear.

7. Click **Next**. The SNMP Community dialog is displayed. Enter the SNMP community string name to which the new security gateway reports SNMP information. The default is the public community string.

   If an existing security gateway is being added to the VPN, enter the new community string name to which the security gateway is to send its SNMP information.

8. Click **Next**. The Listen/Learn dialog is displayed.

   • In the **Routes** area, select **On**, if you want the security gateway to dynamically build a routing table using RIP updates.

   • In the **Default Route (IRDP)** area, select **On** if you want the security gateway to run IRDP to determine the address of a default router to the Internet. If IRDP is On, from the **Learn from** list, select the port or ports to listen for IRDP updates.

   Note: If both Route and Default Route (IRDP) are on, Routes takes priority over IRDP

   • In the Static Route area, click **Config** to configure the static route destination address. Select **Add** to enter the *IP address of the Next Hop* for the static route. Up to 32 network address/mask paris can be configured for the destination network. Click **Ok**.

9. Click **Next.** Select either to **Setup Now** or to **Setup Later**. Set up later sends the configuration information to the directory server, but not to the security gateway.

10. Click **Finished**. When the *...Setup completed successfully* dialog is displayed, click **Done** to exit the Setup Wizard.

11. Click **Save** to save your work.

When you want to send configurations to one or more security gateway, click **Update Devices** from the Configuration Console window or use the Action tab to send the configuration to the security gateway.

Using Device tabs to configure the security gateway

After the security gateway is set up, the VPNmanager displays the tabs you can use to make changes to the security gateway configuration.

This section describes the features to configure a basic device. See Establishing security and Using advanced features for a description of the other tabs that can be configured.

The tabs displayed are dependent on the VPNos release for the device. Table 5 lists the tabs by release.

Table 5. Device tabs by release

Tab	All VPNos Releases	VPNos 4.0 and earlier	VPNos 4.2	VPNos 4.3
Actions	X			
Advanced	X			
Advanced Action		X		
Connectivity	X			
Denial of Service			X	X
DNS	X			
Failover				X
General	X			
High Availability		X		
Interfaces				X
Memo	X			
Policies	X			
Private port			X	

Table 5. Device tabs by release (Continued)

Tab	All VPNos Releases	VPNos 4.0 and earlier	VPNos 4.2	VPNos 4.3
Resilient Tunnel		X		
Routing/RIP		X		
SNMP	X			
Static Route	X			
Upgrade	X			
VoIP			X	X
Directory Servers		X		

General tab

The Device General tab, Figure 17, displays information specific to the security gateway highlighted in the Contents list. From the General tab you can change the IP address VPNmanager uses to communicate with the security gateway. All other information that is displayed is view only.

Figure 17. Device General tab
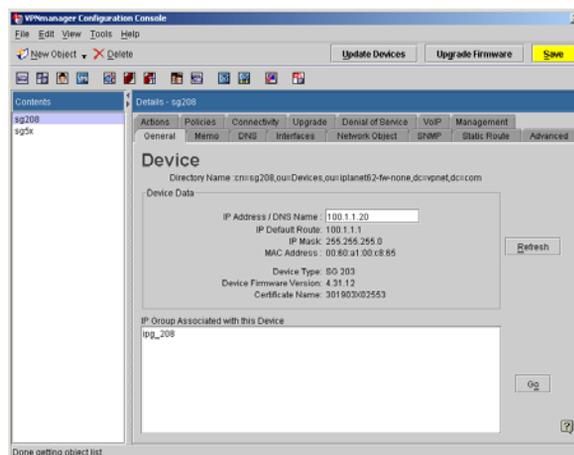
Directory Name The directory name is the location of the security gateway in the directory tree structure.The security gateway name is unique within the VPN domain to which it is assigned.

IP Address/DNS Name VPNmanager uses the address to communicate with the security gateway. This address does not change the security gateway's address. You change the security gateway's address and subnet mask from the security gateway console.

IP Default Route. IP default route is the IP address to the gateway router on the wide area network (WAN).

IP Mask. This is the address mask for the security gateway.

MAC Address. Security gateway MAC Address

Device Type. This shows the model number for the device.

Device Firmware Version. This is the version of firmware running on the device.

Certificate Name. Name of the certificate issuer.

Associated IP Groups area. This area lists the names of the IP groups associated with this security gateway. You can select an IP group from the list and click **Go** to go to the IP Group tab to view the group information.

For VSUs running VPNos 4.0 or earlier, the following additional information is shown.

Export Type. Export type indicates the level of encryption used.

Flash Version. Version of the currently executing NOS from one of two possible flash chips.

FIPS Mode. Federal Information Processing Standards (FIPS) mode indicates if the security gateway is running in the normal or FIPS Level 2 mode. It is recommended that this mode be used only if an organization's policy requires FIPS 140-1 Level 2 certification for cryptographic devices.

The following are not supported in FIPS mode:

- SKIP VPNs

- VPNremote Clients

- Any algorithm other than DES or 3DES,

- Any authentication algorithm other than SHA-1.

RAS. For VSU-100R only. This option is used when dial-in VPNremote users are going to access a security gateway-100R. When enabled, this feature allows the security gateway-100R to support remote clients using VPNremote remote access client software as shipped from the factory. The feature is either enabled or disabled.

Memo tab

The Memo tab is used to record notes about the security gateway, such as change history, physical location, firmware version, etc. This information is stored only in the database and is not downloaded to the security gateway.

To create a memo:

1. From the *Contents* column, select the security gateway you want to configure.

2. Click the **Memo** tab to bring it to the front.

3. In the **Memo** text box, type in any information about the security gateway.

4. When finished, click **Save**.

DNS tab

Use the DNS tab to define where to forward the Domain Name Service (DNS) name resolution requests from the IP devices on the private side of the security gateway.

Figure 18. DNS tab
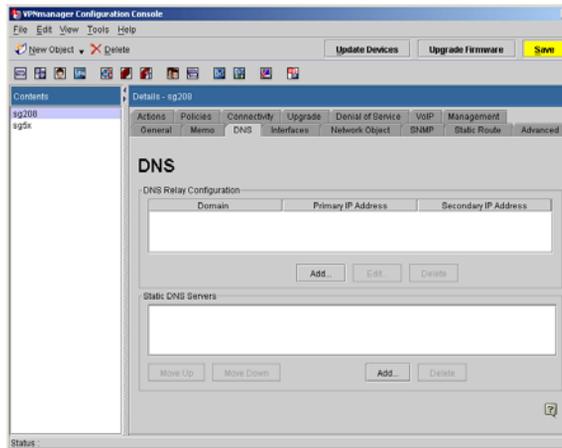
Configuring the DNS tab for security gateways at 4.3 or later

The security gateway includes a DNS name server, and accepts DNS queries from devices on the private side. DHCP devices on the private side receive access to the DNS service automatically. Non-DHCP devices must be manually configured to identify the security gateway as their DNS server. The security gateway server maintains a DNS database on all DHCP clients on the private interface. Non-DHCP clients have no DNS identity.

Note: The security gateway performs DNS relay functionality only for the private zone.

To resolve DNS queries, the security gateway first consults its own database. If this is unsuccessful, the query is forwarded through the public interface. If DNS Relay Configuration domain entries exist, the security gateway tries to find the match of the DNS request domain with the entries' domains. If a match is found, the security gateway only forwards the query to name servers associated with that domain. If no match occurs, the security gateway sequentially forwards the query to the specified static DNS servers. If no static DNS servers exist, queries go to Internet name servers. Note that once static DNS servers are added, Internet root name servers are no longer referenced.

When a DNS server is selected to send the DNS query, and no response is received within a short time, another DNS server is selected by continuing the process as described in the previous paragraph. But if the previous server replies to the DNS query, another DNS server is not selected, regardless of whether response is positive or negative.

By default, when a DHCP client in the private zone sends requests for an IP address and the private zone DHCP server is being used, the DHCP server on the private zone sends its interface IP address as the DNS server in the DHCP response. In this way, all of the DNS queries are automatically forwarded to the security gateway

To add a DNS Relay

To set up DNS Relay Configuration and the static DNS servers. The maximum number of DNS relay rules is 100. You cannot configure Dynamic DNS servers.

*Note: the **Delete, Move Up** and **Move Down** buttons in the DNS Relay Configuration area apply to the IP Address that is currently highlighted.*

1. From the *Configuration Console Contents* column, select the security gateway to be configured. Click the **DNS tab** to bring it to the front.

2. In the **DNS Relay Configuration** area, click **Add**.

3. Enter the **Domain** name and the **Primary IP** address of the DNS server. The secondary IP address is optional.

Figure 19. Add DNS relay configuration
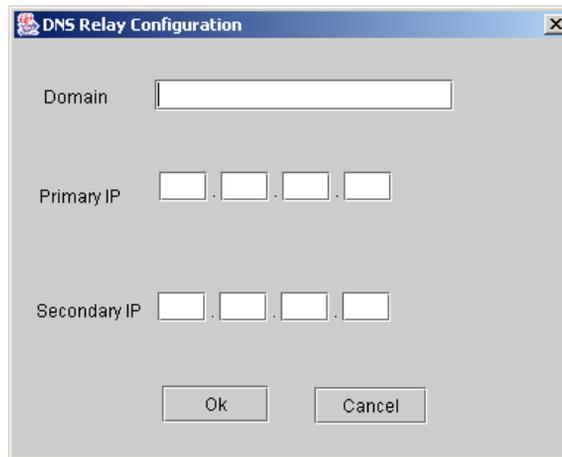
4. Click **OK**.

To add a static DNS server

1. From the *Configuration Console Contents* column, select the security gateway to be configured. Click the **DNS tab** to bring it to the front.

2. In the **Static DNS Servers** area, click **Add**. Enter the IP address of the DNS server and enable the back-up link, if required.

   The backup link is the DNS server that is used when backup ethernet is in use. Only one of the interfaces, either public or public-backup can be in use at the same time.

3. Click **OK**.

   The maximum number of Static DNS servers is four.

Configuring the DNS tab for VSU at VPNos 4.2 or earlier

The VSU can resolve addressing for traffic using the Domain Name Service (DNS). However, the security gateways must know the DNS Server IP address. Up to three server addresses can be referenced by a security gateway. DNS servers can be edited or deleted.

To add a DNS server address

Use Add to enter the initial or backup DNS server(s). Enter the IP address of the DNS server in the "Resolve DNS name with this address" field so that the targeted security gateway can register itself with the DNS server. Click **Apply** to add the new DNS server entry.

1. From the *Contents* column, select the VSU you want to configure.

2. Click the **DNS** tab to bring it to the front.

3. Click **Add** to open the *Add DNS Rule* dialog box.

4. Type the IP address.

5. Click **Apply** to add the IP address to the DNS servers list.

6. Click **Close** to return to the *DNS* tab, or **Apply** to add another address.

7. When finished, click **Save**.

8. When you want to send the configuration to one or more VSUs, click **Update Devices**.

To edit an existing server address:

1. From the *Contents* column, select the security gateway you want to edit.

2. Click the **DNS** tab to bring it to the front.

3. From the *Current DNS Servers list*, select the address you want to change.

4. Click **Edit** to open the *Add DNS Rule* dialog box.

5. Change the IP address.

6. Click **Apply** to add the edited IP address to the DNS servers list. The *Add DNS Rule* dialog box automatically closes.

7. Click **Close** to return to the *DNS* tab. Clicking close ignores any changes made in the *Add DNS Rule* dialog box.

8. Click **Save** to save the change.

9. When you want to send the configuration to one or more VSUs, click **Update Devices**.

To delete a DNS server address:

1. From the *Contents* column, select the security gateway you want to delete.

2. Click the **DNS** tab to bring it to the front.

3. From the *Current DNS Servers list*, select the address you want to delete.

4. Click **Delete** to remove the address.

5. Click **Save** to save the change.

6. When you want to send the configuration to one or more VSUs, click **Update Devices**.

Interfaces tab

For security gateways with VPNos 4.31 or later, the Interface tab is used to edit the configuration of the media interfaces on a security gateway.

When you select the Interfaces tab, the screen displays the available media interfaces, with a summary of their configuration and current status. Scroll to see all the information.

- The name of the media interface
- The zone that is assigned to the media interface
- The IP configuration mode
- The status. Status identifies if the physical link is up or down, and if the interface is being used by network applications
- The IP address
- The mask
- The default route, if relevant
- The MAC address

Figure 20. Interface tab
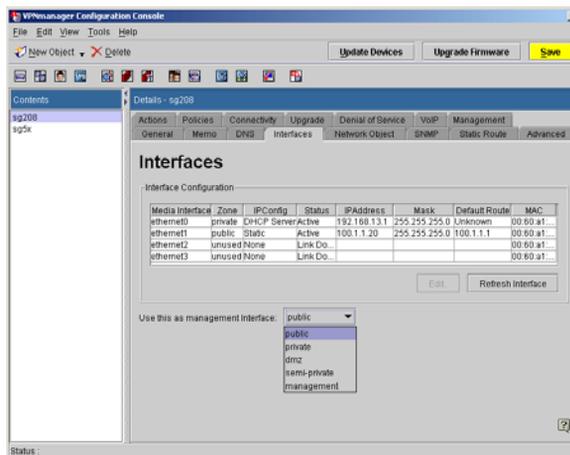
Config Media interfaces can be assigned to one of six different network uses, called *zones*. The number of zones that can be configured depends on the security gateway model (Table 6). Ethernet0 and Ethernet1 are present in all models and are assigned to the public and the private zones. The media interfaces that remain are unused and can be configured as required.

Table 6. Network zones

Media type	SG5 and SG5X	SG200	SG203	SG208
Ethernet0	Public	Public	Private	Private
Ethernet1	Private	Private	Public	Public
Ethernet2	NA	NA	• Unused • Public backup • Private • Semiprivate • DMZ • Management	• Unused • Public backup • Private • Semiprivate • DMZ • Management
Ethernet3 to Ethernet5	NA	NA	• Unused • Public backup • Private • Semiprivate • DMZ • Management	• Unused • Public backup • Private • Semiprivate • DMZ • Management

The following section describes the six network zones.

Public. The public network interface provides connection to the Internet, usually by way of a wide area network (WAN). When VPNmanager is used, the security gateway must be configured with a static IP address. Only one public zone is configured on the security gateway and the configuration for this zone cannot be changed from VPNmanager.

Public-backup The public-backup network interface is used in conjunction with the Failover function on some security gateway models, see "Using advanced features" chapter, Failover to configure failover. If a public-backup network interface is configured, and the public primary network interface cannot reach the Internet, the failover module

deactivates the public primary interface, activates the public-backup interface, and then redirects all encrypted traffic to this link. Only one public-backup zone can be configured on the security gateway.

Note: If the public zone and the public-backup zone are both configured, only one zone can operate at a given time.

Private. The private network interface usually provides connection to your private local area network (LAN) or your corporate LAN. The private network interface can be configured with Static, DHCP Server or DHCP Relay.

Semi-private. The semi-private network interface provides connection to a network whose equipment can be made physically secure, but whose medium is vulnerable to attack, such as a wireless network used within a corporation's private network infrastructure). Traffic on the semi-private interface is usually encrypted. Only one semi-private zone can be configured on the security gateway.

DMZ. The demilitarized zone (DMZ) network interface is usually used to provide Internet users with access to some corporate services without compromising the private network where sensitive information is stored. A DMZ network contains resources such as Web servers, FTP servers, and SMTP (e-mail) servers. Because DMZ networks are vulnerable to attack (that is denial of service), corporations usually add additional security devices such as intrusion detection systems, virus scanners, and so on. Only one DMZ zone can be configured on the device.

Management. The management interface connection can be configured to simplify network deployments, to eliminate enterprise network dependencies on switches or routers. The management network interface is usually used as an access point for a dedicated VPNmanager management station or as a dedicated interface for dumping log messages to a syslog server.

Options for IP addressing for interface zones

You can configure each zone with different addressing options and the private port can be configured as a DHCP server or DHCP relay used to obtain IP addresses from the DHCP server (Table 7). This section explains the options in detail.

Table 7. Type of IP addressing available by zone

	Public	Private	Public-backup	Semi-private	DMZ	Management
Address assigned						
Static	X	X	X	X	X	X
DHCP Client			X			
PPPoE			X			
Server modes						
None	X	X	X	X	X	X
DHCP Server		X		X	X	
DHCP Relay		X		X		

Static addressing

Use static addressing if a dedicated IP address should be assigned to the public interface of the security gateway. To configure static addressing, complete the following information:

Field	Description
IP Address	The public IP address that is assigned to the security gateway
Network Mask	The subnet mask
Route	The IP address of the gateway router to the Internet

DHCP addressing

Use DHCP addressing if the gateway obtains its IP address dynamically from the internet service provider (ISP). This can be configured for public-backup.

Point-to-Point Protocol Over Ethernet (PPPoE) Client

Use PPPoE Client addressing as a convenient way to connect the public-backup zone of the security gateway to the Internet, if your ISP supports PPPoE addressing. PPoE Client addressing requires user authentication. To configure PPPoE addressing, complete the following information

Field	Description
PPPoE User ID	Account user name which your ISP assigns
Password	Account password

Note: Avoid resetting the security gateway by power cycling the unit when PPPoE is configured, as this method requires a proper shutdown in order to avoid a lockout condition during reconnection. This lockout period can last for a few minutes (time varies from ISP to ISP).

Local DHCP Server

The local DHCP server private port configuration is the default configuration to support the IP devices that are connected to your LAN. In the local DHCP server mode, the protected devices are automatically provided with an IP address, a default route, a domain name (the security gateway), and WINS.

To configure the local DHCP server, complete the following information:

Field	Description
IP Address	The IP address assigned. The default IP address is 192.168.1.1 for the private interface. If multiple interfaces on a security gateway have DHCP server configured, their IP addresses must be unique.
IP Range From/ To	The range of IP addresses that the DHCP server that runs on the interface assigns to DHCP clients. The default DHCP address range for the private interface is 192.168.1.32 to 192.168.1.127. Each security gateway on the VPN requires a unique DHCP range. In addition, if multiple interfaces on a security gateway have DHCP server configured, the DHCP range on each also must be unique.

Field	Description
Domain Name	The domain assigned to the interface. This is only applicable to the private interface. The default for domain name is "private."
IP Device Configuration	This is configured to add support for additional IP devices to the DHCP Server.
IP Telephony Settings	This is optional to configure IP Telephony when IP telephones are connected to the security gateway. See IP Telephony Configuration below.

When DHCP server is configured, you can configure the IP Device and the IP Telephony settings. Click **IP Devices** to display a list of all IP devices that the DHCP server currently supports. The MAC address and IP address are listed, along with information that relates to IP telephony devices

Note: Changing the DHCP Server IP address can result in losing current connectivity with the security gateway.

IP telephone configuration If you are using the security gateway with the Avaya Definity® series of IP Telephones, you must configure the TFTP server IP, the TFTP file path, the Definity Clan IP and the Definity Clan port (See the Definity documentation for further information). Non-Avaya IP telephones require at a minimum, the TFTP server IP address.

The following IP telephone DHCP options are supported:

- Option 150. Proprietary to Avaya IP telephones. This option is for the TFTP server IP address.

- Option 176. Proprietary to Avaya IP telephones. Definity Clan IP address and port along with optional TFTP server IP address (all four fields in the IP Telephony Configuration section must contain entries).

- Option 66. The standard DHCP option for TFTP server.

Note: When you add an IP device, you must also modify the Device Account User.

DHCP Relay

This functionality allows the DHCP Relay agent to bind to the device's private and semi-private interface zones and forward only DHCP requests from the network behind the device to the DHCP server(s) on the public network. DHCP Relay server can reside on either the private, semiprivate, public zones, or another remote network.

The *DHCP Relay* area on the *Interface Configuration* dialog is used to configure the security gateway to support DHCP Relay functionality.

Note: DHCP relay and DHCP server services are mutually exclusive. When the security gateway acts as a DHCP relay, the security gateway cannot also be a DHCP server at the same time.

When the DHCP relay agent receives DHCP client requests from the private or semiprivate interface zones, the DHCP server(s) creates new DHCP messages and forwards the messages to the DHCP server(s) on the public, private, semiprivate zones, or remote networks. The DHCP servers on the public network send DHCP offer messages that contain the IP addresses to the DCHP relay agent. The agent broadcasts the DHCP offer messages to the DHCP clients.

If the DHCP server resides on the remote network, the DHCP server and the DHCP clients must be part of the VPN so that the client can obtain the IP address from the DHCP server.

Static

When you select **Static**, the security gateway is configured with a static IP address and Mask. This is the default configuration. If Static is selected and the VPNmanager is on the private side, then the IP address of the computer running VPNmanager should be statically or dynamically configured through other DHCP server.

Changing network interfaces

From the VPNmanager Console Device Interfaces tab, you can modify the media settings, change the IP information, add an IP device, and configure IP telephony settings. You can configure any zone but Public.

To change the media interface configuration:

1. From the *Configuration Console Contents* column, select the security gateway to be configured. Click the **Interfaces tab** to bring it to the front.

2. Click on the media interface that you want to modify. Click **Edit**. The **Interface Configuration** dialog is displayed.

Figure 21. Media interface configuration dialog
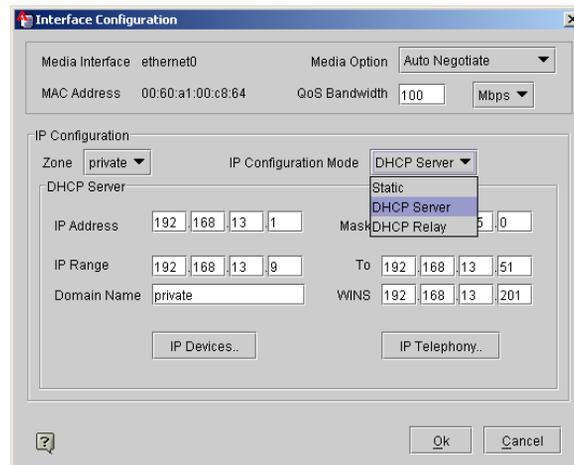
Note: The fields displayed in the screen are based on the type of zone selected.

3. The media option choices depend on the media type selected and the capabilities of the underlying device hardware and driver. QoS is used by the QoS module to restrict the bandwidth of the interface to the upstream limit of the network. For example, to allow QoS to regulate maximum bandwidth of a 100 mbps to 25 mbps, enter 25 mbps.

4. In the **IP Configuration** area, make the required changes.

   • From the **Zone** list, select the zone. Only the zones that apply to that media interface are displayed.

   • From the **IP Config Mode** list, select the IP addressing mode. Depending on your selection, complete the required information.

- If public-backup is selected, complete the **Idle Timer Settings** configuration if failover is enabled.

5. Click **Save** when you finish.

To add an IP device to the security gateway:

1. From the *Configuration Console Contents* column, select the security gateway to be configured. Click the **Interface tab** to bring it to the front property, select the media interface that is configured with private, DHCP Server. Click **Edit**. The Media Interface Configuration dialog is displayed.

2. Click **IP Devices**. The IP Device Configuration dialog is displayed.

3. Enter the following information

   - The MAC address of the IP device. If the device is an Avaya IP telephone, the MAC address is on the back of the telephone.

   - The IP address. This IP address must be within the same subnet as the DHCP server. Avaya recommends that you use an IP address for the device that falls into the DHCP subnet but not in the DHCP range.

4. Click **Add**, and then click **OK**.

5. (Optional) If you are setting up the IP Telephony configuration, click **IP Telephony**. The IP Telephony Settings dialog is displayed.

6. Enter the following information

   - TFTP File Path. The TFTP file path is used when the file path is other than the default path.

   - Definity CLAN Port. The port number for the Definity server. The default port is 1719. The port range is 1 to 65535.

   - TFTP Server IP address. This is the server on which the latest version of the IP telephone firmware is maintained for upgrade purposes.

   - Definity CLAN IP. The IP address of the Definity Clan server.

   - Option 66. The standard DHCP option for TFTP server.

7. Click **OK,** and then click **Save**.

Note: When you configure an IP telephone, secure tunnels are created for TFTP and Definity Clan. However, if only VPN users are connected, the secure tunnels are created on demand. That is, the secure tunnels are created only when traffic exists on the associated tunnel.

Private port tab

For VSUs with VPNos 4.2, the Private Port tab is used to configure of the private IP address. In addition, you can configure the device to act as a DHCP server on the private port or you can configure a DHCP relay.

If a local DHCP server is configured, the security gateway assigns IP addresses to the computers or the IP telephones that are behind the security gateway. If your DHCP server is on the public side, a DHCP relay can be configured to obtain IP addresses from this DHCP server. If the DHCP server is unreachable, the relay can be made to fall back to the local DHCP server.

Figure 22. VSU with VPNos 4.2 Private port
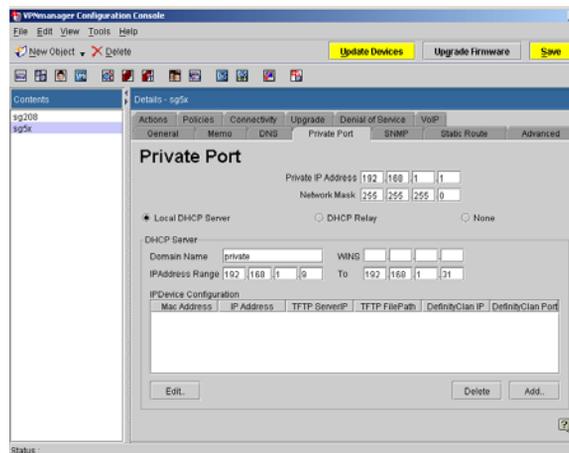
If you plan to use the security gateway's private port local DHCP server capability to support the IP devices connected to your LAN (default), be sure to complete the DHCP setup under the local DHCP Server portion of the screen.

Local DHCP Server. This portion of the screen is used to configure the security gateway as a DHCP server on the private port. The IP Address range must be configured and should fall within the range of the private IP Address subnet. The domain name is provided and the WINS server can be configured.

When deploying the security gateway, you need a unique DHCP range for each security gateway on the VPN.

Note: Changing the DHCP Server IP address may result in losing connectivity to the security gateway, if the VPNmanager is on the private side of the security gateway. Also all active DHCP clients may require renewal through an OS utility (e.g., using winipcfg or ipconfig in Windows), or rebooting.

When changing the DHCP IP address range, execute an ipconfig release and renew command.

IP Devices Configuration. The table displays a list of all IP devices currently supported by the DHCP server. The device MAC Address and IP Address are listed, along with information relating to IP telephony devices, such as the Avaya Definity® IP telephone device information.

Adding an IP Device Configuration

This dialog is used to add IP devices to the virtual DHCP server. The dialog contains a group of fields for IP telephony configuration when IP telephones are connected to the security gateway.

Figure 23. IP Device Configuration
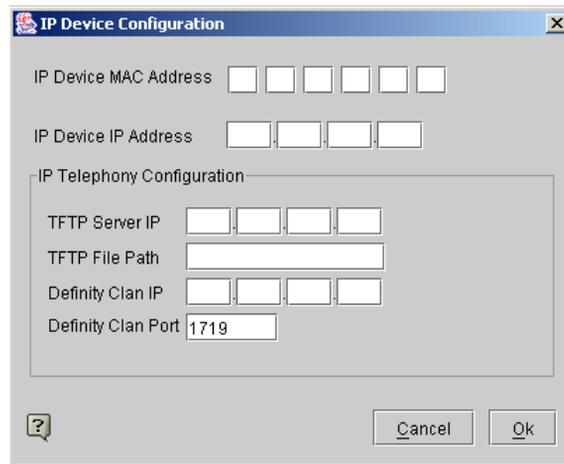
IP Device MAC Address. Enter the MAC address of the IP device. If the device is an Avaya IP telephone, the MAC address can be found on the back of the phone.

IP Device IP Address. This IP address must be within the same subnet as the DHCP server. It is recommended that the IP device address fall in the DHCP subnet, but not in the DHCP range. Also, each IP device should have an unique IP address.

IP Telephony Configuration. This section is used to enter configuration information for an IP telephone connected to the security gateway. This information is sent in response to the IP telephone's DHCP request (this information can also be configured locally in the IP telephone).

The Avaya DEFINITY® series of IP telephones require entries for all four fields (refer to your Definity documentation for further information). Non-Avaya IP telephones require at a minimum, the TFTP server IP address.

Note: The following IP telephone DHCP options are supported:

Option 150: Proprietary to Avaya IP telephones. This option is for the TFTP server IP address.

Option 176: Proprietary to Avaya IP telephones. Definity Clan IP address and port along with optional TFTP server IP address (all four fields in the IP telephony Configuration section must contain entries).

Option 66: Standard DHCP option for TFTP server.

TFTP Server IP. This is the address of the TFTP server on which the latest version of the IP Phone firmware is maintained for upgrade purposes.

TFTP File Path. Used when the file path is other than the default path.

DEFINITY Clan IP. The IP address of the DEFINITY Clan server.

DEFINITY Clan Port. Port number for the DEFINITY server. Default port 1719. Port ranges 1 to 65535.

To add an IP Device:

1. From security gateway Objects, select the **Private Port tab** from the Properties pane.

2. Select the **Local DHCP Server** radio button.

3. Click **Add**.

4. Enter the required information to complete the IP Device configuration.

5. Click **OK**.

6. Click **Save**.

DHCP Relay

Select DHCP Relay to configure the security gateway to support DHCP Relay functionality. This functionality allows the DHCP Relay agent to bind to the device's private port and forwards only DHCP requests from the network behind the device to the DHCP server(s) on the public network.

The IP devices are supported in the case of DHCP relay. To configure the IP devices, from the local DHCP Server configure the IP devices. Return to the DHCP Relay and save.

Note: When the security gateway is acting as a DHCP Relay, the security gateway cannot be a DHCP server at the same time. DHCP Relay and DHCP Server services are mutually exclusive.

When the DHCP Relay agent receives DHCP client requests from the private port, the DHCP server(s) creates new DHCP messages and forwards the messages to the DHCP server(s) on the public network. The DHCP server(s) on the public network sends DHCP offer messages that contain the IP addresses to the DCHP Relay agent. The agent broadcasts the DHCP offer messages to the DHCP clients.

Important: The remote DHCP server(s) and the device's private port IP addresses must be part of the VPN in order for the DHCP Relay process to begin.

The Fallback to Local DHCP Server option allows the DHCP server to revert or fallback to the Local DHCP Server if the DHCP Relay is not functioning.

Note: In order for the security gateway to support the DHCP Relay Fallback feature, Local DHCP Server must be configured. IP Devices are not supported in Fallback mode.

None

Select *None* to configure the security gateway without the Local DHCP Server configuration or the DHCP Relay configuration. None is the default configuration. If None is selected and the VPNmanager is on the private side of the security gateway, then the IP address of the computer running VPNmanager should be statically or dynamically configured through other DHCP servers.

Network Object tab

The *Device>Network Object* tab displays the hosts or networks that are located behind the security gateway. The type of predefined network objects that are listed depends on the type of zones that are configured for the security gateway.

By default, the network object includes the IP address and mask that have been configured for the corresponding zone. Besides this address, you can add additional addresses.

Select a network object and click Add to configure additional IP addresses and mask.

Figure 24. Device Network Objects tab
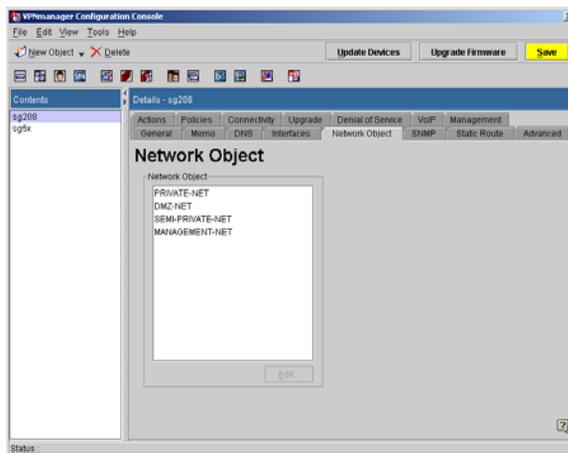
Static Route tab

Static Routes are specified when more than one router exists on a network to which the security gateway must forward either VPN or non-VPN traffic.

The Static Route tab shows the static routes, including the IP address of the hop and the IP address nd network mask pairs for this hop. You can add, modify, and delete routes.

Figure 25. The Static Route tab for a security gateway Object
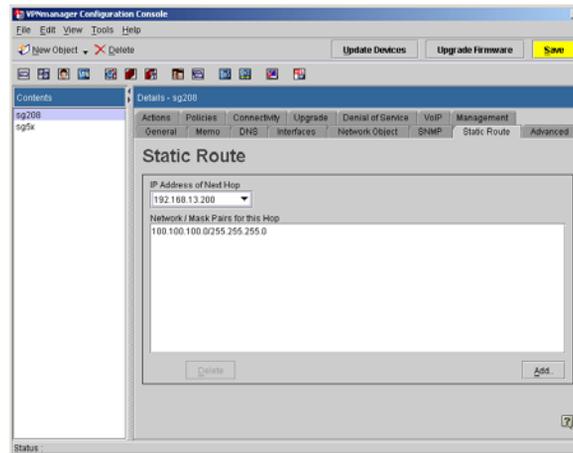
The *IP Address Next Hop* is a list that displays the IP address of the next hop routers (from the security gateway in focus).

The Network/Mask Pairs for this Hop list indicates the static route destination address. You can build a static route table with up to 32 network address/mask paris. This limit allows for any combination ranging from a single router with 32 network address/mask pairs to 32 routers with a single address/mask pairs.

To build a static routing table:

1. From the *Configuration Console Contents* column, select the security gateway you want to configure.

2. Click the **Static Route** tab to bring it to the front.

3. Click **Add** to start the *Static Route configuration wizard*.

4. In the **IP Address of Next Hop** field, type in the address of the next router that leads to your other LANs.

5. Click **Add** to put the router's address into the *IP Address of Next Hop* list box.

6. Click **Next** to move to the *Add Network/Mask Ranges for this Next Hop Address* options.

7. In the **Network** field, type in the network address for the LAN that is beyond the next hop router.

8. In the **Mask** text boxes, type in the subnet mask for the network address.

9. Click **Add to List** to put the address/mask pair into the *Current Network/Mask Pairs for this Hop* list box, which also associates the pair with the IP address of the next hop router.

10. Click **Finished** to return to the *Static Route tab*.

11. Click **Save**.

12. When you want to send the configuration to one or more security gateways, click **Update Devices**.

Routing/RIP tab

If the security gateway is in a network with many routers (gateways) to other TCP/IP networks, there can be more than one possible path to a specific router. In that case, routers are probably building routing tables from the information exchanged by a routing protocol. Security gateways can use such protocols to dynamically build a routing table. Figure 26 shows the Routing/RIP tab for a security gateway Object.

Figure 26. The Routing/RIP tab for a security gateway object

To run dynamic routing:

1. From the *Configuration Console Contents* column, select the security gateway you want to configure.

2. Click the **Routing/RIP** tab to bring it to the front.

3. Configure the *Listen/Learn* options. If desired, both options can be run simultaneously with *Routes* taking priority over IRDP.

   • Routes. Select **On** if you want the security gateway to dynamically build a routing table using *RIP updates*.

   • ICMP Router Discovery Protocol (IRDP). Select **On** if you want the security gateway to run *IRDP* to determine the address of a default router (gateway) to the Internet. The protected network uses that address whenever there is no specific route in the table for a destination address.

   • Learn From. This list becomes active when IRDP is selected. This forces the security gateway to listen for IRDP updates at its public port, private port, or both ports.

4. Configure the *Broadcast* options.

    •   Remote Tunnels. If **On** is selected, the security gateway broadcasts VPN routing information from its private port. The information tells listeners to send packets to this security gateway if the destinations are to remote members of the VPN. The security gateway encrypts the packets then sends them to remote members.

    •   Remote Client Address Pool. If **On** is selected, the security gateway broadcasts routing information about remote client address pools. This information tells listeners to send packets to the security gateway if the address is a mapped address. The security gateway translates the mapped address

    *Note: Select **On** if Remote Client Address Pools are created. For additional information, see "Client IP address pool configuration" on page 137.*

5. When finished, click **Save**.

6. When you want to send the configuration to one or more VSUs, click **Update Devices**.

Policy tab, NAT services

Network Address Translation (NAT) is an Internet standard that allows private (nonroutable) networks to connect to public (routable) networks. To connect private networks and public networks, address mapping is performed on a security gateway that is located between the private network and the public network.

Note: Beginning with the VPNmanager 3.2 and the VPNos 4.2 releases, the VPNremote Client 4.1 is supported behind a NAT device (DSL or Broadband Router).

About NAT types for VPNos 4.31

Beginning with VPNos 4.31, you can set the following three types of NAT mapping on the security gateway:

•   **Static NAT**. With Static NAT, addresses from one network are permanently mapped to addresses on another network. One private IP address can be translated to one public IP address. Static NAT is

bidirectional, that is, for outgoing packets, Static NAT translates the source IP address of the packets. For incoming packets, Static NAT translates the destination address of the packets. You must specify both the original address and the translated address to configure Static NAT.

- **Port NAT**. With Port NAT, addresses from internal, nonroutable networks are translated to one routable address in Port NAT. Port numbers, in the case of TCP/UDP packets and sequence numbers and IDs in the case of ICMP packets, are used to create unique channels. Port NAT is unidirectional. That is, Port NAT translates only outgoing packets and not incoming, but it does translate the replies. On the way out, the source address of the packet is translated. For the replies, the destination address is translated back. You can choose from predefined network objects or user-defined network objects, or you can specify the IP address and the Mask for the original address. You must specify the IP address and the port ranges for the translated address. The port ranges must be in a range from 5000 to 65535.

- **Port Redirection**. With port redirection, addresses from a specific address and a specific port are redirected to another address and port. Port redirection translates the destination address of an incoming packet and the source address of the reply. You must specify the from address, the to address, and the port number.

By default, NAT is enabled, and the *Share public address to reach the internet* feature is selected. NAT affects only clear traffic.

Note: If your network contains any nonroutable addresses, Avaya recommends that you enable the Share public address to reach the internet feature. Any firewall rules that are in use can block translated traffic.

Priority of NAT types

NAT is a rule-based policy, where the priority is based on the NAT type and than the order in which the NAT types appear in the NAT list. NAT types have the following priority:

1. Redirection

2. Static NAT

3. Port NAT

Configuring NAT (VPNos 4.31)

Note: You should understand how NAT works before trying to configure NAT for VPNos. This guide does not explain how NAT works.

The NAT screen displays the following information for each rule. Scroll to see all the information.

- The type of rule. The types are static, port, or redirection.

- The zone to which the NAT rule applies.

- The protocol. Protocols are TCP, UDP, TCP/UDP, or ANY.

- The Original IP address/mask.

- The Translation IP address.

- The Start port.

- The End port.

- The status of the rule. Status is enabled or not enabled.

You can add, modify, and delete NAT rules. You can construct a series of rules, and enable or disable each rule as necessary.

A rule can be moved up or down to change the priority.

Enable NAT. NAT is enabled when this box is checked.

NAT List. Note that this is a rule-based policy, where the priority of the rule is the order in which they appear in the NAT List.

Note: For VSUs with firmware version VPNos 4.x, Dynamic mapping cannot be configured.

To add a NAT rule (VPNos 4.31)

1. From the *Configuration Console Contents* column, select the **Policy** tab to bring it to the front. Select **NAT** from the list.

2. Click **GO**. The **NAT Rules** dialog is displayed and the selected device's name should be visible in the Object Names list.

3. From the **Type** list, select either static, port, or redirection. See "Policy tab, NAT services" on page 93.

   Note The screen displays only the fields that must be configured according to the zone and the translation type that you select.

4. In the **Original** area, complete the available or active areas:

   • Option. Select from the list of predefined network objects and user defined network objects or select *Specified*.

   • IP Address. Type the original/from address

   • Mask. Type the mask

   • Port. Type the from TCP/UDP port number. This port number can be from 1 to 65535.

5. In the **Translation** area, complete the areas that are not grayed out

   • Option. Select from the list.

   • IP Address, Type the translated/to address

   • Start Port. Type in the Start port. This port number can be from 5000 to 65535

   • End Port. Type in the End port. This port number can be from 5000 to 65535

6. To enable this NAT rule, select **Enable Rule**.

7. Click **Save.** Close the Policy Manager dialog.

8. From the **Configuration Console**, click **Update Device** to send the configured information to the security gateway.

To edit a NAT rule

1. From the *Configuration Console Contents* column, select the rule that you want to modify. Click **Edit**. The Edit NAT Rule dialog displays.

2. Change the information, following the steps in To add a NAT rule (VPNos 4.31) section.

3. Click **OK** and then click **Save**.

To delete a NAT rule

1. From the *Configuration Console Contents* column, select the rule that you want to delete. Click **Delete**. An information box appears to verify the deletion.

2. Click **OK,** and then click **Save**.

About NAT types for VPNos 3.X

For VPNos 3.X, you can set the following types of NAT mapping on the VSU.

- **Static Mapping** – Addresses from one network are permanently mapped to addresses on another network. Static mapping works when traffic is initiated either inside or outside of the private network.

- **Dynamic Mapping** – Addresses from one network are temporarily mapped to an address from another network. When traffic is initiated from a client on the private network, its address is temporarily mapped to an address selected from a pool of public addresses.

  When the client traffic is idle for a specified period of time, the mapped address is returned to the pool of available addresses. When all public addresses have been assigned, no other private clients can initiate traffic until a public address becomes available.

  Dynamic mapping works only for connections initiated from the private network.

- **Port Mapping** – This option is similar to dynamic mapping except that only one public IP address is required. The security gateway maps every packet from the private network to the public IP address and a source port selected from a predefined range of TCP and UDP port numbers. When traffic is initiated from a client on the private network it is dynamically mapped to the public IP address and an available port number.

  When the client traffic is idle for a specified period of time, the port number is returned to the pool of available port numbers. When all port numbers have been allocated, no other private clients can initiate traffic until a port number becomes available.

  Port mapping works only for connections initiated from the private network. In addition, port mapping works only for TCP and UDP traffic.

NAT applications

Network administrators may choose to use the NAT mechanism for any of the following reasons:

- **Allow access to the Internet from private networks**. Networks which are assigned private addresses, such as 10.0.0.0 (RFC 1918), or addresses that have not been registered must be mapped to public addresses to allow users access to the Internet.

- **Provide support for more hosts with fewer public addresses**. Address mapping allows network administrators to increase the number of hosts that can access the Internet without needing additional registered network addresses.

- **Hide host addresses for security reasons**. Network administrators may choose to use address mapping to hide actual host addresses from the public.

- **Set up VPNs that include overlapping private addresses**. Address mapping allows network administrators to set up VPNs between two sites that use the same private network addresses. For example, both sites may be using 10.0.0.0 private network addresses.

- **Provide support for multi-gateway network configurations**. Address mapping can be used to ensure that request and reply packets enter and exit the network through the same security gateway.

Accessing the Internet from private networks

Figure 27 shows an example of using NAT to allow hosts on a private non-routable or non-registered network to access the Internet.

Figure 27. Access the Internet from private Networks
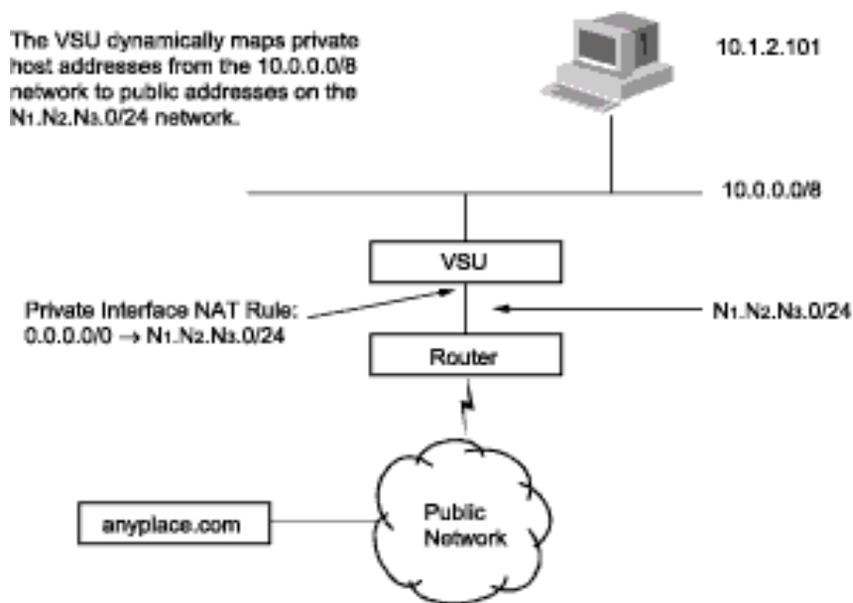
The above example can be used for the following three applications described in the previous section, *NAT applications*:

• Allow access to the Internet from private networks

• Provide support for more hosts with fewer public addresses

• Hide host addresses for security reasons

This configuration allows up to 254 private addresses from the 10.0.0.0/8 network to be dynamically mapped to public addresses from the $N_1.N_2.N_3.0/24$ network.

Each NAT mapping is assigned to an interface. The rules for applying address translations to a packet entering or leaving an interface are:

• When a packet is routed out on an interface (away from the security gateway), the source address of the packet is modified.

• Conversely, when a packet comes in on an interface (toward the security gateway), the destination address of the packet is modified.

In the example shown in Figure 27, when client 10.1.2.101 initially sends a packet to a host on the public network, the security gateway dynamically maps the client's private address 10.1.2.101 to a public address selected from the $N_1.N_2.N_3.0/24$ address pool. Since the packet is going out the public interface, the security gateway changes the packet's source address 10.1.2.101 to its assigned public address $N_1.N_2.N_3.X$.

When the public host receives the packet, it sends a reply to $N_1.N_2.N_3.X$. The reply packet is routed into the security gateway through the public interface, the security gateway changes the packet's destination address back to the client's private address 10.1.2.101 before sending the packet back to the client.

The public address assigned to the client's private address remains in effect until the client traffic is idle for a user-defined period of time. When this idle period is reached, the mapped address is returned to the pool of available addresses. When all public addresses have been assigned, no other private clients can initiate a connection to the public network until a public address becomes available.

One limitation for dynamic mapping is that communication with remote hosts on the public network can only be initiated from clients on the private network. If communication initiated from either the public or private side is required, static address mapping must be used. Static address mapping permanently maps private addresses to their corresponding public addresses, thereby allowing communication between clients and hosts to be initiated from either the private or public network.

Setting up VPN with overlapping private addresses

Figure 28 shows an example of using NAT to set up VPNs between two sites that use the same private network addresses while still allowing private network connections to the Internet. Three NAT rules are applied to each security gateway: one on the private interface, one on the public interface, and one on the VPN tunnel. A DNS entry is also required for each host that can be reached through the tunnel.

The tunnel-mode VPN, named Sales_VPN, provides a secure connection between the SF_Sales_Group and LA_Sales_Group over the public network. Since both sites are using the same private network addresses, NAT mapping must be performed on packets entering and leaving the Sales_VPN tunnel. This is required to ensure that unique host addresses are used on each side of the tunnel.

Communication between a member of the SF_Sales_Group and the server in LA_Sales_Group starts with a DNS lookup of the LA_Sales_Group server address which in this example returns a destination address of 10.0.88.20. The SF_VSU proxy ARPs for 10.0.88.20 by sending its own MAC address in response to an ARP request.

When the packet sent from 10.1.1.17 to 10.0.88.20 enters SF_VSU through the private interface, its destination address is changed from 10.0.88.20 to 172.16.1.20 by applying the NAT rule assigned to the security gateway's private interface.

The SF_VSU performs a VPN lookup and determines that the packet needs to be tunneled to the LA_VSU. Since the packet is leaving the SF_VSU through the Sales_VPN tunnel, the SF_VSU applies the tunnel NAT rule to the packet's source address changing it from 10.1.1.17 to 172.16.0.17. At this point, the packet's source and destination addresses are: 172.16.0.17 -> 172.16.1.20.

The packet is then tunneled across the public network to LA_VSU. Since the packet enters LA_VSU through a tunnel, the NAT rule on the tunnel interface is applied to the packet changing its destination address from 172.16.1.20 to 10.1.2.20, which is the IP address of the LA_Sales_Group server. Before the packet is sent out of the private interface, the NAT rule on the private interface changes the packet's source address from 172.16.0.17 to 10.0.89.17.

Figure 28. Setting Up a VPN with Overlapping private Addresses
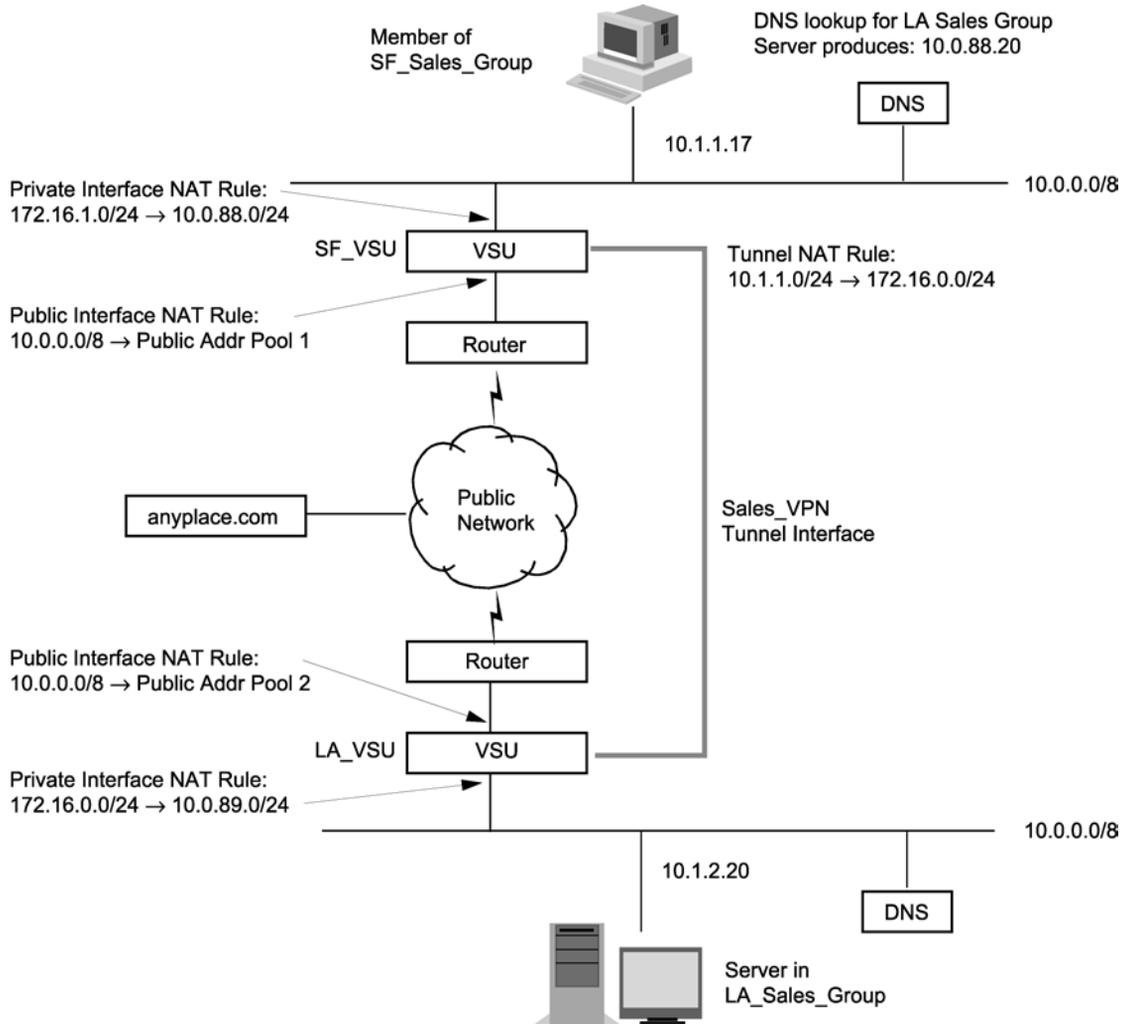
When a reply packet is sent from the LA_Sales_Group server to the LA_VSU, the private interface NAT rule changes the packet's destination address from 10.0.89.17 to 172.16.0.17 and the tunnel NAT rule changes the packet's source address from 10.1.2.20 to 172.16.1.20 before tunneling the packet across the public network to the SF_VSU. At this

point, the reply packet's source and destination addresses are: 172.16.1.20 -> 172.16.0.17.

When the SF_VSU receives the reply packet through the tunnel, the tunnel NAT rule changes the packet's destination address from 172.16.0.17 to 10.1.1.17 and the private interface NAT rule changes the packet's source address from 172.16.1.20 to 10.0.88.20 before the packet is sent out to the SF_Sales_Group client through the private interface.

The NAT rule applied to the public interface on each of the VSUs allows clients on the private networks to access the Internet by mapping their private addresses to public address as described in the previous section *Accessing the Internet from private networks*.

Using NAT to support multiple gateway configurations

Figure 29 shows an example of using NAT to ensure that all replies to packets entering the network through a security gateway exit the network through the same security gateway. The NAT rule applied to the security gateway-B private interface dynamically maps the source IP address of packets sent out the private interface of the security gateway-B to one of 16 addresses assigned to the security gateway-B address pool. Note that the IP address 0.0.0.0/0 matches any packet entering or leaving the security gateway through the designated interface.

When a packet is initially sent from Host A to Host B through the VPN tunnel, security gateway-B dynamically maps the packet source address $(X_1.X_2.X_3.11)$ to an IP address selected from the address pool $(Y_1.Y_2.Y_3.X)$ before sending the packet out the private interface. As a result, reply packets destined for Host A are sent to $Y_1.Y_2.Y_3.X$. security gateway-B proxy ARPs for $Y_1.Y_2.Y_3.X$ by sending its own MAC address in response to an ARP request from Host B. When security gateway-B receives a reply packet on the private interface, it changes the packet's destination address $(Y_1.Y_2.Y_3.X)$ back to the original address $(X_1.X_2.X_3.11)$ before sending the reply to Host A through the VPN tunnel.

A possible alternative to configuring a NAT rule on the private interface of security gateway-B shown in Using NAT to Support Multiple Gateways is to add a static route to the default router which sends packets destined for the $X_1.X_2.X_3.0/24$ network through security gateway-B.

Figure 29. Using NAT to Support Multiple Gateways
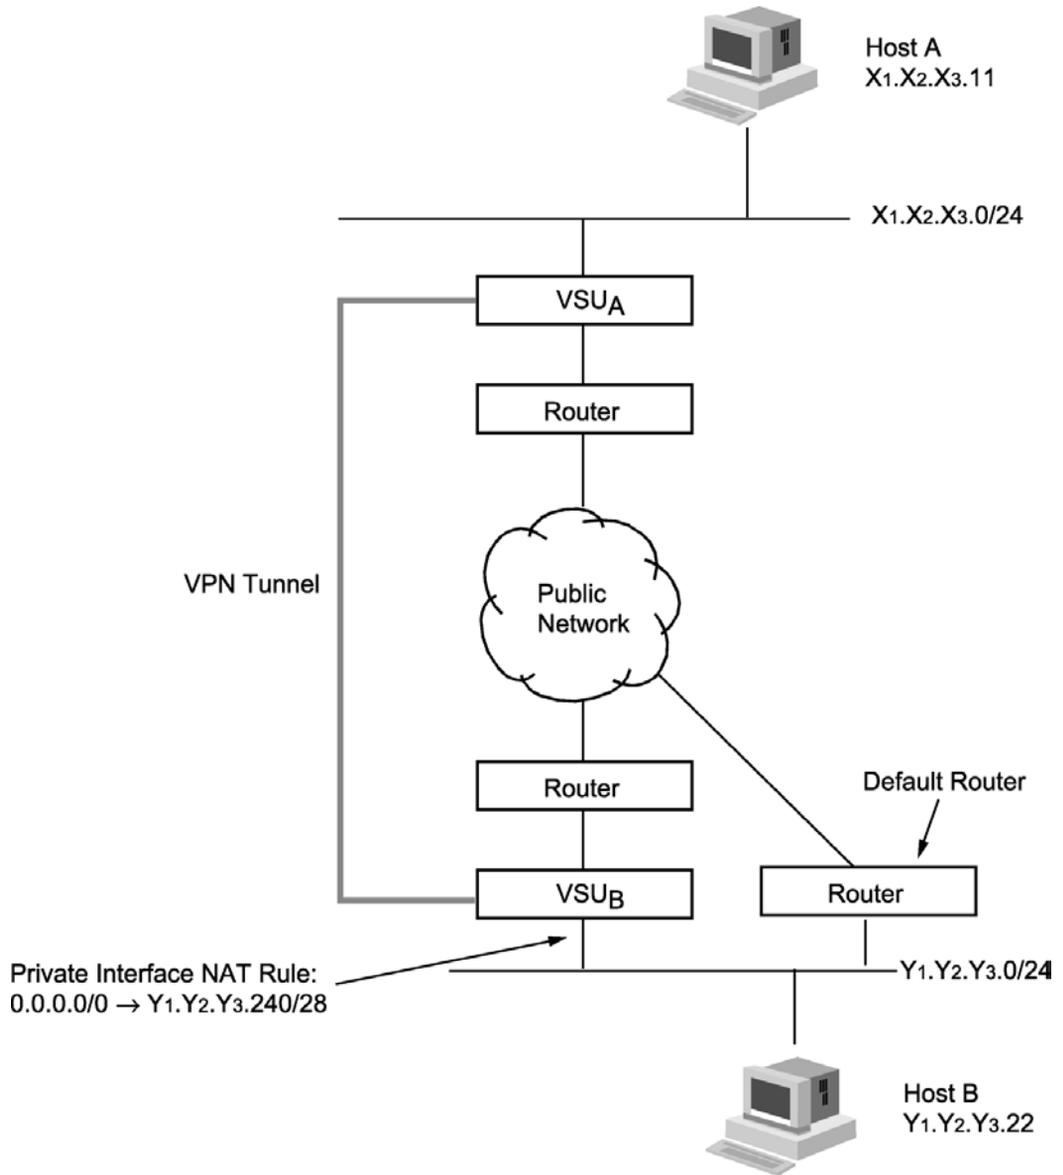
Interface for VPNos 4.2

The following three interface choices are available for devices with VPNos 4.2:

- **Public** – Primarily used to allow clients on a private network to access hosts on the Internet and for transport mode VPNs.

- **Private** – Used to support multiple gateways.

- **Tunnel** – This is a special interface used to support tunneling between overlapping private networks while still allowing connections to the Internet.

Group If you select "Use existing groups," the original address and masks are replaced with the Group selection list.

Original The IP address of the original address and Network/Mask Pair.

Translated Enter the translated address and mask or port range in the Translated fields.

Note: The appropriate fields to use for this translation are enabled based on the Translation Type selected earlier.

Choose where the translation should be inserted in the list on the main Network Translation pane

Add NAT Rule (VPNos 4.2 or earlier)

This function is used to add a new NAT rule to the list.

Translation Type Choices are Static, Dynamic, and Port.

Translation will be applied on Choices are public Interface, private Interface, and Tunnel Interface.

Original

Network/Mask When the Network/Mask Pair selection is made, the IP address of the original address and Network/Mask Pair must be entered.

Translation Enter the Translated Address (and port if the Translated Type is set to port). Enter the Translated Mask.

Locate This Translation Rule Beginning of List, End of List, After Selected Item.

Add this translation rule without enabling it Checking this box allows you to construct a series of rules before actually enabling them.

Memo This area allows you to record notes about this NAT rule in the space provided.

To configure a NAT rule:

1. From the **Configuration Console>Device Contents** pane, select the **Policy** tab to bring it to the front. Select **NAT** from the list. Click **GO**. The NAT Rules dialog is displayed.

2. Click the **Add** to open the *Add NAT Rule* dialog box.

3. From the **Translation Type** list, select a translation type.

4. From the **Translation will be applied on** list, select which interface needs the NAT rule.

5. In the **Original Address** and **Original Mask** text boxes, type in the original address and mask.

6. Do one of the following.
   • In the **Translated Address** and **Translated Mask** text boxes, type in the translated address and mask.
   • If the **Translation Type** is port, type in the **Port Range** in the enabled boxes.

7. From the **Locate This Translation Rule** options, do one of the following.
   • Select **Beginning of List** to put the new rule at the beginning of the NAT Rule list shown in the *Policy Manager for NAT* window.
   • Select **End of List** to put the new rule at the end of the NAT Rule list shown in the *Policy Manager for NAT* window.
   • Select **After Selected Item** to put the new rule after a specific rule that was selected from the NAT Rule list shown in the *Policy Manager for NAT* window.

8. If you want, in the **Memo** text box type in a comment about this rule.

9. If you want to create this rule without making it active, select the **Add this translation rule without enabling it** check box.

10. Click **OK** to return to the *Policy Manager for NAT* window.

11. If you configured a dynamic NAT rule, do the following.

    • From the *NAT Rule* list, select your new rule to highlight it.

    • In the **Translated Address will age out in** text box, type in the number of minutes of undetected traffic that must pass before the assigned translation address is returned to the pool of available addresses.

12. If necessary, use the **Move Down** and **Move Up** buttons to rearrange the position of the new rule in the NAT list.

13. Click **Save**.

14. Close the **Policy Manager** dialog box.

15. From the **Configuration Console**, click **Update Devices** to end configured information to the security gateway.

Chapter 4   Configuring IP Groups

An IP Group is composed of a set of hosts (workstations and servers) that

are located behind a common security gateway. The hosts are defined by their IP address and mask. The security gateway must exist prior to creating IP Groups.

Virtual private networks (VPNs) are made up of IP Groups at multiple locations linked across a public IP network. Assigning workstations and servers to different IP Groups offers a powerful way to limit VPN traffic to specifically designated users.

About IP Groups

Data Terminal Equipment (DTE), such as computers, printers, and network servers, are devices that can be members of a VPN. Two methods are used for creating members. One involves *User Objects*, which is described in "Configuring remote access users" on page 119, but is reserved for creating members that are remote and have to dial into the VPN. The other method involves *IP Group Objects* (or *IP Groups)*, which is reserved for DTEs that are connected to a LAN.

An IP Group contains an *IP address* and *IP mask*. An IP Group can be configured with many of these address/mask pairs. The address/mask pair is used to create an address space (range). Pairs are used for identifying a range of addresses used in a LAN. Therefore, a DTE that has an address within the range of the pair, belongs to a specific IP Group.

IP Groups can be created and edited at anytime. However, since IP Groups are associated with a security gateway, it's recommended that IP Groups are defined after the security gateways is created and configured.

Creating a New IP Group

To create a new IP Group:

1. From the VPNmanager Console main window, click **New Object** and select **IP Group**. The New IP Group dialog is displayed.

2. In the **Name** text box, type in a name for your new IP Group. Any characters can be used, except a comma [,]

   A good practice is to incorporate identifiers in a name so they can be easily managed. For example, a LAN used by an accounting department in San Francisco that is made into an IP Group can be named *SF Accounting LAN*. Using this scheme clearly identifies who are the members of an IP Group.

3. Click **Apply**, then click **Close** to go to the *Configuration Console* window.

   Your new IP Group appears in the Contents column.

4. Click **Save**.

After an IP Group is created, use the *General* and *Memo* tabs to configure it.

New IP Group

The New IP Group screen is displayed when New>IP Group is selected, or when no IP Groups currently exist.

Note: If the Hide directory context field box is unchecked (in the Advanced tab of the Preferences drop-down menu), the Context field is displayed (default = off). This field is used to define where the object is located in the LDAP directory tree.

All VPN components must have unique names. To prevent naming conflicts:

- Add the suffix "group" to the group name.

- Check the names of existing groups to avoid duplication.

- Use department or work group references for group names (for example, "Chicago_Sales_Group" or "Seattle_Engineering_Group") since groups usually represent one or more host devices belonging to employees in a corporate network.

IP Group - General tab

The General tab is used to manage your IP Groups. In addition to displaying a list of all existing IP Groups, it also provides a means of adding new IP Groups and linking the IP Group to a specific device.

Figure 30. IP Group General tab
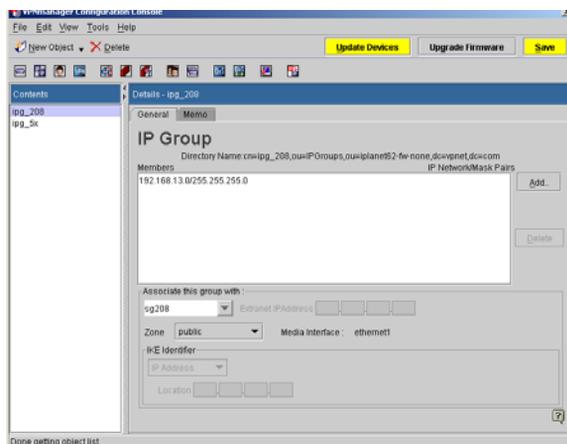
One or more address/mask pairs can be created, and the group can be associated with a specific security gateway. Your new group can even be associated with a security gateway belonging to an extranet, a VPN outside your domain and belonging to another organization, such as suppliers, banks, or customers. This tab includes the following information.

Members IP Network/Mask Pairs and Ranges. This list shows the IP address and Mask Pairs for all the security gateways currently in the IP Group.

Associate this group with. Associating a group with a security gateway means that the hosts corresponding to the IP address/mask pair entered are on a network that is behind or protected by the selected security gateway. The list contains the names of all security gateways in the VPNmanager database, a choice of None, and a choice of Extranet device.

Extranet device. You can create a group associated with a security gateway that is not managed by your company's VPNmanager. You do this to create "extranets" or VPNs between partner companies. In an extranet, each company network uses VPN components that are managed separately by their respective system administrators.

If you are creating an extranet, choose *Extranet device* as the group's associated security gateway. Doing this enables the "Extranet IP Address" entry field. The IP address of our partner company's security gateway is entered here.

The IKE Identifier box is also activated when *Extranet devices* selected.

Zones. This is the zone that is used. The default is public. For Avaya SG203 and SG208 security gateways, if the semi-private zone is configured, it can be selected.

IKE Identifier. Extranet security gateway using IKE key management can be based on the following IKE Identifier types:

- IP Address
- DNS Name
- Directory Name
- Email Name

When one of the above is selected, an appropriate field appears in which the information is entered.

Add IP Group member

The Add IP Group Member dialog appears when **Add** is clicked. New member can be added to the current IP Group list.

Depending on the release of VPNos, two options are available in this pane: IP Network address and Mask, or IP Range. For the IP Range, enter the starting and ending IP addresses.

Table 8. Deriving the Group Mask

To specify a contiguous range of this many addresses:	Start from an IP address that meets these specifications:	And use this mask:
1	###.###.###.### (any IP address)	255.255.255.**255**
2	###.###.###.n (n = multiple of 2); e.g., 130.57.4.**2** or 130.57.4.**4**	255.255.255.**254**
4	###.###.###.n (n = multiple of 4); e.g., 130.57.4.**4** or 130.57.4.**8**	255.255.255.**252**
8	###.###.###.n (n = multiple of 8); e.g., 130.57.4.**8** or 130.57.4.**16**	255.255.255.**248**
16	###.###.###.n (n = multiple of 16); e.g., 130.57.4.**16** or 130.57.4.**32**	255.255.255.**240**
32	###.###.###.n (n = multiple of 32); e.g., 130.57.4.**32** or 130.57.4.**64**	255.255.255.**224**
64	###.###.###.n (n = multiple of 64); e.g., 130.57.4.**64** or 130.57.4.**128**	255.255.255.**192**
128	###.###.###.n (n = zero or 128); e.g., 130.57.4.**128**	255.255.255.**128**
256	###.###.###.0 (n = zero); e.g., 130.57.4.**0**	255.255.**255.0**
512	###.###.n.0 (n = multiple of 2); e.g., 130.57.**2**.0 or 130.57.**4**.0	255.255.**254**.0

Table 8. Deriving the Group Mask (Continued)

To specify a contiguous range of this many addresses:	Start from an IP address that meets these specifications:	And use this mask:
1024	###.###.n.0 (n = multiple of 4); e.g., 130.57.**4**.0 or 130.57.**8**.0	255.255.**252**.0
2048	###.###.n.0 (n = multiple of 8); e.g., 130.57.**8**.0 or 130.57.**16**.0	255.255.**248**.0
4096	###.###.n.0 (n = multiple of 16); e.g., 130.57.**16**.0 or 130.57.**32**.0	255.255.**240**.0
8192	###.###.n.0 (n = multiple of 32); e.g., 130.57.**32**.0 or 130.57.**64**.0	255.255.**224**.0
16384	###.###.n.0 (n = multiple of 64); e.g., 130.57.**64**.0 or 130.57.**128**.0	255.255.**192**.0
32768	###.###.n.0 (n = zero or 128); e.g., 130.57.**128**.0	255.255.**128**.0
65536	###.###.n.0 (n = zero); e.g., 130.57.**0**.0	255.255.**0.0**
Etc.		

Configuring an IP Group

To configure an IP Group that communicates within its own VPN domain:

1. Select the IP Group to be configured. Click the **General tab** to bring it to the front.

2. Click **Add**. The Add IP Group dialog is displayed.

3. Configure the address/mask pair.

   - **New IP Network**. Type in the network address for a LAN.

   - **New IP Mask**. Type in a mask to define the range of addresses that will become members of the IP Group. The larger the mask, the smaller and more focused the address range will be. The method is just like masking a subnet.

   The address/mask pair can be as simple as the network address for a specific LAN and its subnet mask. In that case, all the addresses in the LAN become members of the IP Group. Or, the pair can use the network address, but with a larger mask (more bits) to reduce the range of the address space, so that a smaller range of addresses become members.

4. Click **Apply**, then **Close** to return to the General tab.

   Your new pair appears in the Members list.

5. From the **Associate this group with** area, select a security gateway that the group must be associated with.

   The security gateway selected should be one that is protecting the LAN containing the IP Group.

6. Click **Save**.

7. (Optional) Go to the Memo tab to make a note about this IP Group.

Configuring an IP Group that connects to an extranet

Typically, IP Groups are associated with security gateways that belong to the same VPN domain. However, IP Groups can also be associated with security gateways that belong to other VPN domains. For example, IP Groups can be associated with your organization's customers, suppliers, or to other IKE/IPSec compatible devices.

Note: For a detailed explanation about extranets, see "Exporting a VPN object to an extranet" on page 182

To configure an IP Group that is associated with an extranet:

1. From the **VPNmanager Console main** window, click the IP Group icon from the Icon toolbar. The Contents column displays a list of existing IP Groups.

2. From the **Contents** column, select the *IP Group* to be configured. Click the **General tab** to bring it to the front.

3. Click **Add**. The Add IP Group dialog is displayed.

4. Configure the address/mask pair.

   - **New IP Network**. Type in the network address for a LAN.

   - **New IP Mask**. Type in a mask to define the range of addresses that will become members of the IP Group. The larger the mask, the smaller and more focused the address range will be. The method is just like masking a subnet.

5. Click **Apply**, then **Close** to return to the General tab.

   Your new pair appears in the Members list.

6. From the **Associate this group with** area, select Extranet device.

   The security gateway selected should be one that is protecting the LAN containing the IP Group.

7. In the **Extranet IP Address** box, type the IP address of the security gateway that belongs to the extranet.

8. From the **IKE Identifier** drop-down list select a method for identifying the extranet's device. The device must be an IKE/IPSec compatible device.

   - Select **IP Address** if the extranet's device identifies itself by using an IP address. In the Location text boxes, type in its IP address.

   - Select **DNS Name** if the extranet's device identifies itself by using a DNS name. In the **Name** text box, type in the host name of the device

- Select **Directory Name** if the extranet's device identifies itself by using a Directory Server name. In the **Name** text box, type in the directory name of the device.

- Select **E-mail ID** if the extranet's device identifies itself by using an IP address. In the **Name** text box, type in the e-mail address of the extranet's device.

9. Click **Save** to save your work.

10. (Optional.) Use the Memo tab to make a note about this IP Group.

Delete

Click **Delete** to delete the highlighted IP Group from the Contents list.

Memo

Memo can be used to record notes about the IP Group, such as change history, where the group is located, etc. Information entered here is associated only with the security gateway in focus. This information is stored only in the database and not downloaded to the security gateway.

Chapter 5   Configuring remote access users

VPNRemote™ Client users who log in to the VPN through the security

gateway must have their user authentication configured on the security
gateway. User objects are used for creating remote users. Those remote
users connect to the VPN through an ISP (Internet Service Provider).

Each user is defined by a name, password, and dyna-policy distribution
and authentication method.

As a minimum, you must configure the user name and the password for
each remote user. The dyna-policy can be defined globally for all users on
the VPN or you can define them for individual remote users.

This chapter describes how to:

- Configure a default client configuration

- Create new remote users

- Configuring a dyna-policy, either global or for individual users

- Establish a path to a secure DNS server to resolve client DNS names

- Use Policy Manager to configure client IP address pools, Radius/ACE
  authentication and create a legal notice for users

- Define the type of IKE identifier associated with a user

Default client configuration

When you create a domain with VPNmanager, a default *client configuration download* (CCD) is configured that can be shared by the users. Using the default client configuration makes it faster to configure new user parameters. The default configuration can be changed as required by your specific security and authentication requirements.

The VPNmanager Preferences property includes three tabs, Dyna-Policy Default (Users), Dyna-Policy Default (Global), and Dyna-Policy Authentication that are configured with the dyna-policy parameters. The parameters can be changed any time. This configuration is the default dyna-policy for all users. When you create new users, if the user should not use the CCD, you must check *Do not use default Dyna-Policy* on the User Dyna-Policy tab.

Using dyna-policy

The VPNremote client uses a *Dyna-Policy* when communicating with a VPN. The dyna-policy tells the VPNremote client which authentication and dyna-policy must be used and the topology of the VPN.

A dyna-policy can be configured for either globally for all users on the domain or for individual users.

The global dyna-policy is configured from the VPNmanager Preferences property and is automatically distributed to the VPNremote Client. The automatic distribution method is called *Client Configuration Download (CCD)*. The security gateways distributes the Dyna-Policy when VPNremote Client connects to the VPN.

An individual dyna-policy is configured from the user object, dyna-policy tab and is manually distributed to the VPNremote Client. The manual distribution method involves a three step process.

From within a specific User object, you create a dyna-policy file.

• The file is then delivered, for example, by e-mail, to the user of the VPNremote Client. Although the file can be password protected, the file is encrypted using *DES* (*Data Encryption Standard).*

• The user then runs VPNremote Client to install the dyna-policy file.

The RSA SecurID New PIN and Next Token CCD modes are supported.

Figure 31. User Dyna-Policy tab
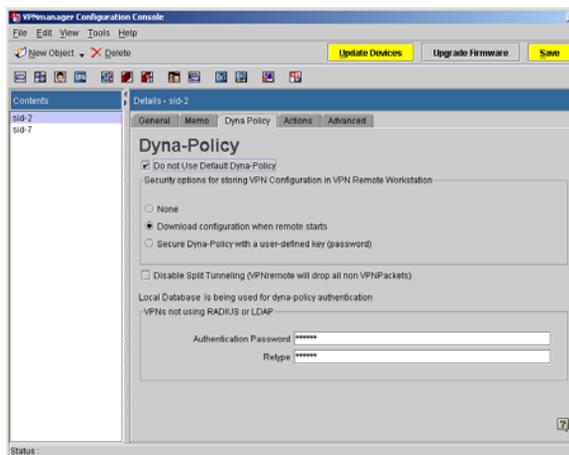
Configuring a global dyna-policy

You configure the global CCD from the *Preferences* property sheet. You should set up the default global CCD before you configure user objects. The parameters can be changed any time.

You configure the following Preferences property tabs to create a global dyna-policy:

- Dyna-Policy Defaults (User)

- Dyna-Policy Defaults (Global)

- Dyna-Policy Authentication

- Remote Client

The following describes each of the tabs. For the procedure to configure a default CCD see "Configure a default CCD with global dyna-policy" on page 128.

Dyna-Policy Defaults (User) tab

The *Preferences Dyna-Policy Defaults (User)* tab is used to define how the remote user's computer handles the dyna-policy configuration data (VPN session parameters).

Figure 32. Preferences, Dyna-Policy Defaults (User) tab
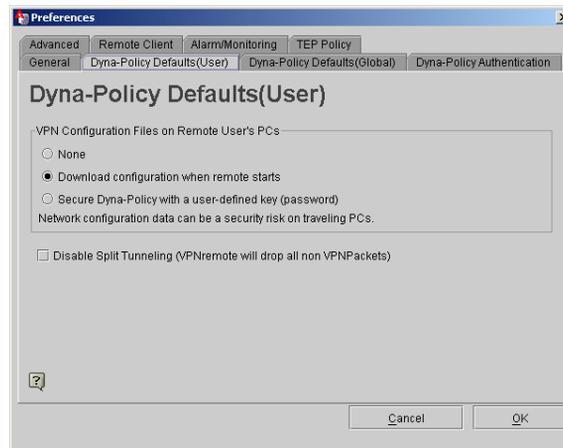
VPN configuration files on remote user's computer

- **None**. The VPN session parameter information is stored locally on the remote users computer. No password is required when VPNremote is subsequently launched.

- **Download configuration when remote starts**. VPN session parameter data is downloaded over the network to the remote computer at the beginning of every session, and purged when the session is terminated (most secure method).

- **Secure Dyna-Policy with a user defined key (password)**. VPN session parameter data resides on the remote users hard disk and are activated by a password at the start of a VPN session. The remote user is responsible for password protecting this data.

Disable split tunneling

Split tunneling allows the VPNremote Client to simultaneously maintain both a VPN (secure) connection and a clear connection. This is the default. You must check the *Disable Split Tunneling* check box to turn the default off. When the default is off, only secure VPN traffic from the VPNremote client computer is allowed.

How you configure this function depends on your corporate security policy. With the default setting to allow split tunneling, a typical application might be when the client wishes to explore a public website while maintaining an email connection on the (private) corporate network. In a security-conscious organization where there is a perceived risk of intrusion into the private network through a remote client's public connection, split tunneling would be disabled.

When split tunneling is disabled, the remote users see a message on their VPNremote Client console indicating Connected - Private access only.

Dyna-Policy Defaults (Global) tab

The *Preferences Dyna-Policy Defaults (Global)* tab is used to define the dyna-policy defaults for the number of times a user can enter an incorrect password before log on fails and the number of minutes that a user is locked out after the password fails.

Figure 33. Preferences, Dyna-Policy (Global) tab
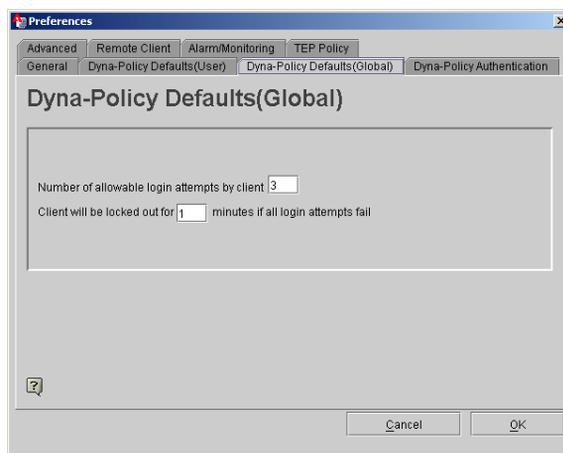
Dyna-Policy Authentication tab

The *Preferences Dyna-Policy Authentication* tab is used to define how user authentication and Client Configuration Download (CCD) are performed. Choices are Local (security gateway-based), RADIUS, or

LDAP. Whichever method you selected becomes the global used across the entire VPN.

Figure 34. Preferences, Dyna-Policy Authentication Tab
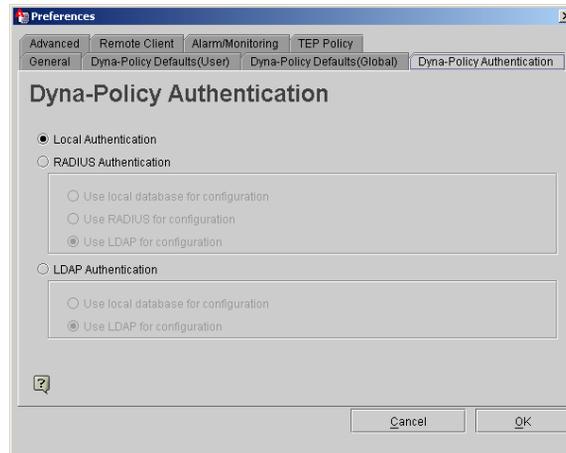
Local authentication

Local authentication is used in non-dynamic VPNs, that is VPNs that are not using RADIUS or a directory server as the authentication database. The user is authenticated from the database stored in the security gateway's flash memory. This is the default.

RADIUS authentication

(VPNos 3.x and VPNos 4.31) RADIUS authentication uses an existing RADIUS database for user authentication.

When this option is selected, you must choose how the remote client configuration download (CCD) is handled:

• by the security gateway, "Use local database for configuration"

• by the RADIUS server, "Use RADIUS for configuration" (VPnos 3.x only)

• by the directory server, Use LDAP for configuration" (VPNos 3.x only)

LDAP authentication

Note: This feature is only available for VPNos 3.x, when iPlanet Directory Server is supported.

LDAP authentication uses the designated directory server database for user authentication.

As with RADIUS, you must also choose how the remote client configuration download (CCD) is handled:

- By the security gateway, "Use local database for configuration"

- By the directory server, "Use LDAP for configuration"

Dynamic VPNs (VPNos 3.x)

Dynamic VPNs is a term given to VPNs that are readily scalable by maintaining the remote client database on a RADIUS server (as opposed to maintaining this data in the local security gateway). This method avoids any size limitations on the number of remote users due to security gateway Flash memory restrictions.

Depending on your security policy, you may wish to have the VPN session client configuration download file (CCD, part of dyna-policy) reside in the security gateway while remote client authentication occurs via the RADIUS database.

Remote Client tab

The *Preferences Remote Client* tab is used to establish a path (tunnel) to a secure DNS server to resolve client DNS names (as opposed to using a public DNS server) and to set the remote client idle time-out period.

Figure 35. Preferences, Remote Client Tab
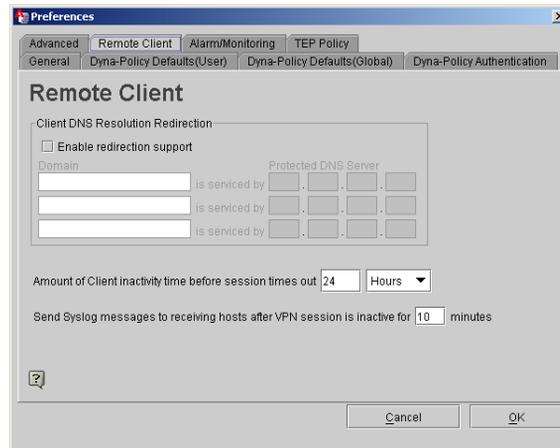
Client DNS resolution redirection

By using the Client DNS Resolution Redirection feature, VPNremote Client-initiated DNS name resolution requests for specific subdomains can be directed to private DNS servers residing on a network protected by a security gateway. This allows VPNremote Clients to use host names in place of IP addresses when accessing corporate network resources without exposing corporate DNS servers and name resolution databases to the public. Thus, a VPNremote Client can use public DNS servers to resolve public resources and private DNS servers to resolve private resources.

Note: DNS name resolution requests are redirected at the user side by VPNremote Client. The remote Client must be running a version of VPNremote Client software which supports Client DNS Resolution Redirection. Check with Avaya Technologies for version support information.

You can enable Client DNS Resolution Redirection and enter up to three subdomain names along with the IP address of the DNS server that will resolve DNS requests for the corresponding subdomain name.

To configure Client DNS Resolution Redirection for all VPNremote Clients:

- Enter a subdomain name in the Domain Name field (for example, finance.mycompany.com).

- Enter the IP address of the DNS server that will resolve DNS requests for the corresponding subdomain name in the Protected DNS Server field.

- Repeat this procedure for up to two additional subdomains, then click Apply.

These settings apply to all Clients in all VPNs. Client DNS Resolution Redirection cannot be set uniquely for each Client.

For proper operation, a VPN protecting the specified DNS servers must be configured between the VPNremote Client and the security gateway. This VPN must contain a Group that includes the IP addresses of the DNS servers defined within the Client DNS Resolution Redirection. The VPN services of the "DNS server VPN" will be applied to any DNS requests made by the Client to the subdomains defined within the Client DNS Resolution Redirection.

Client DNS resolution redirection

Enable Client DNS Resolution Redirection and enter up to three subdomain names along with the IP address of the DNS server that will resolve DNS requests for the corresponding subdomain name.

Remote Client inactivity connection time-out (VPNos 3.x)

You can set the amount of time that a VPNremote Client can be idle before its assigned IP address is returned to the Client IP Address Pool. This is useful if you have VPNremote Client users that typically use TCP-based applications (e.g. Telnet, FTP, Web traffic) and leave those applications idle for long periods of time.

Units can be seconds or minutes. The maximum idle time is 65,535 minutes.

Send Syslog messages . . .

Send Syslog messages to receiving hosts after VPN session is inactive for XX minutes enables you to set the session inactivity time before issuing a Syslog message. The default time is 10 minutes.

Configure a default CCD with global dyna-policy

The following procedure describes how to configure default dyna-policy parameters.

These commands control how CCD automatically delivers dyna-policies to VPNremote Clients. By default, all user adopt these settings, but they can be rejected and custom configured from the Dyna-Policy tab of a specific user.

1. From the VPNmanager Console main window or from the Configuration Console window, select **Edit Preferences** to open the P*references* property sheet.

2. Click the **Dyna-Policy Defaults (User)** tab to bring it to the front. Select how the VPN session parameters are handled on the user's computer.

   • Select **None** to store the VPN session parameters locally on the remote user's computer. The policy is automatically downloaded to the user's computer the first time that the VPNremote Client is initially connected. The policy is not password protected.

   • Select **Download configuration when remote starts** to automatically download the VPN session parameters at the beginning of every session. The policy is removed when VPNremote client is disconnected. This is the most secure method.

   • Select **Secure Dyna-Policy with a user-defined key (password)** to have the VPN session parameters reside on the user's hard disk and be activated by a password at the start of a VPN session. The user is prompted to create a password to protect the policy.

   • Check **Disable Split Tunneling** if users cannot browse the Internet while they are connected to the VPN.

3. Click the **Dyna-Policy Defaults (Global)** tab to bring it to the front.

    • Enter the number of times a remote user can incorrectly login before they are locked out. The default is 3.

    • Enter the number of minutes a remote user is locked out if all login attempts fail. The default is 1 minute.

4. Click the **Dyna-Policy Authentication** tab to bring it to the front.

    Before CCD begins, remote users must have a *user name* and *password* pair to authenticate themselves. From here, you configure the authentication method to use and where the authentication dyna-policy is stored.

    • Select **Local Authentication** to have the security gateway authenticate the users and to store the authentication policy on the security gateway.

    • Select **RADIUS Authentication** to use a RADIUS server to authenticate users. Select a RADIUS method to store the policy.

        • Select **Use local database for configuration** to store the Dyna-Policies on the security gateway. **Note**: This is the only choice for VPNos 4.31

        • Select **Use RADIUS configuration** to store the Dyna-Policies on a dedicated RADIUS server.

        • Select **Use LDAP for configuration** to store the Dyna-Policies on the Directory Server.

    • (Only with VPNos 3.x with iPlanet Directory Server) Select **LDAP Authentication** to use the directory server to authenticate remote users. Select a method to store the policy.

        • Select **Use local database for configuration** to store the Dyna-Policies on VSUs.

        • Select **Use LDAP for configuration** to store the Dyna-Policies on the Directory Server.

5. Click the **Remote Client** tab to bring it to the front. Configure the pat (tunnel) to a secure DNS server to resolve client DNS names and to set the remote client idle time-out period.

- Check **Enable Redirection Support** if remote clients use private domain names, such as *accounting.avaya.com*, for navigating their VPN. Then enter the Domain and Protected DNS server address

- Enter the number of minutes of inactivity before sessions time out. Default is 4 minutes.

- If Syslog services are running, enter the number of minutes the VPN session can be inactive before a Syslog message is sent. The default is 10 minutes.

6. Click **OK** to save your changes.

After the default parameters have been adjusted to meet your VPN's needs, user can be created.

Creating new user object

A user object is built with either a default or a custom CCD. Using a default CCD speeds up the configuration process, but the existing default CCD might not meet all of your users' requirements.

The New User dialog is used to enter information about a new user. Fields are included for the new user's name, password, and confirmation of password. A default user check box is included to create a default user.

Default user

The Default User feature is normally used in conjunction with the default dyna-policy to establish a common template by which a desired VPN policy type is delivered to the remote clients in the domain. Multiple default users can exist in a domain, but only one default user can exist per VPN in a domain. When a remote user is configured as a default user, the user password is not required to log in. Note that the Default User has a unique icon.

To create a new user object:

1. From the VPNmanager Console main page, click **New Object** and select **Users**. The New User dialog is displayed.

2. In the **Name** text box, type the name of a remote user. Any character, except a comma can be used.

Note: If you plan on using RADIUS as an authentication method, this name must match the name used in the RADIUS server.

3. In the **Password** text box, type the user password for the local, RADIUS, and directory servers.

4. In the **Confirm Password** text box, retype the password.

5. Press **Apply** to save the user name.

6. You can continue to add users, or you can click **Close** to return to the *Configuration Console* window.

About creating individual dynamic-policy

You configure the individual user object from *Configuration Console>User Object*.

User - General tab

The User General tab displays information about the user highlighted in the Contents column, including which VPNs and User Groups the user is a member of.

Figure 36. User General tab
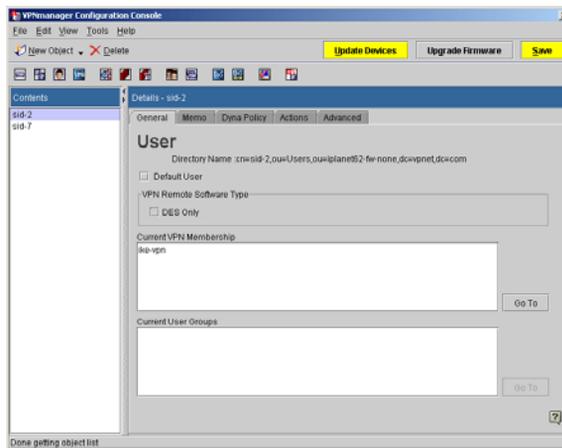
Directory Name. This is the unique users name within the directory structure. It is not duplicated anywhere within the VPN domain to which it is assigned.

VPNRemote Client software type. This area contains a checkbox to select the encryption type on this remote users computer: DES or Any.

When the DES box is unchecked, the encryption being used by the remote client (Any) is accepted. While 3-DES is one of the strongest type of encryption (equivalent to 168 bits), this specific security technology is controlled by the US Department of State and is currently restricted from export, limiting its use to computers used within the U.S. and Canada. Users with laptop computers who occasionally travel out of the country are not subject to this restriction (as long as they return with their laptop).

Checking the box selects 56-bit DES encryption. While not as robust as 3-DES, this encryption standard is not subject to U.S. export restriction.

Note: You cannot place a DES-only user in a VPN that is 3DES.

Current VPN Membership. This section lists VPNs to which the currently highlighted user is assigned membership.

Current User Groups. This displays a list of the User Groups to which the user belongs.

Memo tab

Memo can be used to record notes about the user, such as change history, specific computer type, etc. Information entered here is associated only with this user. This information is stored only in the database and not downloaded to the security gateway.

Dyna-Policy tab

The Dyna-Policy tab is used to define an individual remote user's dyna-policy to specify the security options for how the VPN configuration information is handled on the user's computer.

See "Dyna-Policy Defaults (User) tab" on page 121 for how to configure.

Actions tab

The User Actions tab is used for non-dyna-policy alternatives.

Figure 37. User's Action tab
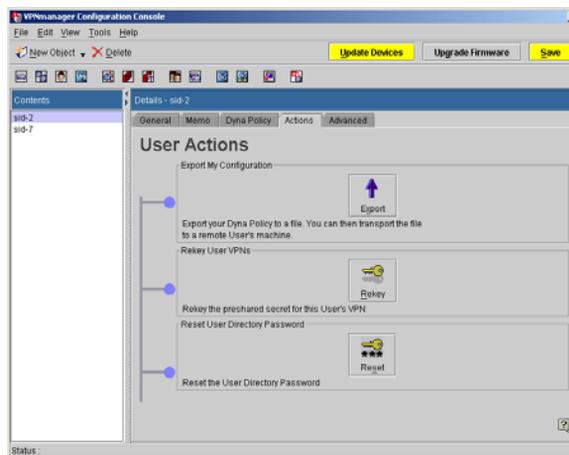
Export My Configuration. Exports your dyna-policy to a file for conveyance to the remote user's machine. Enter a password and retype the password.

Note: If Default User is configured, this button is disabled.

Rekey User VPNs. Clicking the Rekey button causes the preshared secret to be rekeyed for this users VPNs.

Reset User Directory Password. The user's password is reset.

Note: If Default User is configured, this button is disabled.

Advanced tab

The Advanced tab allows you to define the type of IKE identifier associated with the user currently highlighted. Internet Key Exchange (IKE) is a protocol by which a security association (secure tunnel) is established between the security gateway and the remote client.

Figure 38. User Advanced tab
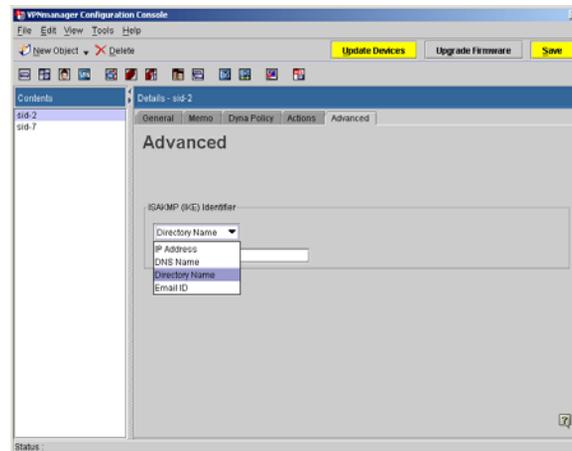
Four types of identifiers can exist in the certificate generated for the remote user.

- Directory Name

- IP Address

- DNS Name

- Email Name (RFC 822)

Configuring a remote user object

If you remote users use the default CCD, you only need to complete steps 1 through 5. If a individual dyna-policy should be created continue with step 6.

1. From the *Configuration Console* window, click **Users** to list all User Objects in the *Contents* column.

2. From the **Contents** column, select the User Object that needs to be configured.

3. From the **General** tab, select the **DES** check box if the VPNremote Client is limited single *DES (Data Encryption Standard)*.

Note: A remote user using single DES encryption can only connect to a VPN using single DES encryption.

4. (Optional) Click the **Memo** tab to bring it to the front, then in the **Memo** text box, type in some information about the user. For example where the user will be dialing from or the location their headquarters.

5. Click the **Dyna Policy** tab to bring it to the front. If you do not want the default Dyna-Policy settings, select **Do Not Use Default Dyna-Policy**. Then configure a customized method for storing the VPN configuration for the user.

   - Select **None** to store the VPN session parameters locally on the remote user's computer. The policy is automatically downloaded to the user's computer the first time that the VPNremote Client is initially connected. The policy is not password protected.

   - Select **Download configuration when remote starts** to automatically download the VPN session parameters at the beginning of every session. The policy is removed when VPNremote client is disconnected.

   - Select **Secure Dyna-Policy with a user-defined key (password)** to have the VPN session parameters reside on the user's hard disk and be activated by a password at the start of a VPN session. The policy is automatically downloaded. The user is prompted to create a password to protect the policy.

   - Check **Disable Split Tunneling** if users cannot browse the Internet while they are connected to the VPN.

6. If *Local Authentication* is used for authentication method, in the **Authentication Password** text box, type in the a password for this VPNremote Client user.

   Note: These text boxes are not available if the RADIUS or LDAP authentication is used. For more information about authentication methods, see "Dyna-Policy Authentication tab" on page 123.

7. If the User object can communicate with an extranet, click the **Advanced** tab to bring it to the front.

   If the method used to identify a remote user is different than within your VPN, use the *IKE identifier* options to configure a method which is used in the extranet. See "Exporting a VPN object to an extranet" on page 182 for information about connecting to an extranet.

After configuring a User object, the user name and password pairs must be given to the VPNremote Client user.

Information for VPNremote Client users

Users who receive their Dyna-Policies by the *Client Configuration Download (CCD)* method must have a user name and password pair. When trying to connect, they use the pair to authenticate themselves. After passing authentication, CCD is used to send the Dyna-Policy to the VPNremote Client. Which pairs to use depends on the authentication method used.

Using local authentication

If the security gateway uses authenticating remote users for CCD, deliver the following pairs to the respective users.

- NAME: The name created in Step 2 of page 130.

- PASSWORD: The password created in Step 3 of page 131

Using RADIUS authentication (VPNos 3.X and VPNos 4.31)

If a RADIUS server is used for authenticating remote users for CCD, deliver the following pairs to the respective users.

- NAME: The name created in Step 2 of page 130

- PASSWORD: The password stored in the user's record of the RADIUS server.

Using LDAP authentication (VPnos 3.X only)

If the directory server is used authenticating remote users for CCD, deliver the following pairs to the respective users.

- NAME: The name created in Step 2 of page 130.

- PASSWORD: The password created in Step 3 of page 130

Using Policy Manager for user configuration

From the VPNmanager Policy Manager property, you can configure the client IP address pool for the remote users and define to users when they log on. (VPNos 3.x and VPNos 4.31 only) You can configure the RADIUS/ACE services.

Client IP address pool configuration

Access control devices (ACD), such as firewalls, guard the networks from unauthorized users. Analyzing source addresses is one method ACDs use to decide which packets can enter the network. ACD is a problem for VPNremote Client users. The addresses which ISPs dynamically assign to VPNremote Clients is naturally blocked because it is impossible to know ahead of time which address is assigned. The security gateway solves this problem by using *Client IP Address Pools*.

A *Client IP Address Pool* is a range of *source* IP addresses that is recognized by an ACD. The pool is stored in the security gateway, so when it recognizes an inbound packet from a VPNremote Client, it swaps the source address with one from the pool. When the security gateway recognizes an outbound packet having a pooled address, it changes the destination address to the remote client's address.

A security gateway can be configured with multiple pools. When selecting a list of *source* addresses to pool, choose ranges that are not used by the *destination* network.

Figure 39. Policy Manager - Client IP address pool
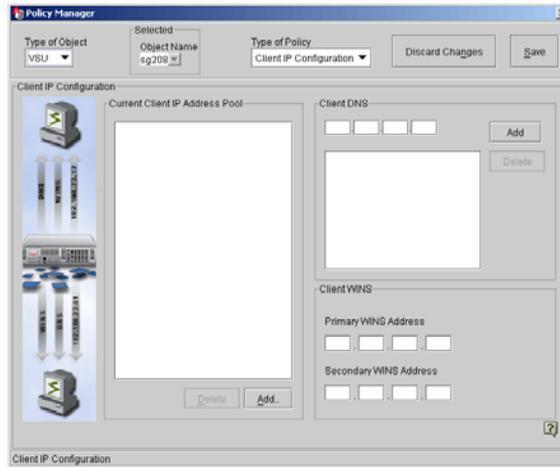
Add Client IP address pool

From the *Policy Manager* properties you select Client IP Configuration to make add new client IP addresses. At the top of the screen is the target security gateway to which this address pool resides.

For VPNos 4.2 and earlier, you enter the starting address of the range in the Client IP Address, Range Start field, followed by the ending address of the range in the Range End field. Up to 20 non-contiguous IP address ranges of any size may be entered (depends on security gateway memory available).

For VPNos 4.31, you enter the IP address and mask.

Add Client DNS

The Client DNS address entered here is sent to the security gateway that is used for the VPNremote virtual adapter configuration. This information is then sent to the VPNremote Client through CCD. Three Client DNS addresses can be configured in the VPNmanager.

Add Client WINS

The Client WINS address entered here is sent to the security gateway that is used for the VPNremote virtual adapter configuration. This information is then sent to the VPNremote Client through CCD. Two Client WINS address can be configured in the VPNmanager.

To configure the Client IP configuration

1. From the *Configure Console* window, go to **Tools>Policy Manager**.

2. From the **Select Object Name** list, select the security gateway to be configured.

3. From the **Type of Policy** list, select **IP Client Configuration**.

4. In the *Current Client IP Address Pool Policy* area, click **Add**.

5. In the **Range Start** text boxes, type in the address for *lower* boundary of the address pool.

6. In the **Range End** text boxes, type in the address for *upper* boundary of the address pool.

7. Click **Apply**. The contents are then cleared from the Add screen allowing for the next entry. Repeat the process until you have entered all required Client IP Address.

8. Click **Close** to return to the *Policy Manager* for *Client IP Address Pools* window.

   The new pool is seen in the *Current Client IP Address Pool* list.

9. (Optional) If a client DNS address should be configured, in the *Client DNS* area enter the DNS address and click **Add**. Up to three client DNS addresses can be configured.

10. (Optional) If Client WINS should be configured, enter the WINS address to use for VPNremote virtual adapter configuration. Two Client WINS addresses can be configured.

11. Click **Save**.

12. Click **Close** to return to the *Configuration Console* window.

13. When you want to send the configuration to the security gateway, click **Update Device**.

Configuring client attributes

From *Policy Manager Client Attributes* property, you can configure a message that remote users see every time they log in and specify the brand name used for VPNRemote Client.

Creating a message

The message you create can be a legal message about company policy for using the network or any other type of message to communicate information when remote users log in. This message can be configured so that remote users are required to accept the message before the log in is complete.

Figure 40. Policy Manager for client attributes
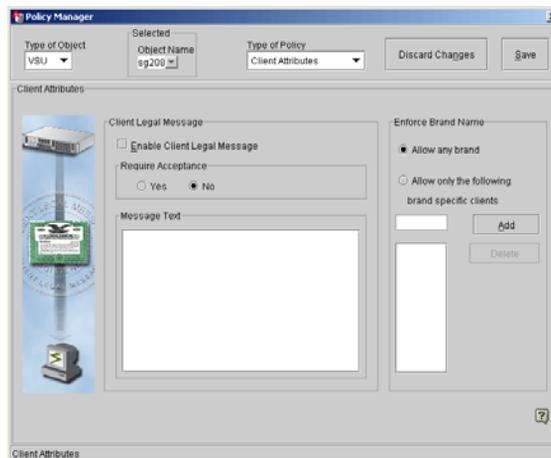
Enable Client Legal Message. The check box is used to enable the Client legal message. The default is disabled.

Require Acceptance. Select **Yes** to require the remote user to accept the message before log on is authenticated. Select **No** if the message is to be displayed, but the remote user is not required to accept the message to authenticate to the security gateway. The default is No.

Message Text. In the Message Text box, type the message that should be displayed. Default messages are not included in the VPNmanager software.

Enforce brand name

VPNmanager allows administrators to restrict access to remote users by specifying client brands. The default is *Allow any brand*. The Administrator can allow any brand name or can restrict access by specifying a brand name. However, in order for this feature to work correctly the brand name must be specified in VPNmanager and in the Avaya VPNremote Client. To customize the Avaya Remote Client, contact your sales representative.

- **Allow Any Brand** allows any brand client to be authenticated by the security gateway during CCD. This radio button is the default.

- **Allow Only the Following Brand Specific Clients** allows clients that have registered brand names with the security gateway to be authenticated during CCD. The Administrator can enter up to five brand specific names for the Client Legal Message to be displayed.

RADIUS/ACE Services

(VPNos 3.x and VPNos 4.31 only)

Note: If a RADIUS server is used, the name assigned to a VPNremote Client must be identical to the one used in the RADIUS server.

A popular tool for managing authentication and accounting for remote access has been *Remote Authentication Dial-In User Service (RADIUS)*. Use the *Policy Manager* for *RADIUS/ACE* if you want to use one or more RADIUS servers to authenticate remote users. A security gateway can query up to three RADIUS servers, where two of the servers is recognized as backups.

Figure 41. The Policy Manager for RADIUS/ACE

Note: The security gateway must authenticate itself to the RADIUS server with a "shared secret" before they can exchange information. Therefore, the RADIUS server must be configured with a shared secret for the security gateway.

Enable RADIUS/ACE

When checked, RADIUS is enabled as the authentication and configuration database.

Rank in group of this particular RADIUS server.

IP Address IP Address of the RADIUS server.

UDP Port UDP port of the RADIUS server. The default value is 1645.

Settings

RADIUS attempts before assuming failure Integer from 1 to 10 indicating the number of attempts the security gateway makes before timing out with a failure. The default is 3.

RADIUS time-out before assuming failure Time in seconds from 10 to 500. This value is the total number of seconds that the security gateway waits for a response from any specified RADIUS server before timing out with a failure. The default is 6 seconds.

RADIUS concepts

For additional user authentication, the VSUs support the Remote Authentication Dial-In User Services (RADIUS) protocol, thus providing stronger Client authentication and accounting mechanisms via third-party products such as Ascend Access Control™ and RSA Security ACE/ Server™ AccessManager.

Using RADIUS, remote users must pass the RADIUS server's authentication mechanism in order to connect to a corporate network. This authentication process is summarized as follows:

- First, the user initiates communication with a VPN member.

- The VPN traffic is processed by VPNremote and then sent to the target security gateway.

- The security gateway identifies then incoming traffic as new VPN traffic and initiates a request to the RADIUS server for user authentication requirements.

- The RADIUS server responds to the security gateway indicating authentication is required.

- The security gateway challenges the user to provide the required authentication information.

- The user enters the required authentication information via a prompt displayed by VPNremote. This challenge response is sent back to the security gateway.

- The security gateway forwards the challenge response to the RADIUS server.

- The RADIUS server decides if the user has met the challenge, and if so, informs the security gateway that the user is authorized. The RADIUS server also forwards the user configuration details, known as user attributes, to the security gateway. These attributes specify VPN-specific information, including the cryptographic keys used for encryption.

- The security gateway then allows VPN traffic to flow between the VPNremote Client and the VPN members.

Two methods of user authentication—simple passwords and "one-time" passwords based on two-factor authentication mechanisms—can be used to meet a variety of security, cost, and convenience requirements. All RADIUS implementations support standard password authentication, and many can be used in conjunction with RSA Security ACE/Server for SecurID™ Token requirements.

The RADIUS protocol

The RADIUS protocol is documented in an Internet Engineering Task Force (IETF) Request for Comment (RFC), specifically RFC 2058.

- **Client/Server Model** – A Network Access Server (NAS) operates as a client of RADIUS. The client is responsible for passing user information to designated RADIUS servers and then acting on the response that is returned. RADIUS servers are responsible for receiving user connection requests, authenticating the user, and then returning all configuration information necessary for the client to deliver service to the user. A RADIUS server can act as a proxy client to other RADIUS servers or other kinds of authentication servers.

- **Network Security** – Transactions between the client and RADIUS server are authenticated through the use of a shared secret, which is never sent over the network. Additionally, user passwords are sent encrypted between the client and RADIUS server to eliminate the possibility that someone snooping on an unsecure network could determine a user's password.

- **Flexible Authentication Mechanisms** – The RADIUS server can support a variety of methods to authenticate a user; when given the user name and the original user password, it can support PPP, PAP or CHAP, UNIX login, and other authentication mechanisms, some of which include the use of cryptographically strong tokens. These tokens use a two-factor approach to authentication: the first is a Personal Identification Number (PIN); the second is a value taken from the token. An example of a two-factor authentication mechanism is the SecurID™ token card and ACE/Server AccessManager by RSA Security.

Some RADIUS server implementations use several files to manage the database of information needed to provide Client authentication. A number of these files must be modified to use the VSUs as an NAS within a RADIUS environment.

Add (RADIUS/ACE server)

Authenticating (secret) password

Enter the authenticating password followed by a retype.

RADIUS server data

IP Address Enter the IP address of the RADIUS/ACE server.

UDP Port Enter the UDP port of the server. The default value is 1645. Check your RADIUS server documentation to verify the value for this field.

Use this as my: Select the role you wish this server to perform: Primary Server, Secondary Server, or Tertiary Server.

To add a RADIUS server:

1. From the *Contents* column, select the security gateway you want to configure.

2. Click the **Policies** tab to bring it to the front.

3. From the *drop-down list*, select **RADIUS/ACE**, then click **GO** to open the *Policy Manager* for *RADIUS/ACE*.

4. Select the **Enable RADIUS/ACE** check box so the security gateway uses RADIUS services.

5. Click **Add** to open the *Add RADIUS/ACE* dialog box.

6. In the **Password** text box, type in the *shared secret* that the security gateway uses to authenticate itself to the RADIUS server.

   Note: This value is also entered later in the RADIUS server Client file. Check your RADIUS server documentation for valid password length and allowed characters.

7. In the **Confirm Password** text box, type in the shared secret to confirm it.

8. In the **IP Address** text boxes, type in the address of the RADIUS server.

    Note*: An IP address must be entered (domain names are not valid). There must be an IP route between the security gateway and the target RADIUS server.*

    To verify that a valid IP route exists, use the security gateway proxy ping function (security gateway tab/Connectivity) and enter the target RADIUS server's IP address as the ping target.

9. In the **UDP Port** text box, type the port number for the server.

    The default number is usually 1645, but use the RADIUS server's documentation to confirm the number.

10. From the **Use this as my** options, assign a query order to the server. If backup servers are being used, here is where they can be identified.

    • Select **Primary Server** if no backup servers are used, or if this is the server primarily used if backup servers are running.

    • Select **Secondary Server** if this server operates as a backup to the primary server.

    • Select **Tertiary Server** if this server operates as a backup to the secondary server.

11. Click **OK** to return to the *Policy Manager* window.

12. From the list of servers, select the new server.

13. From the *Settings* options, use the following to configure the connection expiration times for the server.

    • **RADIUS Attempts**. The number of times a RADIUS server is contacted before failure is assumed and the next RADIUS server is used.The default is 3 attempts.

    • **Time to assume failure**. The time that should pass when a RADIUS server is not responding and the next RADIUS server is used.The default value is 6 seconds.

- **Designated RADIUS attribute for policy**. Designates the VPN Policy to the security gateway that is delivered to the remote client when the remote client authenticates throughout the security gateway to the RADIUS Server.

  The VPNmanager provides the following attributes for the remote client to choose from:

  - Filter ID

  - Replay Message

  - Class (default set by Administrator)

  - Vendor Specified

  - User Defined

- **User-defined RADIUS attribute ID**. ID text field is enabled and the user provides the attribute ID. If the user does not provide the ID, an error message is displayed. This field can be used with less common attribute IDs.

- **Use this tag for RADIUS attribute**. The tag must contain the letters a to z or A to Z. The tag can be up to 15 characters in length.

14. Click **Close** to return to the *Configuration Console* window.

15. Click **Save**.

16. When you want to send the configuration to the security gateway, click **Update Devices**.

Chapter 6 Configuring user groups

The User Group function is used to setup and maintain logical groups in which the individual VPN users reside.

User groups have a single-level hierarchy - you cannot have a user group within another user group.

A *User Group Object* is a method for simultaneously managing many user objects (remote users). For example, all remote users, who are in sales, can be consolidated into a single user group. Then that group can be associated with one or more VPN objects. Without user groups, remote users would have to be individually associated with a VPN object.

User groups are easy to create and configure. You give them a name then populate them with user objects.

Users can belong to more than one user group. When this is the case and policy conflicts exist, permit wins over deny (user group settings override individual user settings).

User groups can be created at anytime. But since they are configured with user objects, you should configure users before configuring user groups.

New user group

To create a user group:

1. From the VPNmanager console main page, Click **New Object** and select **User Group**. The New User Group dialog is displayed.

2. In the **Name** text box, type in a name for the new group. Any characters can be used, except a comma [,].

3. If you want to create more groups, press **ENTER**, then type in another name.

4. Click **Apply**, then **Close** to return to the *Configuration Console* window.

5. Click **Save**.

You now configure your new user group.

Note: Renaming user groups is not currently supported.

User Group - General tab

The User Group General tab is used to manage your users and their respective user group assignments.

Figure 42. User Group, General tab
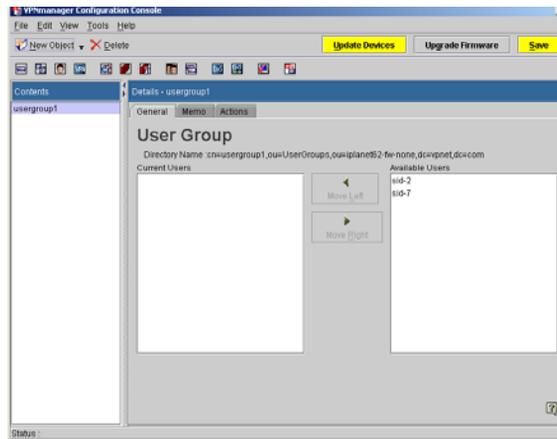
All existing user groups are displayed in the Contents list. The highlighted user group is displayed in the General tab window.

Directory Name. This is the unique User Group name. It is unique in that it is not duplicated anywhere within the VPN domain to which it is assigned.

Current Users. This area contains the names of all individual Users currently assigned to this User Group. A second pane, titled Available Users, lists all existing VPN users. The left and right arrows are used to move the highlighted users from one column to the other.

Available Users. This pane is a list of all available users. The highlighted user may be moved into the Current User Members list by using the left arrow. Only one default user can be added to a User Group.

User Group - Memo tab

Memo can be used to record notes about the User Group, such as change history, function of this group (such as all administrators, etc.). Information entered here is associated only with this User Group. This information is stored only in the database and not downloaded to the security gateways.

User Group - Actions tab

The Actions tab is used to control authentication for specific user groups.

Figure 43. User Group, Actions Tab
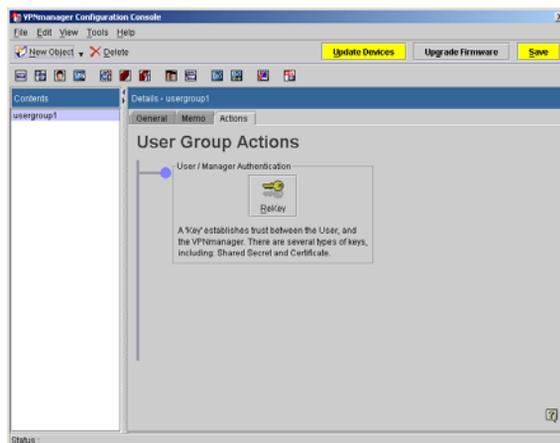
User/Manager authentication Rekey is used to change the key of the highlighted user group. You should change the key regularly to ensure maximum security.

Only SKIP and Preshared Secret IKE VPNs can be manually rekeyed. In the case of SKIP, rekeying generates and distributes a new master key to all security gateways associated with the VPN. This SKIP master key is used to generate session keys used for cryptographic functions. In the case of Preshared Secret IKE VPNs, rekeying generates and distributes a new negotiation key to all security gateways associated with the VPN. This negotiation key is used to provide authentication during IKE negotiations, in which the actual session key is dynamically generated. Manual Keyed VPNs can be rekeyed by manually editing the relevant keys.

Configuring a user group

To configure a user group:

1. Move to the *Configuration Console* window.

2. From the Icon toolbar, click **User Group** to list all the user groups in the *Contents* column.

3. From the Contents column, select the user group that needs to be configured.

4. Use the **General** tab to populate the group with specific users.

   • From the **Available Users** column, select one or more users. To select multiple users which are listed adjacently, hold the SHIFT key. To select multiple users which are not adjacently listed, hold the CTRL key.

   • Click **Move Left** to move your selected users to the **Current Users** column.

5. (Optional) Click the **Memo** tab to bring it to the front, then type in a message about the group, such as its purpose, or who it serves.

6. Click **Save**.

Chapter 7   Configuring VPN objects

A VPN object is the method used for linking security gateways, remote

terminals, and LAN terminals in a fully configured virtual private network. To create a VPN, you name the VPN, select a key management method, and optionally, designate it as the Default VPN. After that you can configure the VPN using VPNmanager, using the tabs associated with the created VPN. When you configure the VPN, you add users and user groups and further define the IKE, IPSec, and SKIP security protocols for VPN traffic.

Types of VPN objects

Two types of VPN objects can be built.

- *SKIP* based VPN

- *IKE* based VPN

Both types use *IP Security Protocol (IPSec)* for encrypting and decrypting VPN traffic. The main difference between the two VPN types are the methods used for creating the encryption key. When you create a VPN object, you select which protocol to use.

SKIP VPNs

Note: SKIP VPNs are supported in VPNremote Client 2.5 only.

Simple Key-management for IP (SKIP) is a protocol that stores authentication and security information in every packet. SKIP VPNs can operate in *Tunnel* or *Transport* modes. Tunnel mode involves encrypting the entire original IP packet before it goes out to the public networks.

Transport mode involves encrypting only the payload of the original packet. Also, SKIP VPNs can be manually rekeyed at any time.

IKE VPNs

Note: IKE VPNs are supported in VPNremote Client 3.0 and later.

An IKE VPN can run in *certificate* or *preshared secret* authentication mode. Also, IKE VPNs always operate in tunnel mode, which means the entire original packet (header and payload) is encrypted and inserted in the payload of an IPSec packet before it goes out to the public networks.

Certificate mode involves the exchange of X.509 *public-key certificates* between endpoints of a VPN tunnel to authenticate VPN tunnel end points. A certificate belonging to a specific endpoint is authenticated by a third party certificate called an issuer's certificate. Certificates can be obtained from a third party Public Key Infrastructure (PKI). See for more information about using a PKI. Certificate based VPNs cannot be manually rekeyed.

Preshared Secret mode involves the Diffie-Hellman algorithm for creating a *shared secret key* that is used for authenticating VPN traffic. Large prime numbers and modular arithmetic equations are exchanged between endpoints. Each endpoint uses the equations and numbers to calculate the same shared secret key. The tunnel endpoints then use the shared secret key to authenticate each other's traffic. Even if the prime numbers and equations become publicly known, the protocol still protects the shared secret key. As an added security measure, preshared secret can be manually rekeyed at any time.

VPN packet processing modes

There are two ways to process packets when forming VPNs: transport mode and tunnel mode.

In transport mode, IP packets sent between VPN members are secured by applying VPN services to the IP packet payload, leaving the original addressing header unchanged.

Original IP Packet

Source Address	Dest. Address	Payload

Source Address	Dest. Address	IPSec/SKIP Overhead	Payload with Applied VPN Services

Transport Mode
Secured VPN IP Packet
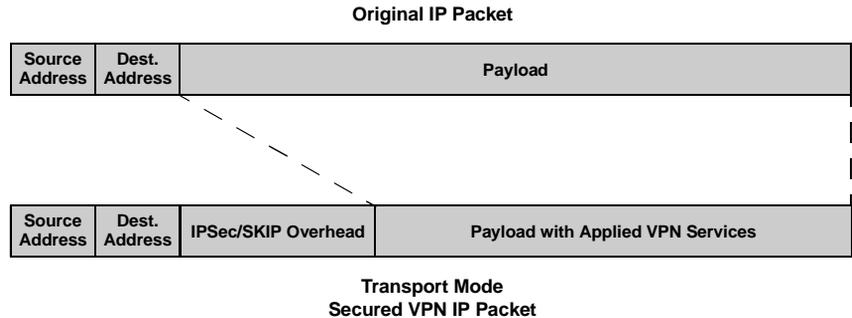
In tunnel mode (security gateways and VPNremote Client only), IP packets between members are secured by encrypting and authenticating the entire packet, including the addressing header. The encrypted and authenticated packet is then used as the payload of a new packet with a new addressing header. This new addressing header specifies the IP addresses of packet's source and destination, whether they be two security gateways or a VPNremote Client and a security gateway.

Original IP Packet

Source Address	Dest. Address	Payload

Src VSU or Client Address	Dest VSU or Client Address	IPSec/SKIP Overhead	IP Packet with Applied VPN Services

Tunnel Mode
Secured VPN IP Packet
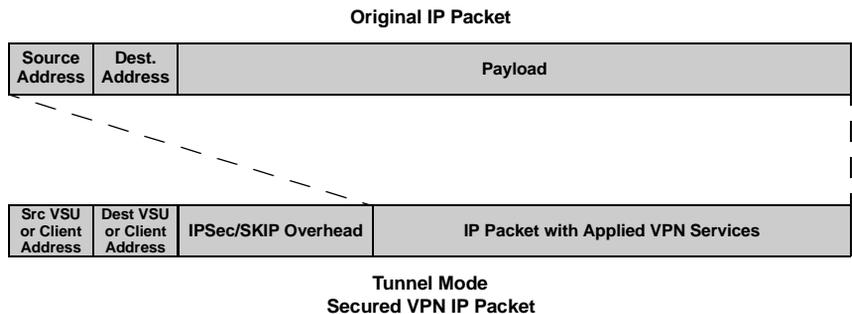
The choice between using transport and tunnel mode involves many factors, including the use of private IP addresses for Groups and security concerns about the visibility of member workstation IP addresses.

The following key management and packet mode combinations are supported:

- SKIP in Transport or Tunnel mode.

- IKE in Tunnel mode only.

Default VPN policy

Default VPN applies only to the IKE VPN and is used in conjunction with RADIUS authentication. Only one VPN can be the default VPN in a domain. When you create a VPN, you can enable this function.

Default Policy is an alternative method of external user authentication. This feature is suited for large IKE-based VPNs where hundreds or even thousands of users are authenticated, or where the ability to scale the VPN to large numbers of authenticated users is required. This default VPN policy is applied to any remote user authenticated successfully by the external RADIUS server.

When a remote user requests CCD from the security gateway, the security gateway's RADIUS client contacts the RADIUS server to authenticate the user. Upon successful authentication, the CCD serer provides the default VPN policy to the user.

Creating a new VPN object

To create a new VPN object:

1. From the VPNmanager Console main window, click **New Object** and select **VPN**. The New VPN dialog is displayed.

2. In the **Name** text box, type in a name for your new VPN Object. Any characters can be used, except a comma [,].

3. From the *VPN Type* options, do one of the following.

   • Select **SKIP** to create a SKIP VPN Object.

   • Select **IKE** to create an IKE VPN Object.

4. Click **Apply** to create the object.

5. If you want to create another object, repeat step 2 and step 3.

6. Click **Close.** The *Configuration Console* appears and the details pane displays a series of tabs.

7. Click **Save** to save your work.

Creating a default VPN

To create a default VPN within a selected domain:

1. Add the security gateway(s). Add an IPGroup(s) and associate this group with this security gateway.

2. Create a default user or default user group in the VPNmanager.

3. Create new VPN Object, see "Creating a new VPN object" on page 156 and check Default VPN checkbox.

4. Add default user and IPGroup(s) to the new VPN.

5. Configure the RADIUS Server using the Policy Manager.

6. Enable the RADIUS Authentication/Local Configuration or LDAP/Local Configuration from the Preferences screen on the VPNmanager. You can go to the preferences screen by clicking the Edit/Preferences menu on the First screen of VPNmanager Console.

7. Update this configuration to the security gateway(s). The security gateway(s) should now have a default VPN in its configuration.

8. On the RADIUS server, add a user. Enter the user credentials.

9. On the LDAP server, a local server or an external server with a different context, add user. Enter the user credentials.

10. Log in to the security gateway through the VPNremote client using the credentials entered in the RADIUS/LDAP server. The user should be authenticated successfully by the RADIUS/LDAP server. The RADIUS/LDAP server returns the VPN name to the security gateway. The user then gets the default VPN policy from the security gateway.

Creating a designated VPN

RADIUS attributes enable the VPN administrator to define what VPN policy is delivered to the remote client by the security gateway during the authentication process.

To set up a designated VPN within a selected domain, perform the following steps:

1. Add the security gateway(s). Add an IPGroup(s) and associate this group with this security gateway.

2. Create a default user or default user group in the VPNmanager.

3. Create a new VPN Object, see Creating a new VPN object.

4. Add the default user and IPGroup(s) to the new VPN.

5. Use the Policy Manager to configure the RADIUS Server, Attributes, and Settings.

   The RADIUS attributes and setting can remain as default.

6. Click **Edit>Preferences** on the main screen of VPNmanager Console and enable **RADIUS Authentication/Local Configuration**.

7. Update this configuration to the security gateway(s). The security gateway(s) should now have the designated VPN in its configuration.

8. On the RADIUS server, add a user. Enter the user credentials and the attribute type & tag to match to the one you entered in the VPNmanager for that security gateway.

9. Now login into the security gateway through the VPNremote client using the credentials entered in the RADIUS server. The user should be authenticated successfully by the RADIUS server. The RADIUS server returns the VPN name to the security gateway. The user then gets the designated VPN policy from the security gateway.

Using the VPN tabs

After you have created a VPN object, you can use the VPN tabs to change the default settings or modify configuration. The tabs displayed are dependent on the VPNos release for the device.

General tab

The General tab provides high-level control of the VPN. A check box enables the VPN. This allows VPNs to be built before being activated.

The contents of this screen depends on what VPN type you have selected, IKE or SKIP.

General tab with IKE

If the VPN type selected is IKE, the following General tab appears:

Figure 44. VPN General Tab (IKE)
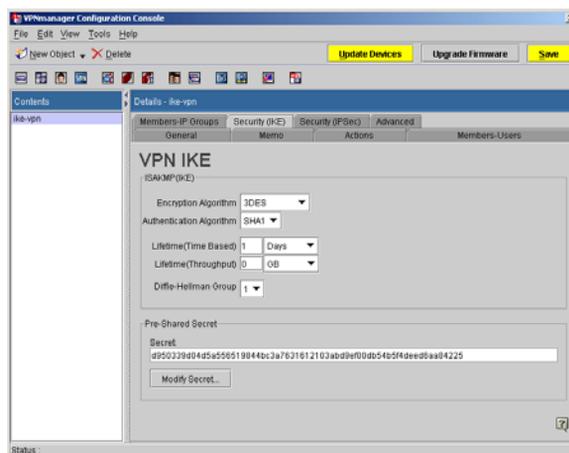
From the General tab you can configure the following information:

Certificate Based. Certificate based authentication is the most secure key management method used to construct a VPN, but requires greater setup effort than with the Preshared Secret method.

Preshared Secret. Preshared Secret authentication is the simplest key management method used to construct a VPN. Authentication key exchanges between security gateways in the VPN are based on a single pre-shared secret known to all security gateways in the VPN.

Enable VPN. When this box is checked and the security gateway has been updated, the VPN is active. Unchecking the box disables the VPN and is typically used during the troubleshooting process.

Default VPN. When this box is checked, this VPN is the default VPN for the domain. Only one VPN can be the default VPN in a domain. Default

VPN is an alternative method of user authentication suited for large IKE-based VPNs.

Directory Name. In the VPN information area the unique VPN name is displayed along with the directory server context.This area also shows the security key exchange protocol that the VPN uses globally.

General tab with SKIP

If the VPN type you selected is SKIP, the following General tab appears.

When SKIP is selected, from the General tab you can configure the following information:

Tunnel. Select the tunnel mode if IP packets between members are secured by encrypting and authenticating the entire packet including the addressing header.

Transport. Select the transport mode if VPN services are applied to the IP packet payload sent between VPN member. The original addressing header is unchanged.

Enable VPN. When this box is checked and the security gateway has been updated, the VPN is active. Unchecking the box disables the VPN and is typically used during the troubleshooting process.

Directory Name. In the VPN information area the unique VPN name is displayed along with the directory server context.This area also shows the security key exchange protocol that the VPN uses globally.

Memo tab

The Memo tab can be used to record notes about the VPN, such as change history, VPN type, etc. Information entered here is associated only with this VPN and is stored in the database.

Members-Users tab

The Members-Users tab is used to establish the user membership of the VPN. A list of currently assigned users appears in the Current VPN Members list. All available users appear in the Available list. Use the right and left arrows to move the users to the desired column.

Figure 45. VPN, Members [Users] tab
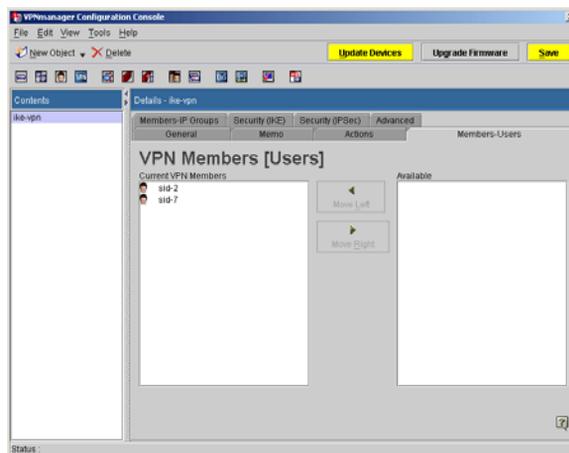
Note: *When a remote user is removed from a VPN and the security gateway is updated, all non-RADIUS enabled security gateways that are affected by the removal of the remote user are updated. For RADIUS enabled security gateways, the remote user is not removed from the VPN until the configuration record is removed from the RADIUS database.*

Members-IP Groups tab

The Members-IP Groups tab is used to establish the IP Group membership for this VPN. A list of currently assigned members appears in the Members list while all available IP Groups appear in the Available list. Use the right and left arrows to move the IP Groups to the desired column.

Figure 46. VPN, Members [IP Groups] Tab
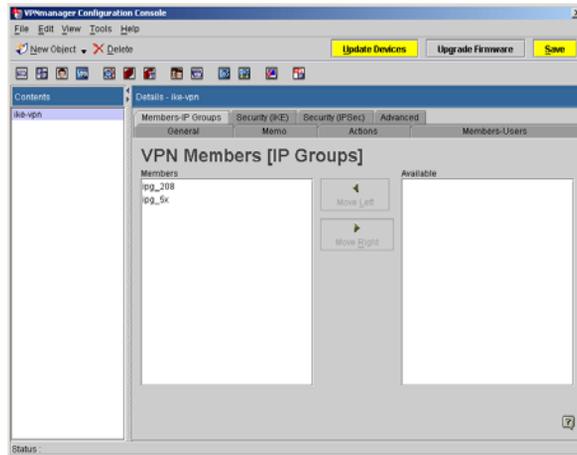
Security (IKE) tab

The *Security (IKE) tab* is used for configuring the encryption and authentication algorithms used at the end-points of a VPN tunnel. The configuration procedure involves setting a lifetime for public-keys, and a specific Diffie-Hellman Group for automatically generating keys of a specific strength. For additional protection, unique (new) keys are automatically generated and exchanged between all security gateways and VPNremote Clients in the VPN Object based on their lifetime.

In the ISAKMP (IKE) area you set up the key-exchange parameters that you want used for the VPN.

Field	Description
Encryption Algorithm	Select one of the following types: • DES. A common encryption algorithm that is not subject to export regulations. • 3DES. A robust encryption algorithm. 3DES is subject to government regulation. Contact Avaya for a current list of controlled and uncontrolled application and territories. • Any. Accepts any encryption proposal that is made by the device on the other side. IKE VPNs use ESP to encrypt IP packets as defined in RFC2406. You can choose either DES-CBC or 3DES-CBC (Domestic U.S./Canada only) encryption. *Note: The use of 3DES is subject to government regulation. Contact Avaya, VPN Support, for a current list of controlled and uncontrolled applications and territories.*
Authentication Algorithm	Select one of the following types: • MD5 (RFC1321) • SHA1 • Any. Accepts any authentication proposal that is made by the device on the other side. IKE VPNs use either an ESP trailer as defined in RFC2406, or AH as defined in RFC2402 to authenticate IP packets.

Field	Description
Lifetime	Payload key lifetime defines the extent to which a single set of cryptographic keys is used when applying VPN services to IP packets. Lifetimes are either time based or based on throughput. **Time-based** lifetimes are based on the amount of time that the keys are used without a key change. **Throughput** lifetimes are defined by the amount of data that is acted on by a set of keys. The more often a key is changed, the "more secure" the system. However, frequent key changes can affect system performance.
	Enter a numerical value and select a unit of measure for both time-based and throughput lifetimes. Whichever occurs first triggers the new key.
	Note: For time-based lifetime, the following are the minimum values in each category: Day = 1, Minutes = 1, and Seconds = 60.
Diffie-Hellman Group	Diffie-Hellman Groups define the mathematical parameters that are used during IKE negotiations. Group 1 specifies the use of a 768-bit modulus. Group 2 specifies the use of a 1024-bit modulus , and thus Group 2 is "more secure". See RFC2409 for more information on Diffie-Hellman Groups.

Pre-Shared Secret

The Pre-Shared Secret area appears only when the VPN type is IKE with Preshared Secret selected. The preshared secret appears in the Secret field as either ASCII or hexadecimal.

Select **Modify Secret** to change the preshared secret. Both the local and the remote security gateway must have the identical preshared secrete text, or a secure tunnel cannot be established between them.

Enter the secret character string, up to 64 hexadecimal characters or 16 ASCII characters.

Use **Autogenerate** to generate a random character sequence.

Select either ASCII or hexidecimal to display the secret.

Security (IPSec)

In IKE VPNs, VPN traffic flows in tunnel mode. Therefore, the *Security (IPSec) tab* is used for configuring the parameters used for encapsulating the original packet (header and payload) into the payload of an IPSec packet.

Packet-level security involves establishing an agreement between security gateways about which IPSec protocol configurations to use. The Security (IPSec) tab has two sets of options. The *IPSec options* control packet alteration, and the *IPSec Proposal options* are used for creating up to four different proposals for payload encryption and authentication.

Security gateways must use the same IPSec Proposal. An *IPSec Proposal* dialog box is used for creating different proposals in cases where the proposal is unknown.

Figure 47. VPN, Security (IPSec) Tab
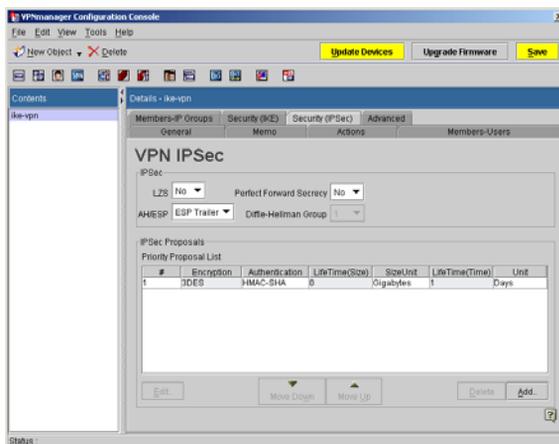
In the IPSec area you set up the IPSec protocol information that you want the VPN to use

LZS. This refers to Lempel-Ziv-Stac hardware date compression technique used prior to encryption. Yes/No enables or disables its use.

AH/ESP. This is the Authentication Header (AH)/Encapsulation Security Payload (ESP). IKE VPNs authenticate IP packets using either an ESP trailer as defined in RFC2406, IP Protocol 51, or AH as defined in RFC2402, IP Protocol 52.

Perfect Forward Secrecy. Perfect Forward Secrecy defines a parameter of IKE that discloses long-term secret keying material that does not compromise the secrecy of the exchanged keys from previous communications. Enabling Perfect Forward Secrecy is more secure, but involves more overhead. It is recommended that your VPN use this option if your VPN encryption algorithm is DES. See RFC2409 for additional information on Perfect Forward Secrecy.

When enabled (Yes), a Diffie-Hellman Group number must be selected.

Diffie-Hellman Group. Diffie-Hellman Group defines mathematical parameters used during IKE negotiations. Group 1 specifies use of a 768-bit modulus, Group 2 specifies use of a 1024-bit modulus (Group 2 is more secure). See RFC2409 for additional information on Diffie-Hellman Groups.

IPSec Proposals

The IPSec proposals area displays a list of all currently defined proposals ranked by priority of negotiation. You can add, edit or delete new IPSec proposals and you can relocate them in the list. A maximum of four IPSEC proposals are allowed in the IPSEC Proposal Priority Proposal list.

An extranet is an example of when several proposals are desirable. By having several choices, the odds of a finding a mutually common proposal on both sides is increased. Another example is where international security gateways (DES only) and a domestic security gateways (DES or 3DES) are part of the same VPN. Having a DES proposal establishes a common ground for the two security gateways to communicate.

Add IPSec proposal

You can add up to four IPSec proposals. You determine the encryption method, the authentication methods, how long a single set single set of cryptographic keys is used when applying VPN services to IP packets and the order this proposal is in the list.

Field	Description
Encryption	Select one of the following types: • DES. A common encryption algorithm not subject to export regulation. • 3DES. A robust encryption algorithm. • AES-128. The advanced encryption standard that uses a 128-bit block to help resist large attacks. • Any. Accepts any encryption proposal made by the device on the other side.
Authentication	Select one of the following types: • Any. Accepts any authentication proposal that is made by the device on the other side. • None • HMAC-MD5 • HMAC-SHA

Field	Description
Lifetime	Payload key lifetime defines the extent to which a single set of cryptographic keys is used when applying VPN services to IP packets. Lifetimes are either time based or based on throughput. **Time-based** lifetimes are based on the amount of time that the keys are used without a key change. **Throughput** lifetimes are defined by the amount of data that is acted on by a set of keys.
	Enter a numerical value and select a unit of measure for both time-based and throughput lifetimes. Whichever occurs first triggers the new key.
	Note: For time-based lifetime, the following are the minimum values in each category: Day = 1, Minutes = 1, and Seconds = 60.
Locate This IPSec Proposal	Establishes the IPSec proposal rank in the negotiating list. The first proposal in the list is the first attempted to be negotiated with the device on the other side. • Beginning of list • End of list • After Selected Item

Actions tab

The Actions tab is used to export the VPN (without keys) and to change the VPN security key (Rekey).

Figure 48. VPN, Actions tab
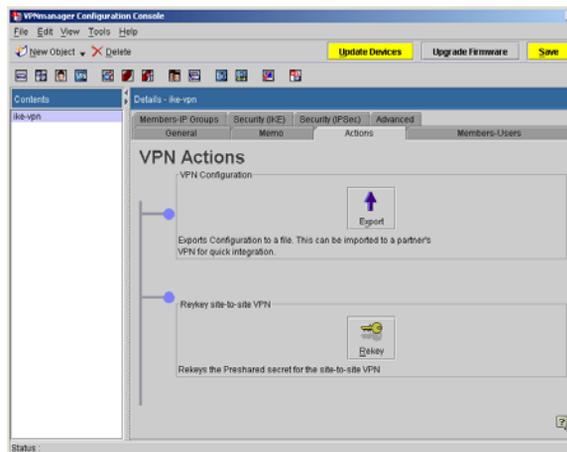
VPN configuration

Export

Exports the VPN to another VPN domain without the keys. Typically used to create an extranet.

Creating an extranet of Avaya VPN devices is a cooperative effort between system administrators running independent copies of VPNmanager and involves the same steps as creating any other VPN: create the device, then the groups and users, and finally the VPN.

The names chosen for VPN components must be unique and synchronized within each corporation's VPNmanager. This requires close coordination between the system administrators during the VPN component creation process.

Once OK is clicked, a password screen appears. After the password is entered and OK is clicked, the Save screen that is displayed shows the destination location and filename.

Rekey site-to-site VPN

Rekey

Used to change the preshared secret key of a site-to-site VPN. This should be done regularly to ensure maximum security.

Only SKIP and Preshared Secret IKE VPNs can be manually rekeyed. In the case of SKIP, rekeying generates and distributes a new master key to all security gateways associated with the VPN. This SKIP master key is used to generate session keys used for cryptographic functions. In the case of Preshared Secret IKE VPNs, rekeying generates and distributes a new negotiation key to all security gateways associated with the VPN. This negotiation key is used to provide authentication during IKE negotiations, in which the actual session key is dynamically generated. Manual Keyed VPNs can be rekeyed by manually editing the relevant keys.

Advanced VPN tab

The Advanced tab is used to set up advanced VPN options. Generally, the defaults do not need to be changed.

Figure 49. VPN Advanced tab
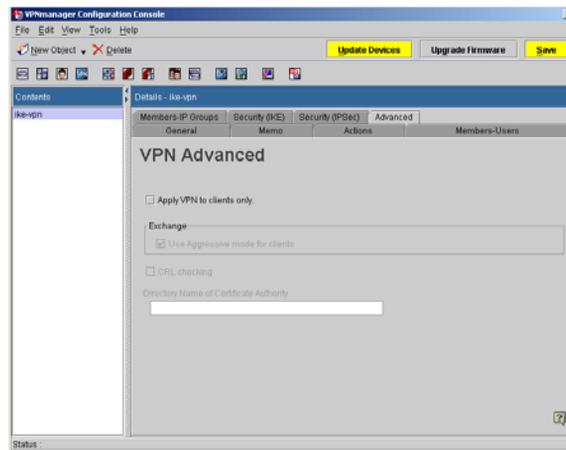
Apply VPN to clients only provides VPN access to users and ignores the site-to-site "mesh" or relationships between security gateways. This is a usability feature that can be used in VPNs with complex rules to only mesh the users.

In a normal VPN, the IP Groups are meshed together and the users are meshed with the groups. When the "Apply VPN to clients only" check box is check, only the users are meshed.

In the **Exchange** area, check **Use Aggressive mode for clients** to enable the IKE Aggressive mode between a user and then security gateway, which accomplishes the same goals as Main mode, only faster.

Note: Aggressive mode must be used when Preshared Secret is being used for the remote client users. When certificate-based key exchange is used, either Main mode, or Aggressive mode may be used.

CRL checking enables *certificate revocation list* checking, which looks to a directory server to obtain a CRL to validate a newly arrived certificate.

In the **Directory Name of Certificate Authority** box, enter the DNS name of the CA server.

Configuring a SKIP VPN

Note: security gateways at each end of a tunnel must use the same SKIP settings.

To configure a new SKIP VPN object:

1. Move to the *Configuration Console* window.

2. From the Icon toolbar, click **VPN** to list all VPN Objects in the *Contents* column.

3. From the Contents column, select the VPN Object that needs to be configured.

4. Click the **General** tab to bring it to the front.

5. Select one of the following to control how VPN traffic must be protected.

   • Select the **Tunnel** radio button so entire IP packets (header and payload) are encrypted and put it into the payload of a VPN packet.

- Select the **Transport** radio button so only the payload of IP packets is encrypted, and the entire IP packet is put into the payload of a VPN packet.

   Note: If you plan on defining the VPN Object with IP Group Objects, Transport mode must be used.

6. (Optional) Click the **Memo** tab to bring it to the front, then type in a note about this specific VPN Object.

7. If you want to add *User Objects* or *User Group Objects* as members of this VPN Object, do the following.

   - Click the **Members-Users** tab to bring it to the front.

   - From the **Available** list, select specific User Objects and User Group Objects. User Group Objects are always located at the bottom of the list.

   *Tip: Hold the **Shift** key to simultaneously select many adjacent items, or hold the **Crtl** key to simultaneously select many non-adjacent items.*

   - Click **Move Left** to move the selected items to the **Current Members** list.

8. If you want to add *IP Group Objects* as members of this VPN Object, do the following.

   - Click the **Members-IP Groups** tab to bring it to the front.

   - From the **Available** list, select specific IP Group Objects.

   - Click **Move Left** to move the selected items to the **Current Members** list.

9. Click the **Security (SKIP)** tab to bring it to the front.

10. From the **Encryption Algorithm** list, do one of the following.

    - Select **Triple DES** to divide VPN traffic into 64 bit blocks and encrypt each block three times with three different keys.

    - Select **DES** to divide VPN traffic into 64 bit blocks and encrypt each block with a 56-bit key.

    - Select **NONE** to not encrypt VPN traffic.

11. From the **Authentication Algorithm** drop-down list, do one of the following.

    - Select **Keyed MD5** if you want VPN tunnel end-points to authenticate themselves using the Message Digest 5 hash function.

      Tunnel end-points are security gateways and VPNremote Clients.

    - Select **NONE** if you do not want tunnel end-point to authenticate themselves.

12. From the **Compression Algorithm** list, do one of the following.

    - Select **Stac** if you want the payloads of VPN packets to be compressed using the STAC Lempel-Zif standard compression. Since encryption is time-consuming, compression speeds up the entire process.

    - Select **NONE** you do not want payloads of VPN packets to be compressed.

13. Click **Save** to save your work.

Configuring an IKE VPN

Note: security gateways at each end of a tunnel must use the same IKE settings.

To configure a new IKE VPN Object:

1. Move to the *Configuration Console* window.

2. From the Icon toolbar, click **VPN** to list all VPN Objects in the *Contents* column.

3. From the **Contents** column, select the VPN Object that needs to be configured.

4. Click the **General** tab to bring it to the front.

5. Select one of the following to control how tunnel end-points must authenticate themselves. End-points are defined as security gateways and VPNremote Clients.

    - Select **Certificate Based** to use X.509 public-key certificates.

    - Select **Preshared Secret** to use shared secret keys.

6. (Optional) Click the **Memo** tab to bring it to the front, then type in a note about this specific VPN Object.

7. To add *User Objects* or *User Group Objects* as members of this VPN Object, do the following.

    • Click the **Members-Users** tab to bring it to the front.

    • From the **Available** list, select specific User Objects and User Group Objects. User Group Objects are always located at the bottom of the list.

    • Click **Move Left** to move the selected items to the **Current Members** list.

8. To add *IP Group Objects* as members of this VPN Object, do the following.

    • Click the **Members-IP Groups** tab to bring it to the front.

    • From the **Available** list, select specific IP Group Objects.

    • Click **Move Left** to move the selected items to the **Current Members** list.

9. Click the **Security (IKE)** tab to bring it to the front.

    Configuring the encryption and authentication algorithms used at the end-points of a VPN tunnel.

    Use the **Encryption Algorithm** list to select a specific type of encryption algorithm that each security gateway and VPNremote Client must use for this VPN Object.

    • Select **Any** if you want the security gateways to automatically negotiate which algorithm to use.

    • Select **DES** to divide VPN traffic into 64 bit blocks and encrypt each block with a 56-bit key.

    • Select **3DES** to divide VPN traffic into 64 bit blocks and encrypt each block three times with three different keys.

10. Use the **Authentication Algorithm** list to select a specific type of algorithm that each security gateway must use to authenticate each other.

   • Select **Any** if you want the security gateways to automatically negotiate which algorithm to use.

   • Select **MD5** if you want each security gateway to authenticate each other using the *Message Digest 5 (MD5)* hash function.

   • Select **SHA1** if you want each security gateway to authenticate each other using the *Secure Hash Algorithm-1 (SHA-1)*.

      SHA1 is considered to be a stronger hash function than MD5, and may be required for US Federal applications that do not require a digital signature.

11. From the **Lifetime** text boxes and lists to configure the time limit for creating and exchanging a new set of unique keys.

   If the *Time-based* value expires before the *Throughput* value, key creation and exchange is performed, and likewise, if *Throughput* expires before the *Time-based* value.

12. Click **Modify Secret** to open the *Modify Secret* dialog. Create a shared secret for authenticating security gateways and members of the VPN.

   • To manually create a secret, type in an alphanumeric string in the text box

   • To automatically create a secret, click **Auto-generate**.

13. Click **OK.**

   Note: Modify Secret is only available when creating a VPN based on Preshared Secret.

14. Click the **Security (IPSec)** tab to bring it to the front.

   The Security (IPSec) tab is used to set up the desired IPSec protocol information (parameters relating to payload) that the VPNs use.Two sets of options are available. The *IPSec options* control packet alteration, and the *IPSec Proposal options* are used to create up to four different proposals for payload encryption and authentication.

15. Use the **LZS** list for applying compression to packet payloads.

According to RFC 2395, *"IP Payload Compression using LZS,"* experiments have shown that the LZS algorithm compressed a 64-byte file to 85% of its original size, while a 16384-byte file was compressed to 47% of its original size. Whether or not your network benefits from compression, depends on what is typically transported; for example, video and sound traffic are already compressed, so additional compression has little effect and may load the security gateway.

- Select **Yes** to apply compression.
- Select **No** to not apply compression.

16. Use the **Perfect Forward Secrecy** list to control key creation.

Perfect Forward Secrecy (PFS) is a key-creation method used for assuring that a new key is not related to any previous keys. This is done by using key creation values which are independent of past values.

- Select **Yes** to use PFS.
- Select **No** to not use PFS.

17. Use the **AH/ESP** list to create packets containing IPSec headers. The payloads contain the entire original packet (header and payload).

- Select **AH Header** to authenticate the entire packet.

  This inserts an *Authentication Header* and *Encapsulating Security Payload (ESP) Header* into packets and perform encryption on the payload.

- Select **ESP Trailer** to authenticate the entire packet, except for the IP header.

  This will insert an *ESP Header* and *ESP Trailer* into packets and perform encryption on the payload.

18. Use the **Diffie-Hellman Group** list to select which *modulus* to use for the keying algorithm.

- Select **1** to use a 768-bit modulus.
- Select **2** to use a 1024-bit modulus.

For detailed information about Group 1 and Group 2 algorithms, see section 6.2 of IETF RFC 2395.

19. Use the **IPSec Proposals** options to create one or more proposals.

A proposal defines which IPSec parameters all the security gateways of a VPN must use. If all the security gateways are of the same type, only one proposal needs to be created.

If an extranet (a VPN belonging to another organization) is going to connect to your VPN, and its proposal is different, or unknown, additional proposals can be added to the *Proposal List* to accommodate that unique security gateway. The security gateways will automatically go through the list and negotiate on which proposal to use at the appropriate time.

- Click **Add** to open the *Add IPSec Proposal* dialog box.

- From the **Encryption** drop-down list, select the type of encryption to be app**lied to packet payloads.**

  - **Null.** Payload is not encrypted, but AH/ESP headers are included. Used by engineers for packet analysis.

  - **DES Single.** DES encryption is applied to the payload.

  - **3DES Triple.** DES encryption is applied to the payload.

  - **AES-128.** AES-128 advanced encryption is applied to the payload.

  - **RC5.** Applies RC5 encryption.

  - **Any.** Let the security gateways negotiate which encryption method to use.

- From the **Authentication** drop-down list, select the type of authentication to use.

  - **None**. Packets are not authenticated.

  - **HMAC-MD5**. Packets are authenticated using the *Hash-based Message Authentication Code (HMAC)* coupled with the *Message Digest 5 (MD5)* hash function.

  - **HMAC-SHA**. Packets are authenticated using the *Hash-based Message Authentication Code (HMAC)* coupled with the *Secure Hash Algorithm (SHA)*. SHA is considered to be a stronger authentication algorithm than MD5.

  - Any. The security gateways negotiates which encryption method to use.

Configuring VPN objects   **August 2003   177**

- Use the **Lifetime** text boxes and lists to control the period for creating and exchanging a new set of unique keys.

  If the *Time-based* value expires before the *Throughput* value, key creation and exchange is performed, and likewise, if *Throughput* expires before the *Time-based* value.

- Use the *Locate this Proposal* options to select where to put your new proposal in the *Priority Proposal List*. Security gateways always start from the top of the list when making a query.

20. Click the **Advanced** tab to bring it to the front.

21. Select **Apply VPN to clients only** if you have created a VPN Object where *User* and *User Group Objects* can communicate with *IP Group Objects*, but IP Group Objects cannot communicate with each other.

Note: This is an advanced control, used for a rare case. The default setting will apply to most configurations.

22. Select **Use aggressive mode for clients** if you want to speed-up the time needed for VPNremote Clients to establish a secure connection with the VPN.

23. Select **CRL Checking** if you want to automatically track certificates that have been revoked by a specific Certificate Authority (CA).

Note: This control is only available for certificate based VPNs.

Tunnel endpoints (VPNRemote Clients and security gateways) that use certificates shown by a *Certificate Revocation List (CRL)* are denied access to the VPN. To use this feature, you must obtain a CRL from your Certificate Authority then manually install it in the directory server on a periodic basis. See "Enabling CRL checking" on page 179 for more information.

24. If you use *CRL Checking*, in the **Directory Name of Certificate Authority** text box, type in the distinguished name (DN) of the *certificateauthority* object located in directory server. The object is where the CRL is located.

25. Click **Save**.

Enabling CRL checking

For certificate-based VPNs using IKE negotiation, a security gateway must verify the other certificate of the VSU. When *Certification Revocation List* (CRL) Checking is enabled, the VSU validates the certificate revocation list downloaded from the VPNmanager using the *Certificate Authority* (CA) certificate. The VSU checks the certificate against the validated CRL. If the CRL locates a revoked certificate, the IKE negotiation is cancelled.

To manually install a CRL into Directory Server from the CA's LDAP server:

1. From the CA's LDAP server, obtain the CRL that is associated with your installed issuer certificate.

2. Save the CRL as **crl content.txt**.

3. Open the **crl content.txt** file to extract the necessary CRL information.

4. To extract the necessary CRL information, open the **crl content.txt** file.

5. Locate the dn header with the organization unit (ou) that corresponds to the CRL. For example, dn: ou=vpnet VSU, o=Avaya Inc., c=US

6. Locate the paragraphs starting with **cacertificate;binary** and **certificaterevocationlist;binary**.

   For example,
   cacertificate;binary::MIICKzCCAZSgAwIBAgIQRTP4LaWmlSRKYLv
   86Cphk
   .
   .
   .
   ygPDgMZlQq4oQoNyy26HRAV0yJ==
   certificaterevocationlist;binary::MIIC2zCCAkQwDQYJKoZIhvc
   NAQEEBQAw

7. Copy the **cacertification;binary** and **certificaterevocationlist;binary** paragraphs to a new file.

8. Save the new CRL as **crl.ldif**.

9. Add a certificate dn header to the **crl.idif** file. Use the following dn header format:

   dn: cacertificate=IssuerCRL, ou=VPN Domain, o=DNS Domain objectclass: certificationAuthority

Note: dn specifies where the CRL file is filed.

10. Import the **crl.ldif** file by opening the *Netscape Console* login dialog box.

    • Solaris OS: In the server root, enter **./startconsole**.

    • Windows NT: From the windows Taskbar, click **Start**/**Programs/ Netscape Server Family**/**Netscape Console**.

11. In the **User ID** text box, type in the *Administrative ID* string used during the server installation procedure.

12. In the **Password** text box, type in the *Password* string used during the server installation procedure.

13. Click **OK** to open the *Console* window.

14. From the left pane, select the directory server containing your VPN data.

15. Double click to open the console window for that server.

16. Click the **Configuration** tab to bring it to the front.

17. From the left pane, select **Database.**

18. .Click the right mouse button to select **Import** to import the crl.ldif file.

19. In the **Import Database** window, browse to locate the **crl.ldif** file.

20. Click **Open** to import the **crl.ldif** file.

21. The Import Database message box appears upon successful import.

22. From the VPNmanager Console, click **Config**.

23. From the left pane, click **VPN** then the **General** tab to bring it to the front.

24. In the **General** tab, click **Certificate Based** to enable certificate based VPN checking.

25. Click the **Advanced** tab to bring it to the front.

26. In the **Advanced** tab, click **CRL checking** and enter the CRL dn in the Directory Name of Certificate Authority field.

Note: For the CRL dn, use the same dn used in Step 9.

27. From the left pane, click **Device** then the **Servers** tab to bring it to the front.

28. Add the **Directory Server** IP address and port number. Default clear = 389, default SSL:636

29. Click **OK**.

30. From the **Configuration Console**, click **Save**.

31. From the **Configuration Console**, click **Update Devices**.

During IKE negotiations, the CRL is uploaded to the VSU for CRL checking. The CRL is held in the memory of the VSU.

If the Directory Server has been updated using a new CRL, the cached CRL must be manually removed from the VSU console.

To remove the CRL from the VSU:

1. From the VSU Console, enter **3** for the Utilities menu.

2. From the **Utilities** menu, enter **18** to Show CRL information.

3. After selecting **18** from the **Utilities** menu, a list of serial numbers appear on the screen.

4. Enter **Y** to delete the CRL list.

5. From the VPNmanager main menu, click **Config**.

6. Select **Device**.

7. From the Content pane, select the security gateway that includes the CRL list.

8. Click the **Advanced** tab.

9. Clear the **CRL checking** box.

10. Click **Update Devices**.

Exporting a VPN object to an extranet

Exporting a VPN object is a feature used for interconnecting VPN domains. Each domain views other domains as extranets.

Figure 50. Exporting a VPN Object to an Extranet

DomainA created the VPN Object that was exported to an extranet (DomainB). This method allows members of VPN ObjectA and VPN ObjectB to privately share network resources and communicate.

VPN ObjectA is built with IP GroupA and IP GroupB. IP GroupA is configured with IP address masks for terminal devices in DomainA, and IP GroupB is configured with IP address masks for terminal devices in DomainB.

Domain$_A$
VPN
IP Group
IP Group
Device Object$_A$
Extranet

VPN ObjectA is exported to DomainB.

IP Group ObjectA is configured with Device ObjectA, but Device ObjectA does not get exported to DomainB.

IP Group ObjectB is configured with an Extranet Device. The device is configured with the IP address of Device ObjectB.

Domain$_B$
VPN
IP Group
IP Group
Device O$_B$
Extranet

Device ObjectB is configured from DomainB.The Extranet Device and Device ObjectB have the same IP addresses, therefore, traffic to DomainA will automatically use Device ObjectB for VPN

VPN Object export checklist

Table 9 lists what to do before you export a VPN Object. The terms used by Figure 50 are used for orientation.

Table 9. VPN Object Export Checklist

	Task
☐	For certificate based IKE VPNs, administrators of $Domain_A$ and $Domain_B$ assure that all security gateways which are participating in the extranet connection are using the correct certificates (page 272).
☐	Administrators of $Domain_A$ and $Domain_B$ agree that $Administrator_A$ create the VPN Object that is exported to $Domain_B$.
☐	$Administrator_B$ creates security gateway $Object_B$ and supplies the IP address of that object to $Administrator_A$.
☐	$Administrator_A$ creates IP Group $Object_B$ (page 114) and configures it with an *extranet device* (page 116) having the IP address supplied by $Administrator_B$.
☐	$Administrator_A$ creates security gateway $Object_A$(page 63).
☐	$Administrator_A$ creates IP Group $Object_A$ (page 110) and configures it with security gateway $Object_A$.
☐	$Administrator_A$ creates VPN $Object_A$ (page 156) and configures it with IP Group $Object_A$ and IP Group $Object_B$.
☐	$Administrator_A$ exports VPN $Object_A$ data to $Administrator_B$ (page 182).
☐	$Administrator_B$ imports VPN $Object_A$ data into $Domain_B.$ (page 185)

Export procedure

Exporting a VPN Object involves copying the object data to a file, then sending the file to the extranet administrator, who will import the file into their VPN Domain.

To export a VPN Object:

1. Move to the *Configuration Console* window.

2. From the Icon toolbar, click **VPN** to list all VPN Objects in the *Contents* column.

3. From the Contents column, select the VPN Object that needs to be configured.

4. From the Tools menu, select **Export VPN** to open the Export VPN dialog .

5. From the list box, select the VPN Object you want to export.

6. Click **OK** to open the Export VPN password dialog.

7. In the Password text box, type in a password to protect the exported data.

   From 1 to 16 characters can be used.

8. In the Retype text box, type in your password to confirm it.

9. Click **OK** to open the Save dialog.

10. Use the controls in the Save dialog to select a location for the VPN Object data file.

11. In the File name text box, type in a name for the file, and use VPN as the file name extension.

12. Click **Save** to create the file.

You can now deliver the data file, using e-mail, floppy disk, or FTP, to the extranet administrator. The extranet administrator can use the instructions described by the Importing a VPN object from an extranet section to import the data file.

Importing a VPN object from an extranet

To import a VPN Object data file:

1. Copy the VPN Object data file (created during the Export procedure) into the computer running the VPNmanager Console.

2. Open the *Configuration Console* window.

3. From the Icon toolbar, click **VPN** to list all VPN Objects in the *Contents* column.

4. From the Tools menu, select **Import VPN** to open the Export VPN password dialog box.

5. In the Password text box, type in the password created during Step 7 of the Export procedure.

6. Click **OK** to open the Open dialog.

7. Use the controls in the Open dialog box to navigate to the VPN Object data file.

8. Select the data file, then click **Open** to import the data file.

9. After it is imported, the extranet VPN Object appears in the Contents column.

Rekeying a VPN object

Use the *Rekey* command to create a new key that SKIP VPN tunnel endpoints (security gateways and VPNremote Clients) must use for encryption tasks.

To rekey a SKIP VPN Object:

1. Open the *Configuration Console* window.

2. From the Icon toolbar, click **VPN** to list all VPN Objects in the *Contents* column.

3. From the Contents column, select the VPN Object that needs to be rekeyed.

4. Click the **Actions** tab to bring it to the front.

5. Click **Rekey** to create the new key and open the Rekey message box.

6. Click **OK** to return to the *Configuration Console* window.

7. Click **Update security gateways** to send the key to all security gateways in the VPN.

Chapter 8   Establishing security

This chapter describes the VPNmanager security measures you can configure to establish a secure domain. Included in this chapter is how to set up the following:

- Firewall rules set up (4.2 and later)

- Denial of Service (4.X)

- Services

- Voice Over IP controls (4.X only)

- QoS policy and QoS mapping (4.31)

- Packet filtering firewall (3.x only)

Firewall rules set up

Use the Firewall Rules feature to manage the firewall rules that the domain and the security gateway uses. VPNmanager firewall policy management minimizes configuration complexity and increases scalability. The firewall policy allows deployment of a secure network infrastructure in a relatively short amount of time.

The security gateway uses a rules-based method of packet inspection, where the priority of each rule is determined by its position in the list (highest is top priority). The first match determines the fate of the packet: permit or deny. If no matching rule is found, the default action is to permit the packet.

Note: For devices with VPNos 4.1 and earlier, domain level rules and firewall templates are not available. See "Voice Over IP" on page 202.

Levels of firewall policy management

The Firewall Rules tab is used to manage the firewall rules both at the domain level and at the individual device level in the domain. You can view the Firewall rules and add or edit rules from the VPNmanager Configuration *Console>View>Firewall* command. Firewall policy management includes domain firewall rules, device firewall rules, and firewall templates.

Figure 51. Firewall tab
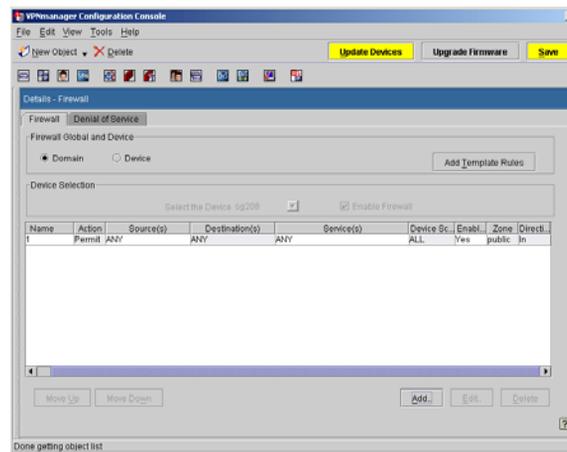
At the domain level, firewall policy management allows the network administrator to set rules across the domain. These rules are referred to as domain level firewall rules. These rules can be applied to all, or some of the devices in the domain. Rules can also be set for specific devices in the domain.

At the device level, firewall policy management allows the network administrator to set rules for a specific device. These rules are referred to as device level firewall rules.

For convenience, you can select from three predefined sets of general firewall rules or templates. Which set of rules you select depends on the interface zones that are configured and your general network requirements. The firewall templates can be used in their default state or as the basis from which a user-defined template can be created.

A brief description of the firewall templates is included in this chapter. For additional information regarding the predefined templates.

Firewall rules

Firewall rules can be defined at the domain level for all devices, for a specific device, or for a device group.

The security gateway uses a rules-based method of packet inspection, where the priority of each rule is determined by its position in the list (highest is top priority). The first match determines the fate of the packet: permit or deny. If no matching rule is found, the default action is to permit the packet.

Domain level firewall rules

Domain, or global, level firewall rules apply to all devices, to device groups, and specific devices within the domain.

You select *View>Firewall* to add domain firewall rules. You can apply common rules to all or some of the devices within the domain when firewall rules are added at the domain level. When firewall rules are applied at the domain level, they can be applied to several devices at the same time which can reduce the complexity of defining security for each device.

To create domain level firewall rules:

1. From the *Configuration Console* window, select **View>Firewall**.

2. From the Firewall tab Firewall Global and Device area, click **Domain**.

3. Click **Add** to start the Firewall Policy wizard.

4. Complete the Firewall Wizard dialog

   - In the **Name** text box, type a unique name that identifies the rule.

   - By default, the **Status** is *Enabled* and the **Action** is *Permit*. Change these if they are not the correct settings.

   - In the Memo area, type notes to describe the firewall (optional)

5. Click **Next** to display the Device dialog. Select the devices to apply the rule. Click **Move Left** to move the selected members to the **Device(s) for this Rule** column.

6. Click **Next** to display the Source dialog. Select the sources; click **Move Left** to move the selected source to the Source column. Click **Next**.

7. From the **Available Destination(s)** column, select the destination; click **Move Left**. Click **Next**.

8. From the **Available Service** column, select the services; click **Move Left**. Click **Next**.

9. The Firewall Wizard Configuration dialog is displayed. From the **Zone** list, select the zone to which you want to apply this rule. For maximum flexibility and capability, the firewall rules for the security gateway can be specified for a particular zone. The packets are checked against the firewall rules at the interface where they are defined.

10. In the **Direction** list, select **In** or **Out**. The direction is in respect to the security gateway.

11. If you want this rule to be logged. select **Enable Log**. If you do not select Enable Log, this rule does not appear in the Monitor>Firewall Log display.

12. If the filter rule set for the intended traffic is also to be applied to the reply packets, select **Keep State**. This function can be applied to TCP, UDP, and ICMP packets.

    If you want to change the default time-out settings for the TCP state, UDP state, or ICMP state, click **Advanced**.

    Note: Keep State sets up a state table, with each entry set up by the sending side. Reply packets pass through a matching filter that is based on the respective state table entry. A state entry is not created for packets that are denied.

    *Although UDP is connectionless, if a packet is first sent out from a given port, a reply is expected in the reverse direction on the same port. **Keep State** "remembers" the port and ensures that the replying packet enters in the same port.*

13. Select the position of the firewall policy in the template.

14. Click **Finish** to return to the Firewall tab.

Device level firewall rules

Device level firewall rules apply to specific devices within the domain. Along with the device-specific rules, the security gateway also inherits the firewall rules that are defined at the domain level. If firewall rules are defined on the security gateway, these device level rules have the highest priority and will take precedence over domain level firewall rules.

To create device level firewall rules:

1. From the *Configuration Console* window, select **View>Firewall**.

2. In the Firewall tab's **Firewall Global and Device** area, click **Device**.

3. Click **Add** to start the Firewall Policy wizard.

4. Complete the Firewall Wizard dialog

   - In the **Name** text box, type a unique name that identifies the rule.

   - By default, the **Status** is *Enabled* and the **Action** is *Permit*. Change these if they are not the correct settings.

   - In the Memo area, type notes to describe the firewall rule (optional)

5. Click **Next** to display the Device dialog. Select the devices to which the rule is applied. Click **Move Left** to move the selected members to the **Device(s) for this Rule** column.

6. Click **Next** to display the Source dialog. Select the sources; click **Move Left** to move the selected source to the Source column. Click **Next**.

7. From the **Available Destination(s)** column, select the destinations; click **Move Left** to move the selected destination to the destination column. Click **Next**.

8. From the **Available Service** column, select the services; click **Move Left**. Click **Next**.

9. The Firewall Wizard Configuration dialog is displayed. From the **Zone** list, select the zone to which you want to apply this rule. For maximum flexibility and capability, the firewall rules for the security gateway can be specified for specific zones. The packets are checked against the firewall rules at the interface where they are defined.

10. In the **Direction** list, select **In** or **Out**. The direction is in respect to the security gateway.

11. If you want this rule to be logged. select **Enable Log**. If you do not select Enable Log, this rule does not appear in the Monitor>Firewall Log display.

12. If the filter rule set for the intended traffic is also to be applied to the reply packets, select **Keep State**. This function can be applied to TCP, UDP, and ICMP packets.

If you want to change the default time-out settings for the TCP state, UDP state, or ICMP state, click **Advanced**.

Note: Keep State sets up a state table, with each entry set up by the sending side. Reply packets pass through a matching filter that is based on the respective state table entry. A state entry is not created for packets that are denied.

*Although UDP is connectionless, if a packet is first sent out from a given port, a reply is expected in the reverse direction on the same port. **Keep State** "remembers" the port and ensures that the replying packet enters in the same port.*

13. Select the position of the firewall policy in the template.

14. Click **Finish** to return to the Firewall tab.

Priority of Firewall rules versus NAT rules

When packets pass through zones that have both Firewall rules and NAT rules set up, NAT rules are applied before the firewall rules are applied. Depending on the type of NAT rule: static, port NAT, or redirection, either the source IP address or the destination IP address of packets are changed. When you set up your firewall rules, you need to consider the type of NAT configured, as you must create the firewall rule to filter on the translated IP address and ports, not on the original address and ports.

Setting up firewall rules for FTP

FTP and Firewall/NAT Operation

The File Transfer Protocol (FTP) uses two TCP connections, one for control, and another for data. The primary methods for establishing the data connection are passive-FTP and active-FTP. In the passive-FTP case, the FTP client makes the data connection to an IP address/port the FTP server has specified. An active-FTP data connection is initiated by the FTP server using information specified by the FTP client.

If the FTP client and FTP server are separated by a firewall, control and/or data connections will normally be blocked. For FTP to function properly, state must be maintained for control and data connections to complete. Typically, a wide range of ports behind the firewall also must be exposed to the external network in order for an external FTP client (passive-FTP) or external FTP server (active-FTP) data connection to be established. So, the location of client/server, as well as mode of operation (active/passive-FTP) dictates the type of firewall issues.

Active-FTP is beneficial to the FTP server administrator, but detrimental to the client side adman. If the FTP server attempts to make connections to random high ports on the client, these packets would almost certainly be blocked by a firewall on the client side. Passive-FTP is beneficial to the client, but detrimental to the FTP server adman. Even if the client makes both connections to the server, the one random high port would almost certainly be blocked by a firewall on the server side. Typically, administrators running FTP servers will need to make their servers accessible to the greatest number of clients, so they will almost certainly need to support passive-FTP. Applications do not consistently use passive-FTP or active-FTP. Modern FTP clients and Internet browsers support a variety of choices.

There are additional problems when the FTP client and FTP server are located on opposite sides of a NAT gateway. Active-FTP clients attempting to gain access to FTP servers from behind a NAT gateway will fail because the data connection received from the FTP server has no address mapping. For example, FTP server attempts to connect to external address of NAT gateway.

Security Gateways and FTP

Two different approaches are available for supporting FTP within the SG environment. One allows the administrator to individually manage each control/data connection through the firewall (FTP-Ctrl, Active-FTP, Passive-FTP services). The other, recommended, uses the FTP-Proxy service.

The first approach allows the administrator to restrict the direction, inbound/outbound, and types of allowed FTP traffic, but does have the potential to expose a large number of ports behind the firewall to outside snooping. An example of a fairly safe configuration would be that of allowing FTP clients on the private zone network to perform passive-FTP. For example, two outbound firewall permit rules, one for FTP-Ctrl and the other for Passive-FTP. Both control and data connection are initiated from within the protected network. An unsafe configuration would be to allow unprotected, external, FTP servers to initiate Active-FTP connections (one outbound FTP-Ctrl firewall permit rule, and one inbound Active-FTP firewall permit rule); in this case Active-FTP allows the full range of ports within the protected network to be accessed by the outside network.

FTP-Proxy service can be incorporated into a firewall rule to concurrently support both passive/active-FTP for protected FTP clients or FTP servers. Configuring an FTP-Proxy rule actually creates one firewall rule to allow the initial FTP control connection and a second redirection rule for the FTP control channel. Upon receiving FTP traffic, the proxy intercepts the control channel exchanges and discovers the type of data connection to be established. It then dynamically creates the appropriate firewall pinhole rule to restrict the protected network ports to which a data connection can be established. The firewall pinholes are removed within a short period of time after the data connection. Thus, FTP-Proxy significantly improves network security as compared to the Passive-FTP (protected FTP server) or Active-FTP (protected FTP client) service cases. It is important to remember that the FTP-Proxy service is applied to a specific zone interface. If network address translation or filter rules are applied to other zone interfaces on the SG that are the source or destination of the FTP traffic, these rules can impact the ability of the proxy to function.

FTP-Proxy does have some issues when operating within a NAT gateway. A protected FTP server must have a routable address, and the router on the unprotected side of the gateway must have static route to it the security gateway interface address is the route. Because this is a proxy application, FTP (TCP) packets destined for external FTP servers

or clients will typically have as source address the address of the interface to which the FTP-Proxy rule was applied. This shows that FTP-Proxy employs some internal address translation.

Note: FTP-Ctrl, Active-FTP, Passive-FTP, and FTP-Proxy services are intended for use with the 'keep-state' firewall rule option.

To add a new firewall rule for FTP-control or passive FTP

1. Complete step 1 through step 12, for adding a new rule. Enter the required firewall information in the wizard.

   *Note: Be sure to define the firewall rule at the interfaces and directions that the FTP server opens a data connection to the client. For example, if the FTP client is on the private side of the security gateway and the FTP server is on the public side of the security gateway, define the interface and direction as **Public/In** or **Private/Out**.*

2. Click **Next**, to display the Source Network Objects dialog. Select FTP Client.

3. Click **Next** to display the Destination Network Objects dialog. Select the FTP Server.

4. Click **Next** to display the Services dialog. Select FTP Control and select Passive FTP.

5. Click **Finish**, to complete the set up of the firewall rules. Click **Save.**

To add a new firewall rule for active FTP

1. Complete step 1 through step 12, for adding a new rule. Enter the required firewall information in the wizard.

2. Click **Next**, to display the Source Network Objects dialog. Select FTP Server.

3. Click **Next** to display the Destination Network Objects dialog. Select the FTP Client.

4. Click **Next** to display the Services dialog. Select Active FTP.

5. Click **Finish**, to complete the set up of the firewall rules. Click **Save.**

Firewall templates

VPNmanager includes predefined firewall templates; high, medium, and low; allowing network administrators to conveniently build secure policies and use the templates as the security foundation in many different network locations.

Administrators can also create their own user-defined templates.

Predefined templates

The predefined templates can be used as a basis for user-defined templates, however; the predefined templates cannot be modified.

For detailed information regarding the predefined templates, see Firewall rules template.

Figure 52. Predefined firewall template, high
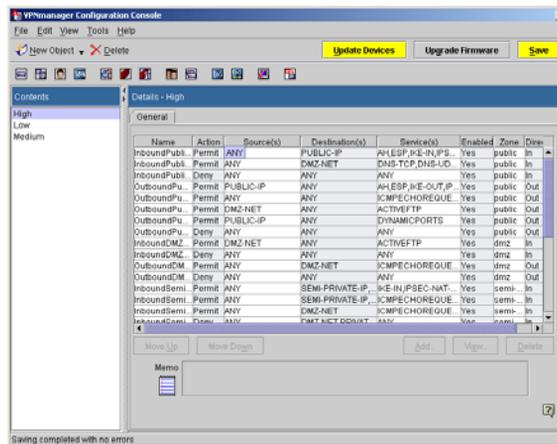
User defined templates

The VPNmanager firewall templates can be used as a general rule set or as a starting point for creating a customized firewall policy, or user-defined template, that conforms to the corporate security requirements. The template rules are enforced on the public interface (the interface through which the security gateway directly connects to the outside world).

To create a user-defined firewall template:

1. Move to the Configuration Console window.

2. From the **Objects** column, select Firewall Template.

3. Click **New Object** to start the New Firewall Template wizard.

4. In the **Name** text box, type in a name for your new firewall template.

5. Select **Template**, **Device**, or **None**.

Parameter	Description
Template	The user-defined template is created using a predefined template – high, medium, or low. Select the template from the drop-down list.
Device	The user-defined template is created using an existing security gateway firewall configuration. Select the existing security gateway from the drop-down list. Using an existing security gateway configuration is also know as cloning the configuration.
None	The user-defined template is created without using a predefined template or an existing security gateway firewall configuration.

6. Click **Apply**.

7. To create a user-defined firewall template, type in a name for your new firewall template, otherwise click **Cancel**.

8. Confirm that the correct user-defined firewall template is selected in the Contents column.

9. Click **Add** to open the Firewall Policy wizard.

10. Type a name for the new rule in the Name text box.

11. Select **Enabled** or **Disabled** in the Status drop-down list to enable or disable the new rule.

12. Select **Permit** or **Deny** in the Action drop-down list to control the flow of packets for this rule.

Permit allows all packets of the selected traffic type to pass.

Deny blocks all packets of the selected traffic type.

13. Click **Next**.

14. Select the set of sources from the available source list.

15. Click **Next**.

16. Select the set of destinations from the available destination list.

17. Click **Next**.

18. Select the set of services from the available services list.

19. Select the **Interface** from the drop-down list.

For maximum flexibility and capability, the firewall rules can be specified on each interface: Public, Private, or Tunnel. The packets are checked against the firewall rules at the interface where they are defined.

20. Select the **Direction** from the drop-down list.

Direction is in respect to the security gateway: in or out.

21. If this rule is to be logged, select the **Log Enable** check box.

22. If this rule is to keep state, select the **KeepState Enable** check box.

The keepstate function allows a rule set for the intended traffic to also be applied to the reply packets. The function can be applied to TCP, UDP, and ICMP packets.

Keepstate sets up a state table with each entry set up by the sending side. Reply packets pass through a matching filter based on the respective state table entry. A state entry is not created for packets that are denied.

23. Click **Advanced** to change the default keepstate values to TCP, UDP, or ICMP.

24. Click **Finish** to return to the Firewall Template General Tab.

Services

The Services property provides a list of predefined traffic types and user-defined traffic types that facilitate the definition of the firewall and Quality of Service (QoS) rules. For instance, you can add a user-defined service for use in firewall rules that allows or blocks a specific type of traffic.

Figure 8-53   Services property
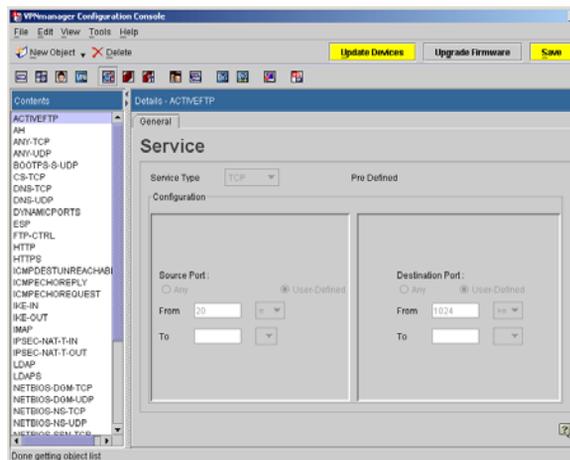
The VPNmanager provides predefined services. The supported predefined services are listed in the Contents column of the Services object.

The predefined services can be used as a general service set or as a starting point for creating a customized service, or user-defined service, that is required for use in the firewall definition. The service types IP, TCP, UDP, and ICMP are provided and parameters for each of these types can be specified in the user-defined service. A comprehensive suite of UDP, TCP, and ICMP filter options are provided.

One or more predefined service can be specified in each firewall rule using the firewall wizard.

Note: The predefined services can be used as a basis for user-defined services, however; the predefined services cannot be modified. To create a user-defined service, click New Object>Services.

Device Group

Device groups help to minimize firewall configuration complexity by allowing network administrator's to create groups of devices that share a common firewall configuration.

To create a device group object:

1. Move to the *Configuration Console* window.

2. From the **Objects** column, select Device Group.

3. Click **New Object** to start the New Device Group Wizard.

4. In the **Name** text box, type in a name for your new Device Group.

5. Click **Apply**.

6. To create another Device Group, type in a name for your new Device Group otherwise click **Close**.

7. Select the devices to be included in the Device Group from the Available Members column.

8. Click **Move Left** to move the selected members from the Available Members column to the Group Members column.

Denial of Service

For servers running VPNos 4.2, configure the DOS to protect the security gateway from attacks by hackers.

A domain has default Denial of Service (DOS) configuration settings that apply to all the devices in the domain. These settings can be seen from the Firewall Object's Denial of Service tab.

The security gateway objects Denial of Service tab is used to change the settings for specific devices. Changing the settings here overrides the domain level settings for that category. When devices are updated, the DOS categories at the device level and the remaining DOS categories from the domain level are sent to the device.

Figure 54. Denial of Service
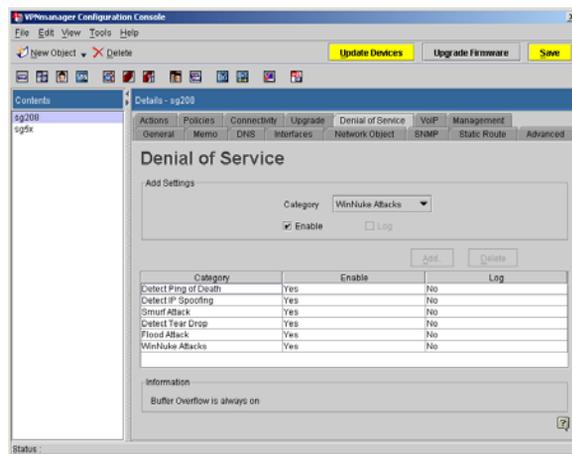
You can enable protection for the following seven areas of attack:

Ping of Death. The ping of death sends packets with invalid lengths. When the receiving system attempts to rebuild the packets, the system crashes because the packet length exhausts the available memory.

IP Spoofing. This attack sends an IP packet with an invalid IP address. If the system accepts this IP address, the attacker appears to reside on the private side of the security gateway. The attacker is actually on the public side, and bypasses the firewall rules of the private side.

Smurf Attack. This attack floods the system with broadcast IP packet pings. If the flood is large enough and long enough, the attacked host is unable to receive or distinguish real traffic.

Tear Drop. This attack sends IP fragments to the system that the receiving system cannot reassemble and the system can crash.

Flood Attack. This attack floods the system with TCP connection requests, which exhausts the memory and the processing resources of the firewall. Flood attacks also attack the UDP ports. This attack attempts to flood the network by exhausting the available network bandwidth.

Note: When you enable Flood Attack, you must also enable the Keep State feature in the Firewall Rules Setup in the Security tab.

WinNuke Attack. This attack attempts to completely disable networking on computers that are running Windows 95 or Windows NT. This attack can be swift and crippling because it uses common Microsoft NetBIOS services. WinNuke attacks ports 135 to port 139 on platforms that are based on Windows 95 and Windows NT.

Buffer Overflow. This attack overflows the internal buffers of the application by sending more traffic than the buffers can process. This attack can contain a program at the end of a packet which can run and attack the system.

To select or deselect DOS categories

1. To set DOS rules, from the Configure Console window, select **View>Firewall**. Select the DOS tab.

2. Select the rules that should be enabled and select to log details about attack attempts, if the log function is available.

3. Click **Save**.

Voice Over IP

For servers running VPNos 4.2 or later, use the VoIP tab to enable or disable Voice over IP (VoIP) and to configure the gatekeeper properties. Definition of the gatekeeper location is with respect to the internal or external firewall definition. Beginning with VPNos Feature Pack 4.31, use the VoIP property to configure the IP trunking properties. You can add, modify, or delete IP trunking configurations.

Voice over IP uses the Internet Protocol to transmit voice as packets over an IP network. So VoIP can be achieved on any data network that uses IP (Internet, Intranets, and Local Area Networks). Here the voice signal is digitized, compressed and converted to IP packets and then transmitted over the IP network. Signaling protocols are used to set up and tear down calls, carry information required to locate users and negotiate capabilities. One of the main motivations for Internet telephone is the low cost involved.

When you add IP trunking, include the IP Trunk name, call model selection, service port, timeout connection in seconds, source zone, and destination zone.

To enable VoIP and add IP Trunking:

1. From the *Configuration Console* window, select **View>Device**. Click the VoIP tab to bring it to the front.

2. Click **Add**. The **Add VoIP Rule Settings** dialog is displayed.

3. In the **Name** field, enter a descriptive, unique name to identify the IP trunk.

4. In the **Call Model** field, select IP Trunking from the drop-down menu.

5. In the **Service Port** field, enter specify the H.323 protocol port. The default is 1720.

6. In the **Timeout** field, specify the idle timeout for the connection. Timeout is the number of seconds that the security gateway allows for inactivity on the connection. If the inactivity continues beyond the specified timeout, the connection is closed. The default is 90 seconds.

7. Click **Next**. The source endpoints dialog appears.

   - In the Zone field, select the zone which the source endpoints are connected to. For example, if the calling trunk endpoints are connected to the public zone, select public zone for this field.

   - In the Network Objects field, specify the source endpoint network object. This should be defined as a network object or network objects with IP addresses equal to the calling trunk endpoints.

   Note: If the network object does not exist, cancel the configuration and create one.

8. Click **Next**. The Destination Endpoints dialog appears.

   - In the **Zone** field, select the zone which the destination endpoints are connected to. For example, if the called trunk endpoints are connected to the private zone, select private zone for this field.

9. Click **Add**. The **Add Destination Endpoint** dialog appears.

10. In the **Endpoint** IP field, specify the IP address of the called trunk endpoint.

11. In the **Proxy IP** field, specify the public IP address that is being shared.

---

12. In the **Proxy Port** field, enter the proxy port. The default is 1720. If this is a Gatekeeper routed call, the default is 1719.

13. Click **Finish**.
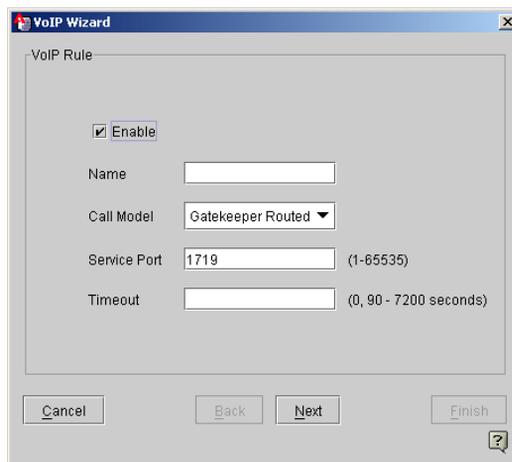
Figure 55. Voice over IP tab
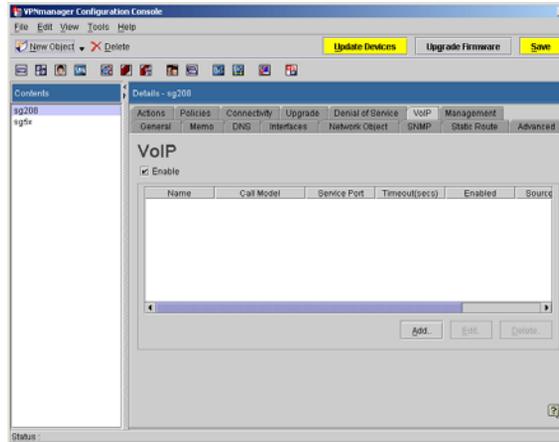
Add gatekeeper settings

When you add a gatekeeper, you include the gatekeeper name or IP address, the location of the gatekeeper with respect to the firewall, the registration, authentication, status protocol, and time-out. Click Add to configure gatekeeper settings for the VoIP configuration. Only one gatekeeper can be configured for a device.

Figure 56. Add gatekeeper setting for VoIP

To enable VoIP and add gatekeeper settings

1. From the *Configuration Console Contents* column, select the device to be configured. Click the VoIP tab to bring it to the front.

2. Click **Add**. The Add Gatekeeper Settings dialog is displayed.

3. In the **Name** field, enter a descriptive, unique name to identify the gatekeeper. Once the name is saved, the name cannot be changed.

4. In the **Location with respect to firewall** field, specify the location of the gatekeeper with respect to the firewall.Select **Internal** if the gatekeeper is behind the firewall. Select **External** if the gatekeeper is outside the firewall.

5. In the **IP Address** field, specify the registration, authentication, and status IP address. The RAS IP address is also the IP address of the gatekeeper, to which the phones to register and send the H.225/RAS messages.

6. In the **Port** field, specify the H.225/RAS protocol port. The default is 1719.

7. In the **Time-out (seconds)** field, specify the idle time-out for the connection. Time-out is the number of seconds that the security gateway allows for inactivity on the connection. If the inactivity continues beyond the specified time-out, the connection is closed. The default is 90 seconds.

8. Click **OK** and then click **Save**.

QoS policy and QoS mapping

The Quality of Service (QoS) function allows the administrator to classify and prioritize traffic based on a DSCP value and/or TCP/IP services and networks. The bandwidth available to a class of traffic can be restricted or rate-limited to a specific percentage of the total upstream bandwidth. This restriction or rate-limiting of bandwidth is only applicable to upstream or outgoing traffic on the interface.

A QoS policy can be created with up to four classes, highest, high, medium, and low. Attributes that can be assigned to these classes are percentage of bandwidth allocation, type of services, network objects, DSCP, and burst.

QoS policies can be mapped to public, public-backup, and semi-private zones. By default, QoS is enabled and VoIP is given the highest priority and there is no restriction of bandwidth or rate-limiting. In the default configuration, VoIP is identified solely by IP precedence values of three and five. This corresponds to the following DSCP values: 24-31 and 40-47.

If QoS is disabled, all traffic receives the same priority. VoIP is treated the same as data traffic.

QoS Policy

This property allows you to add, modify and delete QoS policies. Each policy can include up to four configurable classes, highest, high, medium and low.

You can configure each class according to how network traffic should be prioritized. Each class can contain data, voice or both. Within each class the following is configured:

- **Bandwidth allocation**. Percentage of bandwidth to be allocated to the class. The sum of all allocations for a QoS policy should be 1 to 98%.The remaining 2% is internally allocated by default to ICMP, IGMP, and RSVP. The excess bandwidth not specified in the sum of allocations of the policy is reserved for all other traffic not defined in the classes. Therefore, it is not necessary to create a class for all other traffic. If 0% is allocated, the class is removed from the existing configuration.

  Note: When the media interface is configured, the total upstream bandwidth can be specified in Media Settings and this setting is partitioned to the specified classes.

- **Whether Burst is enabled**. For each class, the burst capability value can be set to Yes or No. The default is No. If bursting is configured for a class, when this class becomes over-limit, it tries to borrow from the unused bandwidth of other classes. If no unused bandwidth is available, the packets are dropped when the class becomes over-limit.

  Caution: Allowing bursting in classes that do not contain voice traffic can affect the availability of bandwidth to voice traffic.

- **DSCP values are assigned**. The valid range of values is 0-63. The default value is 0. This indicates that DSCP is not used for classification. Non-zero DSCP values must be unique among all the classes for one zone because the DSCP value is the only distinguishing factor once a packet is encrypted and sent of the VPN. For example, if DSCP value 10 is assigned to the High class for media interface Ethernet0, DSCP value 10 cannot be assigned to Highest, Medium or Low for Ethernet1. It can be assigned to the High class for Ethernet 1.

  When DSCP value of 0 is specified during configuration, the security gateway generates an internal non-zero DSCP value within the range of 1-63. The non-zero DSCP value generated by the security gateway cannot be used in other classes.

- **Source Network Objects**. Traffic originating from specific networks/ hosts can be selected from existing Network Objects. The source network object specifies the source IP address of the IP packets in this class.

- **Destination Network Objects**. Traffic destined to specific networks/ hosts can be selected from existing network objects. The destination network object specifies the destination IP address of the IP packets in this class.

- **Service.** Traffic can be specified by predefined or user-configured services. A service specifies the IP protocol, TCP/UDP source and destination ports to describe the traffic in this class.

  Note:

  ESP or IKE cannot be assigned with a class as these encrypted packets are assigned to all the classes based on the DSCP value of the packet.

  *It is **not** recommended that a user creates a class with DSCP=, Services=ANY, and Networks=ANY because it is an ambiguous configuration. All traffic not assigned to classes is treated as default traffic. Hence it is not necessary to create such a class.*

  *It is **not** recommended to assign similar traffic in different classes. Example: One class containing any FTP and another class containing "ANY TCP". This would be ambiguous because "ANT+YTCP" would include FTP also. Similar cases might cause ambiguity in classification.*

  *It is **not** recommended to use Services containing ICMP or port- ranges. QoS does not support port-ranges.*

When the *View>QoS* is selected, the screen displays the QoS policies that have been created and their configuration.

Figure 57. QoS policy
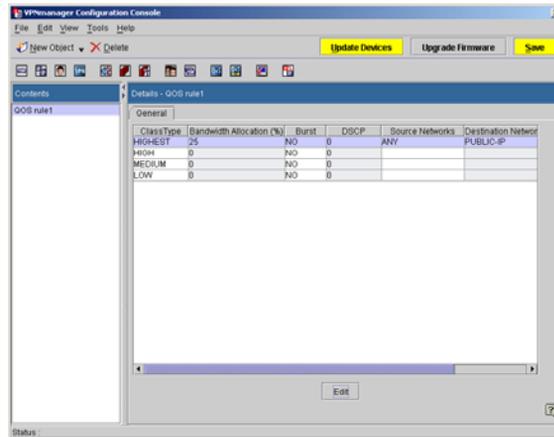
To add a QoS policy

1. Select **NewObject>QoS**. The QoS Policy dialog is displayed.

2. In the **QoS Policy Name** text box, enter a unique QoS policy name. Click **Apply**. Click **Close** to go to the QoS General tab.
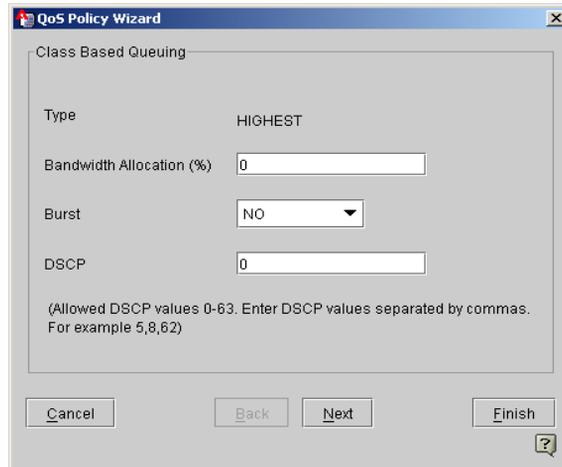
3. Next configure each class setting with associated values. Click the row for the type to be configured. The Class Based Queuing dialog appears.

Figure 58. Modify QoS bandwidth. burst and DSCP value screen

4. Configure bandwidth, burst and DSCP values.

   • Enter the percentage of bandwidth to be allocated for this type.

   When classes are configured, it is recommended that the sum total allocation of all the classes be less than 98% and allow bursting to take advantage of the unused bandwidth. 2% is always internally allocated to control traffic.

   • Burst is set to **No**. Change to Yes if bursting should be allowed.

   If bursting is configured, when this class becomes over-limit, it tries to borrow from the unused bandwidth. If there is no unused bandwidth, then the packets are dropped when the class becomes over-limit.

   • The same DSCP value cannot be assigned in multiple classes for one interface. Do not specify the same DSCP-Services-Network combination in multiple classes.

   If DSCP will not be specified as a criteria in a class, leave the DSCP default value of 0. In this case, it is recommended to assign unique services/networks to this class. Do not assign *ANY* service and *ANY* network objects.

5. Click **Next**. The Source Network Objects dialog appears. Select the network object from the **Available** source and move it to the **Members** column.

6. Click **Next**. The Destination Networks Objects dialog appears. Select the network object from the **Available** destinations and move it to the **Members** column.

7. Click **Next**. The Services dialog is displayed listing the predefined and user defined traffic types. Select the services from the **Available** column and move to the **Members** column.

   Do not assign ESP or IKE as a service within a class as these encrypted packets are assigned to all the classes based on the DSCP field on the packet.

8. Click **Finish**.

9. Complete the configuration of each of the classes from step 3.

10. When the classes have been configured, click **Save**.

QoS mapping

QoS Mapping is the mapping of a QoS policy to a zone. A zone can map to only one QoS policy, but a QoS policy can be applied to multiple zones.

When you map QoS policies consider the following:

- If QoS is configured over multiple interfaces, the DSCP values belonging to a class for a particular zones should not belong to a different class for other zones.

- When QoS is applied over multiple zones, the QoS policies should be identical in definition of classes, DSCP, and service-networks attributes. The only difference in these QoS policies should be in the bandwidth allocation percentage.

Mapping QoS policies

After the QoS policies are created, they can be mapped to either public, public-backup or a semi-private zone at the domain or device level.

1. Select **View>QoS Mapping**. The QoS General tab is displayed. Click **Add.** Select either **Domain** or **Device.** The QoS dialog is displayed.

2. For Domain QoS mapping, select the devices that are to be members for this QoS mapping.

3. Select the **Zone** to be configured.

4. Select the QoS policy that should be applied.

5. Click **OK** and then click **Save**.

Packet Filtering

The Packet Filtering feature is available for devices with VPNos 3.x.

VSUs have a multiprotocol filtering service that analyzes packets (also known as frames) at the Application, Transport, and Network Layers. The headers of the packets can be examined then compared to filter rules organized in an *Access Control List (ACL)*. Filters can be specifically created for inbound and/or outbound traffic, and the state of a connection. Additionally, reports about filtering activity can be sent to a common SNMP manager for viewing.

The ACL can hold up to 200 policies. The default policy is to permit the packet. The default policy is automatically applied if no other policy is configured or if configured policies do not match. The ACL can be organized in a specific sequence so that one policy has a precedence over another. All policies can be customized to meet your needs, and they can be turned off or on at any time.

Policies are semi-automatically created (one at a time) by using the *Packet Filtering Policy Wizard*. As an auxiliary method, policies can also be created at the *VSU Console* (not explained in this guide).

What can be filtered

Table 10 lists the specific types of traffic that can be filtered.

Table 10. Traffic types that can be filtered

User-defined TCP	Exec	Netware-IP/ TCP	VPN-AuthGW
User-defined IP	Finger	Netware-IP/ UDP	VPN-KeepAlive
User-defined UDP	FTP	Nettimep	VPtunnel
AURP	FTP/data	NFS	Who
Bootpc	Gopher	NFS/TCP	WWW-HTTP
Bootps	Gopher/UDP	NNTP	WWW-HTTP/ UDP

Table 10. Traffic types that can be filtered

Bordergw	ICMP	NNTP/UDP	XDMCP
Chargen	IDIRACCP	NWIP-DSS/ TCP	
Chargen/UDP	IPX/TCP	NWIP-DSS/ UDP	
CMD	IPX/UDP	Printer	
Discard	IPrelay	Relaychat	
Domain	IPtunnel	SMTP	
Domain/TCP	Kerberos	SNMP	
Discard/UDP	Login	SNMP-Trap	
Dynamic/TCP	Nameserver	Telnet	
Dynamic/UDP	Nameserver/ TCP	TFTP	
Echo	NetBIOS/TCP	UUCP	
Echo/UDP	NetBIOS/UDP	UUCP-Path	

Packet Filtering and NAT

Network address translation (NAT) and packet filtering services can be run simultaneously. Depending on the direction of the traffic, the VSU automatically determines which sequence the services will run.

For inbound packets (to the WAN), NAT is run first, then filtering. For outbound packets, filtering is run first, followed by NAT.

Figure 59. Policy Manager, Packet Filtering/QoS
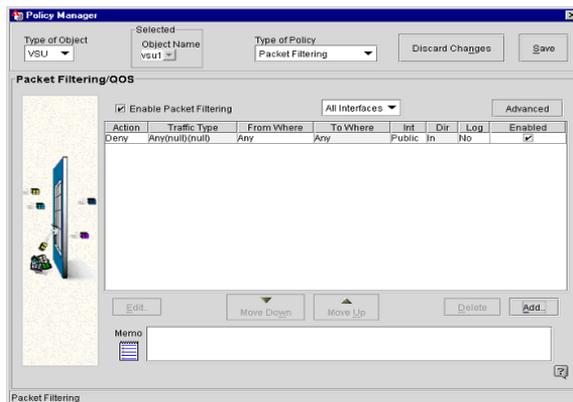
Clicking on the Edit or Add buttons launches a Packet Filtering Policy Wizard that guides you through configuration of the desired packet filtering.

Advanced

The Advanced tab accesses specific types of filters that are activated through checkboxes.

Permit/Deny non-VPN traffic Radio Buttons

The Radio Buttons at the top of the Packet Filter Rule-Advanced screen are set according to your security policy. They include:

- **Permit all non-VPN traffic** - When checked, all non VPN traffic is allowed to pass through the VSU.

- **Deny all IP non-VPN traffic** - When checked, all non-IP traffic is prevented from passing through the VSU. All non-VPN IP traffic is dropped except for the following: ICMP, IGMP, GGP, EGP, IGP, DGP, EIGRP, and OSPF.

    Note: This mode should be used when the VSU dedicated to VPN traffic and is the only device between the private and the public networks.

- **Deny all non-VPN traffic** - When checked, all non-VPN traffic is prevented from passing through the VSU. This mode blocks non-IP traffic and non-VPN traffic including broadcast traffic, IP-multicast traffic and other traffic containing routing information.

  Note: This mode should be used when the VSU is dedicated to VPN traffic and is in parallel with another device (such as a router or firewall) that can resolve ARPs from the private network to the Internet gateway. This mode should not be used when the VSU is the only path between network devices and a router with which those devices need to communicate.

Drop all fragments When checked, discards all non-expected IP packet fragments. Normally used to prevent tiny fragment attacks (RFC1858).

Drop all short packets When checked, this function drops all packets that are not a valid size.

Keep filter statistics (SNMP) When checked, statistics for this filter are reported via SNMP.

Memo Use this area to record comments or notes about your filter.

Add Packet Filtering Policy

This screen performs two basic functions, selection of the desired action, and selection of the traffic type for which a filter is constructed. Additional buttons are provided for Advanced functions, Close, Next, and Finished

Action Two basic actions may be selected: Permit, or Deny. As you would expect, Permit allows all packets of the Traffic type selected to pass, while Deny blocks all packets of the Traffic type selected.

QoS Mark QoS Mark is a drop-down menu of choices used when differentiated levels of priority IP packet routing is used. This allows Quality of Service markings to be placed in the outer IP header when applying the IPSec tunnel mode, thereby allowing "QoS-aware" devices within an MPLS cloud to maintain the desired level of priority in handling the packets. Packets to be marked at the VSU are indicated further specification in the filtering criteria.

A comprehensive list of QoS preset markers are provided in the drop-down menu. For information on the use of these markers, or constructing user defined markers, please refer to the following for details.

- RFC 2474: Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers (http://www.ietf.org/rfc/rfc2474.txt?number=2474)

- RFC 2598: An Expedited Forwarding PHB (http://www.ietf.org/rfc/rfc2598.txt?number=2598)

- You may also wish to check out (http://www.ietf.org/html.charters/diffserv-charter.html) which contains a set of links to relevant related RFC's including 2497 and 2598.

Traffic Type The fields and drop-down lists in this section change according to the IP Protocol type selected. Depending on the traffic type selected (user-defined TCP and user-defined UDP), Source and Destination fields appear to collect additional parameters.

If the Traffic Type selected is user-defined IP, a Protocol ID field appears.

A comprehensive suite of UDP, TCP, and ICMP filter options are provided.

Keep State Appears when user-defined TCP or user-defined UDP traffic type is selected. This function allows a filter rule set for the intended traffic to also be applied to the reply packets. This function can be applied to both TCP and UDP packets.

Keep State sets up a state table, with each entry set up by the sending side. Reply packets pass through a matching filter based on the respective state table entry.

Note: Although UDP is connectionless, if a packet is first sent out from a given port, a reply is expected in the reverse direction on the same port. Keep State essentially "remembers" the port and lets the replying packet enter in the same port.

Source Port Appears when User-defined TCP or User-defined UDP selections are made. Select the Range (Any or User-defined), then enter the from: and to: values. The port range is inclusive. If you want to choose a single port, simply specify the same port as both start and end port.

You can also choose an operator on the port range (= means in the port range and != means out of the port range).

Destination Port Appears when User-defined TCP or User-defined UDP selections are made. Select the Range (Any or User-defined), then enter the from: and to: values. The port range is inclusive. If you want to choose a single port, simply specify the same port as both start and end port.

You can also choose an operator on the port range (= means in the port range and != means out of the port range).

Comparator Permits logical include (=) or exclude (!=) operation on the range entered. For example, if you want to block ports 1024 through 1250, you would enter (Action = Deny) from: 1024 to 1250 and select = as the comparator value.

From/Where

- **Type**. Choices are Network/Mask Pair or Any.

- **IP Network Mask Pair**. Identify the source IP address to which the filter rule applies.

To Where

- **Type**. NetworkMask Pair or Any.

- **IP Network Mask Pair**. Identify the source IP address to which the filter rule applies.

The Filtering Policy in progress

This area presents a dynamically updated summary of the filter parameters based on the current selections.

- **Interface**. Select the private, public, or Tunnel interface of the VSU to which this filter is applied.

- **Direction**. In or Out.

- **Log**. Yes or No. If yes, the maximum number of bytes per entry can be specific

Locating this filtering policy

Establishes the position of this filter rule in the Policy Filter list. Selections are Beginning of List, End of List, and After Selected Item.

The filtering policy in progress

This area presents a dynamically updated summary of the filter parameters currently selected.

When you are satisfied with your filter configuration, click on the Finished button to build the filter. The filter is then automatically placed in the main Packet Filtering window list according to the order indicated by the "Locate This Filtering Policy" radio button.

Running the packet filtering policy wizard

The Packet Filtering Policy wizard is used for creating filtering policies.

To create a filtering policy:

1. Move to the *Configuration Console* window.

2. From the **Contents** column, select the VSU where the new rule has to be located.

3. Click the **Policies** tab to bring it to the front.

4. From the *drop-down list*, select **Packet Filtering**, then click **GO** to open the *Policy Manager for Packet Filtering*.

5. Click the **Add** button to start the *Packet Filtering Policy Wizard*.

6. From the **Action** drop-down list, select **Permit** or **Deny** to control the flow of packets for this policy.

   Note: As you build your policy, its parameters populate the "Filtering Policy in Progress" text box, which is located at the bottom of the wizard.

7. If you want to make a note about this policy, in the **Memo** text box, type in a note.

8. From the **IP Protocol Type** drop-down list, select the type of traffic you want to control.

   Controls appear in the *Traffic Type* options box after you select an item from the list.

9. Use the controls to configure the parameters for the policy.

10. Click **Next** to continue using the wizard until your policy has been built, then click **Finished** to return the *Policy Manager for Packet Filtering* window.

Your new policy appears in the Access Control List.

11. Click **Save** to save your work.

Running the Policy Manager for packet filtering

The Policy Manager for Packet Filtering is used for starting and stopping filtering services, managing the ACL, and for configuring advanced filtering options. Figure 59 shows Policy Manager for packet filtering.

Starting and stopping filtering services

To start or stop filtering services:

1. Move to the *Configuration Console* window.

2. From the **Contents** column, select the VSU where the services need to be started or stopped.

3. Click the **Policies** tab to bring it to the front.

4. From the *drop-down list*, select **Packet Filtering**, then click **GO** to open the *Policy Manager for Packet Filtering* window.

5. Select the **Enable Packet Filtering** check box to start the filtering services, or clear it to stop the services.

6. Click **Save** to save your work.

Managing the ACL

The filtering policies in the *Access Control List (ACL)* can be edited, have their sequence changed, or even deleted. A VSU starts from the top of the ACL when it begins to filter a specific packet. Keep the first policy you want to apply to the packet first at the top of the list.

Note: A packet is filtered against the ACL policies defined in the ACL list in the list order. The packet is matched against policy number 1 first, then policy number 2, then policy number 3, and so on until the packet finds a match or it exhausts the list. If a match is found, the VSU applies the action specified in the policy to the packet. If no match is found, the VSU

applies the default policy to the packet. The default policy is to permit the packet.

To edit, change the sequence, or delete a filtering policy:

1. Move to the *Configuration Console* window.

2. From the **Contents** column, select the VSU where you want to modify the ACL.

3. Click the **Policies** tab to bring it to the front.

4. From the *drop-down list*, select **Packet Filtering**, then click **GO** to open the *Policy Manager for Packet Filtering*.

5. From the ACL, select a specific filtering policy.

6. Use Table 11 for performing specific ACL management tasks.

Table 11. ACL commands

Command	Description
Edit	Use this command to modify the filter policy through the Packet Filtering Policy Wizard.
Move UP	Click this button to move the filter policy higher in the ACL.
Move Down	Click this button to move the filter policy lower in the ACL.
Delete	Click this button to remove the filter policy from the ACL.

7. When finished, click **Save** to save your work.

Configuring advanced filtering options

To configure advanced filtering options:

1. Move to the *Configuration Console* window.

2. From the **Contents** column, select the VSU where the new filtering policy needs to be located.

3. Click the **Policies** tab to bring it to the front.

4. From the *drop-down list*, select **Packet Filtering**, then click **GO** to open the *Policy Manager for Packet Filtering*.

5. Click **Advanced** to open the *Packet Filter Rule-Advanced* dialog box.

6. Use Table 12 for determining which option you want.

Table 12. Packet Filter rule-advanced options

Option	Description
Permit all non VPN traffic	Select this button to permit all non VPN packets.
Deny all IP non VPN traffic	Select this button to block all IP non VPN packets.
Deny all non VPN traffic	Select this button to block all non VPN packets.
Drop all IP fragments	Select this check box to block all IP packets that have been fragmented. See MTU Path Discovery for information about packet fragmentation.
Drop all Short IP Packets	Select this check box to block all IP packets that are unusually small. The following are considered short packets. • IP packets shorter than 20 bytes. • TCP packets shorter than 40 bytes. • UDP packets shorter than 28 bytes. • ICMP packets shorter than 28 bytes.
Keep Filter Statistics	Select this check box if you want to send the "packet filtering log" to a common SNMP manager. The manager that is used is configured in the Static Route tab.

7. Click **OK** to return to the *Policy Manager for Packet Filtering*.

8. Click **Save** to save your work.

Marking packets for differentiated services (QoS)

If your network is running Differentiated Services, a VSU can be configured to mark specific IP packets for specific types of services.

About Differentiated Services

IP packets move from router to router by using *Routing* and *Packet Forwarding* processes. The routing process involves building and maintaining a routing table. The packet forwarding process involves comparing the destination address of a packet with entries in a routing table to determine where to send the packet. Furthermore, there is a component of the forwarding process that can be used for controlling the behavior of a specific type of packet. The component is called *Differentiated Services*, which is also known as *DiffServ* or *Quality of Service (QoS).*

Differentiated Services involves using an *identification system* to mark IP packets. When the marked packet is processed by a router that is running Differentiated Services, the router compares the mark with a list of *Packet Forwarding Behavior (PFB)* rules. If a packet matches a specific rule, the rule is used to the forward the packet. A PFB rule defines specific forwarding characteristics such as minimum bandwidth requirements and the transmission precedence of one type of packet relative to other packets.

The *identification system* involves two kinds of marks: *User Defined* and *Predefined*. The user defined mark is in the form of a number, where the number can be from 0 to 63, and identifies a customized PFB rule. The predefined mark is in the form of alpha numeric characters, and it identifies generic PFB rules that come with your router. A predefined mark is also known as a *Behavior Aggregate*.

Note: For additional information about Differentiated Services, see the following documents.

- Your router's documentation.

- RFC 1812, "Requirements for IP Version 4 Routers"

- RFC 2474, "Definition of the Differentiated Services Field in the IPv4 and IPv6 Headers"

- RFC 2475, "An Architecture for Differentiated Services"

How a VSU marks packets

Before a VSU can run marking services, it's loaded with a list of *Packet Marking Rules*. As packets move through the VSU, it examines the header fields of every packet. The information gathered from the header

is compared to the list of rules. If the comparison results in a match, the *Type of Service (ToS)* field of the header is marked. Marking can be performed on packets entering and/or exiting the VSU. Be aware that only IPv4 packets can be analyzed and marked.

Types of marking rules

Two kinds of packet marking rules can be created.

- A rule can be made to examine the ToS field of a header and copy the existing mark to the TOS field of the new packet, which is entering or exiting the VSU. This is known as inheriting a mark.

- A rule can be made to skip the ToS field, but examine the remaining fields of the header. If a match is made, then the ToS field is appropriately marked.

How to create a packet marking rule

The *Packet Filtering Policy wizard* is used to create a *Marking Rule*. VPNmanager Console is then used to update a specific VSU with the new rule. The different types of marks used in a rule are briefly described in Step 7.

Before marking any packets, you must gather the information described in Table 13. Basically, the type of marks, type of packets, and the direction of packet flow (in and/or out of the VSU) is needed to create a marking rule.

Table 13. IP packet marking information

Item	Description
User defined marks	Identify which user defined marks are being read by your routers.
Packet type associated with a specific user defined mark	Examine the *PFB rule* associated with a specific user defined mark to identify the type of IP packet being marked.
Predefined marks	Identify which predefined marks are being read by your routers.
Packet type associated with a specific predefined mark	Examine the PFB rule associated with a specific predefined mark to identify the type of IP packet being marked.

To create a packet marking rule:

1. Move to the *Configuration Console* window.

2. From the **Contents** column, select the VSU where the new rule has to be located.

3. From the **GO** menu, select **Policy Manager**, to open the *Policy Manager* window.

4. From the **Type of Policy** drop-down list, select **Packet Filtering** to view the *Policy Manager for Packet Filtering*.

5. Click the **Add** button to start the *Packet Filtering Policy Wizard*.

6. From the **Action** drop-down list, select **Permit** to activate the *QoS Mark* drop-down list.

   Note: As you build your Packet Marking Rule, its parameters populate the "Filtering Policy in Progress" text box, which is located at the bottom of the wizard.

7. From the **QoS Mark** drop-down list, do one of the following.

   • Select **Inherit** if you want the VSU to examine the ToS field of packets entering the VSU, then copy the QoS mark to the ToS field of the *payload packet header* (exiting the VPN tunnel) or the *VPN packet header* (entering the tunnel). Which packet depends on the rule being created.

   Note: If you do one of the following, assure that the mark used for the rule matches the mark configured in your router(s).

   • Select **User Defined** if you want to activate the *User Defined* text box, then type a specific mark into the box. The mark must be a number from 0 to 63.

   • Select a specific **CS** mark if you want to use a predefined *Class Selector* mark. Although the specific CS mark used must be the same as the one configured in your router(s), these marks serve as a backward compatibility mechanism for *IP Precedence Marks*, which predate modern QoS Marks.

   • Select a specific **AF** mark if you want to use a predefined *Assured Forwarding* mark. The AF mark identifies which level of precedence the packet must be dropped from the stream if traffic congestion limits are exceeded.

- Select the **EF** mark if your router(s) are marking packets with the predefined *Expedited Forwarding* mark. The EF mark assures that a packet does not get promoted or demoted to a specific packet forwarding behavior.

8. Continue to use the *Packet Filtering Policy Wizard* to define the remaining parameters of your packet marking rule. Some of the parameters are listed in Table 14 the table. Packets which match the values of these parameters get their ToS field changed to the QoS Mark selected in Step 7.

Table 14. Parameters used in a Packet Marking Rule

Parameter	Description
Traffic Type	Use the *Traffic Type* controls to configure which IP protocol the rule must contain.
Source Address	Use the *From Where* controls to configure which source address the rule must contain.
Destination Address	Use the *To Where* controls to configure which destination address the rule must contain.
VSU Interface	Use the *Interface* drop-down list to apply the rule to the VSU public, private, or Tunnel interface.
Direction	Use the *Direction* drop-down list to apply the rule to packets that are entering or exiting the VSU.

9. Continue using any remaining controls in the wizard to complete your new rule.

10. Click **Finished** to return the *Policy Manager for Packet Filtering* window.

   Your new rule appears in the Access Control List.

11. Click **Save** to save your work.

Packet filtering firewall

The security gateway uses a rules-based method of packet inspection, where the priority of each rule is determined by its position in the list (highest is top priority). The first match determines the fate of the packet: permit or deny. If no matching rule is found, the default action is to permit the packet.

Figure 60. Policy Manager for firewalls
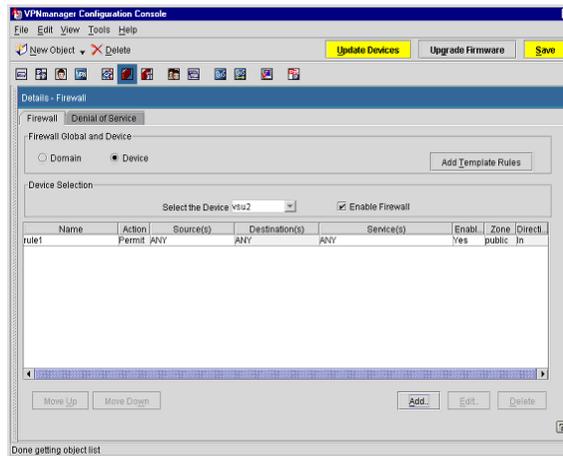
To use the firewall policy management:

1. Move to the Configuration Console window.

2. From the Contents column, select the **security gateway** that the policy is applied.

3. Click the **Policies** tab to bring it to the front.

4. Select **Firewall** from the Policies drop-down list.

5. Click **Go** to open the policy manager for firewall.

Add firewall policy

To add a firewall policy:

1. Click **Add** to open the firewall policy wizard.

2. Type a name for the new rule in the Name text box.

3. Select **Enabled** or **Disabled** in the Status drop-down list to enable or disable the new rule.

4. Select **Permit** or **Deny** in the Action drop-down list to control the flow of packets for this rule.

Parameter	Description
Permit	Allows all packets of the selected traffic type to pass
Deny	Blocks all packets of the selected traffic type

5. Click **Next**.

6. Select the set of sources from the available source list.

7. Click **Next**.

8. Select the set of destinations from the available destination list.

9. Click **Next**.

10. Select the set of services from the available services list.

11. Select the **Interface** from the drop-down list.

    For maximum flexibility and capability, the firewall rules can be specified on each interface: public, private, or Tunnel. The packets are checked against the firewall rules at the interface where they are defined.

12. Select the **Direction** from the drop-down list.

    Direction is in respect to the security gateway: in or out.

13. If this rule is to be logged, select the **Log Enable** check box.

14. If this rule is to keepstate, select the **KeepState Enable** check box.

The keepstate function allows a rule set for the intended traffic to also be applied to the reply packets. The function can be applied to TCP, UDP, and ICMP packets.

Keepstate sets up a state table with each entry set up by the sending side. Reply packets pass through a matching filter based on the respective state table entry. A state entry is not created for packets that are denied.

15. Click **Advanced** to change the default keepstate values to TCP, UDP, or ICMP .

16. Click **Finish** to return to the Policy Manager for Firewall.

Chapter 9   Using advanced features

This chapter explains about the advanced functions of VPNmanager. The following tabs can be used to configure advanced functions for domains and for security gateways:

- Device Advanced tab

- Failover tab

- Advanced Action tab

- Servers tab

- Device High Availability tab

- TEP Policy tab

- Resilient Tunnel tab

- Policy Manager - My Certificates

Device Advanced tab

The Device Advanced tab contains properties that are used to configure security gateway parameters for unique circumstances.

Note: The properties displayed within the Device Advanced tab are determined by the release of VPNos the device is running.

- *VPNos 3.x includes all options described*

- *VPNos 4.0 and 4.1 include MTU Path and Private IP Address*

- *VPNos 4.2 includes MTU Path Discovery*

- *VPNos 4.31 includes MTU Path Discovery, NAT Traversal, and Port for Dyna Policy Download*

The Private IP Address and the local DHCP server IP address are combined beginning with VPNos 4.2. Previously the Private IP Address was located on the Advanced tab.

Figure 61. Security gateway, Advanced tab
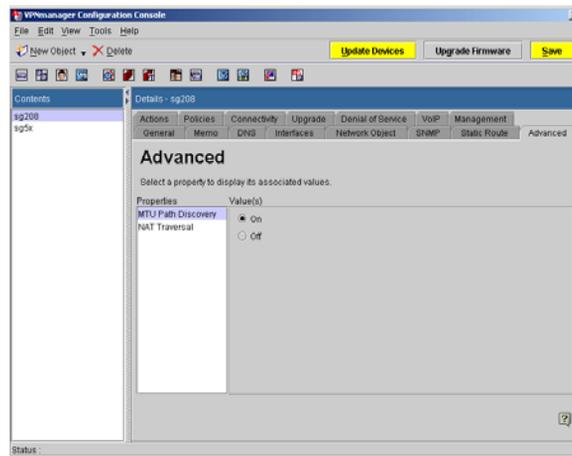
Active VPN sessions (VPNos 3.x)

Active VPN sessions are defined as remote client traffic and security gateway-to-security gateway traffic. An SNMP Agent running on a security gateway is used to collect information about these sessions. This information is private in nature, and can be viewed by VPNmanager Console using its monitoring feature or common SNMP Management software running on another computer. The agent sends the information in clear text. Use *Active security gateway Sessions* for controlling the flow of session information. The direct path would be to the private port. The Active VPN Session options are shown in Figure 61.

Note: If your organization's security policy dictates that this traffic be secure, TEP Policy (in the Main Console Preferences tab) can be turned on to encrypt this traffic.

For additional information about using third party SNMP Manager, see "Using SNMP to monitor the device" on page 277.

To control the flow of session information:

1. From the *Device>Contents* column, select the security gateway you want to configure.

2. Click the **Advanced** tab to bring it to the front.

3. From the *Properties* column, select **Active security gateway Sessions** to display the *Permit* check box.

4. Do one of the following:

   - Select the **Permit** check box if you want session information to be sent by the security gateway SNMP Agent.

   - Clear the **Permit** check box if you *do not* want session information to be sent by the security gateway SNMP Agent.

5. Click **Save**.

6. When you want to send the configuration the security gateway, click **Update Devices**.

ARP

Determines the VSU use of its MAC addresses. In the default mode (Bind one IP address to each port), the Primary IP address is bound to the MAC address of the public port. If a private IP address is configured, that address is bound to the MAC address of the private port of the VSU. In this mode, all packets originating from the VSU destined for the public network uses the public port's MAC address as the packets' source MAC address.

Examples of public network destined traffic are:

- IPSec packets being tunneled to a member VSU

- SNMP Get Responses being sent to a VPNmanager console residing on the public side of the VSU

- Traps sent to a VPNmanager console residing on the public side of the VSU

Also in the default mode, all packets originating from the VSU destined for the private network use the private port's MAC address as the packets' source address.

Examples of traffic destined for the private network are:

• Decapsulated IPSec packets destined for the private network.

• SNMP Get Responses being sent to a VPNmanager console residing on the private side of the VSU

• Traps sent to a VPNmanager console residing on the private side of the VSU

Note: It is important to remember that ARP often works in conjunction with the Advanced Filter setting.

Device in parallel with firewall or router For example, if you setup a VSU in parallel with a network device that provides firewall and routing services and you only want the VSU to:

• send ARPs for addresses in its primary IP address space out the public interface and,

• send ARPs for addresses in its private IP address space out the private interface,

you would then want to:

1. Set the above to "Bind one IP address to each port" and

2. Set the Advanced Filter to Deny all non-VPN traffic. The latter prevents a ARP from going out both interfaces.

Device in One-Arm Mode. Suppose you have deployed the VSU in one-arm mode (which requires that only the private port be plugged into the network) and you have used the *Bind one IP address to each port* setting. This topology requires that the Advanced Filter setting be "Permit all non-VPN packets". This allows ARPs for the VSU's primary IP address that come in the private port (remember it is the only port plugged in) to be resolved.

The "Bind Both Primary and Private IP Address to the Private Port" setting is available for legacy support. In particular, with this setting the VSU always ARPs out both ports independent of the Advanced Filter setting and it always uses the private port's MAC address for all packets

originating from the VSU. Use this setting if you need a VSU running VPNOS 3.1.xx, or later, to support this legacy behavior.

Generally, only if the VSU firmware is earlier than 3.1, *and* the VSU is the only device between the internet and the private network (not in parallel with a firewall), is Bind both Primary and Private IP addresses to private port checked.

Default Gateway for VPN Traffic

The Default Gateway for VPN Traffic policy allows the administrator to specify a gateway that is used for either decrypted traffic, encrypted traffic, or both.

This configuration is commonly applied to a VSU in the following topology:

Figure 62. Common Default Gateway for VPN Traffic topology
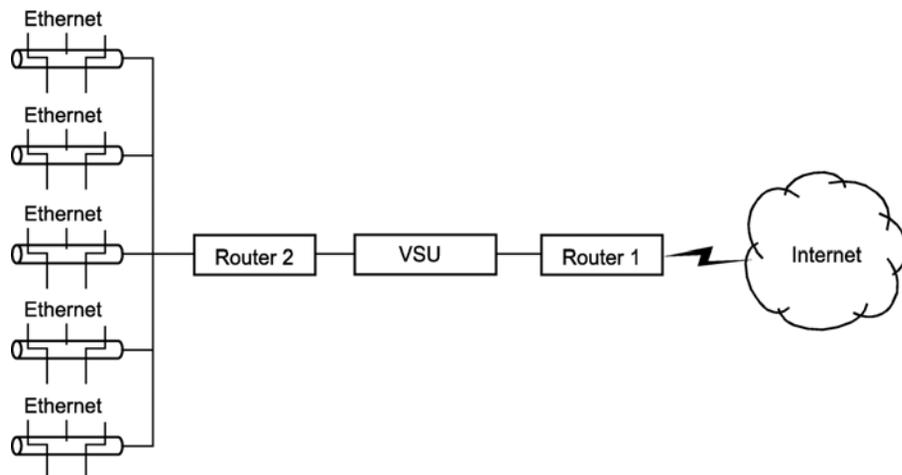
Figure 62 shows the default gateway of the VSU as R_1, the Internet gateway. The VSU is configured to protect several LANs on the other side of R_2, the router on the private side of the VSU.

In this topology, the administrator configures R_1 as the default gateway of the VSU and R_2 as the Default Gateway for VPN Traffic with the decrypted box checked. Using this configuration and checking the

decrypted traffic box, all decrypted VPN traffic would be forwarded to R_2 and all encrypted traffic would be forwarded to R_1. In this application, the Default Gateway for VPN Traffic removes the need for a configured static route on the VSU for each protected LAN.

Note: Configured static routes take precedence over the Default Gateway for VPN Traffic.

To force a default route for secure traffic:

1. From the *Device>Contents* column, select the VSU you want to configure.

2. Click the **Advanced** tab to bring it to the front.

3. From the *Properties* column, select **Default Gateway for VPN Traffic** to display the default route controls.

4. Select the **Enable** check box to activate the property.

5. In the **Default Gateway for VPN Traffic** text boxes, type in the IP address for the default gateway routing device.

6. Select **Encrypted traffic** or **Decrypted traffic** check box to specify the traffic of the routing device.

7. Click **Save**.

8. When you want to send the configuration to the VSU, click **Update Devices**.

MTU Path Discovery

When a device communicates with another network device, it attempts to discover the largest packet it can transmit to the other network device. The largest packet the network can transmit is called *maximum transmission unit (MTU)*.

As a packet is routed through different networks, it may be necessary for a router to divide the packet into smaller pieces because it might be too large to transmit as a single packet on a different network. This may occur at the interfaces of physically different networks.

The MTU of a security gateway passing secure traffic is 1404 bytes, which includes the additional IPSec information. The MTU of a security gateway passing unprotected traffic is 1514 bytes.

If *MTU Path Discovery* is running, a security gateway does not convert the following types of packets into secured traffic, and it uses an ICMP message to ask the source of the packets to fragment them.

- Packets larger than 1404 bytes
- Packets with the *Don't Fragment Bit* set
- Packets being the first fragment in the IP datagram

There are two reasons why you may not want a security gateway to participate in MTU Path Discovery.

- A firewall sits between the security gateway and the source of packets needing VPN services. This would prevent the source from receiving security gateway ICMP messages indicating that fragmentation is needed.
- The source of packets needing VPN services does not fragment packets, even when notified by a security gateway ICMP message.

The symptom of either of these situations would be that a network sniff indicates the security gateway is sending a fragmentation-needed ICMP message, but the traffic initiator is retransmitting the original packet.

To configure the MTU Path Discovery:

1. From the *Device>Contents* column, select the security gateway you want to configure.

2. Click the **Advanced** tab to bring it to the front.

3. From the *Properties* column, select **MTU Path Discovery** to display the *MTU Path Discovery* values.

4. From the Values list, do the following.
    - Select the **On** radio button to run MTU Path Discovery.
    - Select the **Off** radio button to *disable* MTU Path Discovery.

5. When finished, click **Save**.

6. When you want to send the configuration to one or more VSUs, click **Update Devices**.

NAT Traversal

Configurable NAT traversal is available for VPNos 4.31 and later.

Note: For VPNos 3.2, NAT Traversal is enabled by default. You cannot change or disable it.

When a NAT device exists in a network path between security gateways that are part of a VPN, NAT Traversal allows the VPN traffic to successfully pass from one device to another. The default is NAT traversal is enabled.

You can do the following:

• Disable NAT traversal. Avaya recommends that you do not disable NAT traversal even if a NAT device does not exist in the network path of two VPNs.

• Set the value for KeepAlive. The time configured here is used when the security gateway is in the private network of a NAT device. The security gateway behind the NAT device sends a keep alive packet to reserve the dynamic source port. The default is 20 seconds.

Because NAT devices can clear port assignments after a period of inactivity, a still open VPN session may be broken. When a new packet arrives after a certain period of inactivity, a NAT device can assign a new dynamic source port for the packet which causes the VPN connection to fail. To avoid this problem, *keep alive* packets are sent from the VPN peer which is behind the NAT device.

Port for dyna-policy download

If a VSU is configured to receive *dyna-policies* from a remote server instead of storing them locally, it uses a specific port for listening to the remote server. The port uses the *Secure Sockets Layer (SSL)* for protection, and its default number is 1443. The port number can be changed if necessary.

To change the port number:

1. From the *Device>Contents* column, select the VSU you want to configure.

2. Click the **Advanced** tab to bring it to the front.

3. From the *Properties* column, select **Port for Dyna Policy Download** to display the *SSL Port* text box.

4. In the **SSL Port** text box, type in a port number.

5. Click **Save**.

6. When you want to send the configuration to the VSU, click **Update Devices**.

Port for Secure Authentication

Text field for the port number on which the VSU listens for a response from a VPNremote client (over an SSL connection) after the client has been issued an authentication challenge (default port = 2444). A response received on this port is then forwarded to the external LDAP or RADIUS server for authentication.

Private IP Address

A VSU may have two IP addresses assigned to it. The private IP address is used and ARP is set to "Bind one IP address to each port", it is applied to the private port of the VSU, and the public address is applied to the public port. If you specified a private IP address during the VSU Console Quick Setup and the VPNmanager VSU Setup wizard, this address should match that address.

A VSU does not need a private IP address to operate, but some networks may require that a VSU use two addresses. For example, the VPNmanager Console may be running on a machine that is on the private side of the VSU (having a single address). VPNmanager Console-to-VSU communication then has to be routed to the public port of the VSU, which may not be a direct path. The direct path would be to the private port.

A typical use of the private IP address is when the VSU's private side IP network is a different network (different network number and/or mask) from the VSU's public side IP network. For example, when you deploy the VSU in parallel with a firewall or other access device.

If you are using the VSU's primary IP address as the management IP address, use caution when changing it from the VPNmanager. Modifying the private IP address when it is used as the management IP address may cause loss of connectivity between the VSU and the VPNmanager.

Note: The VSU's private (and public) IP address may be used as a gateway IP address for VPN traffic.

To add a private IP address:

1. From the *Device>Contents* column, select the VSU you want to configure.

2. Click the **Advanced** tab to bring it to the front.

3. From the *Properties* column, select **Private IP Address** to display the address controls.

4. Select the **Enable Private IP Address** check box.

5. In the **Private IP Address** text boxes, type in the second address assigned to the VSU.

6. In the **Private IP Mask** text boxes, type in a subnet mask for the address.

7. Select the **Use this address when directly communicating with this device** check box if you want the VPNmanager Console to use this address for communicating with the VSU.

8. Click **Save**, or if you want to send the configuration to the VSU, click **Update Devices**.

Send Device Names

Send VSU Names is an advanced control for managing how remote clients get their *Dyna-Policies.* The Dyna-Policy can be stored locally on one or more VSUs, the Directory Server, or a RADIUS Server. If the policies are stored locally on VSUs, the VSUs in the domain must identify themselves to each other so they can share their database of Dyna-Policies.

To select a VSU name distribution method:

1. From the Device>*Contents* column, select the VSU you want to configure.

2. Click the **Advanced** tab to bring it to the front.

3. From the *Properties* column, select **Send VSU Names** to display the sending options.

4. Select the one of the options.

   - Send all VSU names. Select this option so each VSU in the domain identifies themselves to other VSUs. Use this option if one or more VSUs are storing Dyna-Policies locally.

   - Send VSU(s) names that are involved in CCD only. Select this option if you want the remote client to query only those VSUs that are performing Dyna-Policy services. This is useful if a domain contains many VSUs that are not used for authenticating remote clients. This saves time for the remote client because they don't have to query every VSU to build a complete Dyna-Policy.

   - Send no VSU names. Select this option if a Directory Server or RADIUS Server is used for storing Dyna-Policies. No VSUs are use for locally storing the polices.

   - Customize. Select this option if you wish to specify individual VSU names to be sent.

5. When finished, click **Save**.

6. When you want to send the configuration to one or more VSUs, click **Update Details**.

SuperUser Password (VPNos 3.x)

This function allows you to disable the SuperUser password allowing only LDAP-based communication in the future. Normally used in conjunction with role-based management.

This feature consists of two options for authenticating into a VSU to perform configuration changes:

- VSU/Advanced/SuperUser Password ON (default)

- VSU/Advanced/SuperUser Password OFF

Advanced/SuperUser Password ON (default) - both SuperUser and LDAP authentication are allowed. The VSU attempts to authenticate VPNmanager via SuperUser account first. If this fails the VSU then attempts to authenticate via the VPNmanager user's LDAP account. A successful connection requires that the VSU's authorization provider be set to LDAP user or SuperUser/LDAPuser (default).

When a new configuration is downloaded to the VSU, the VSU authorization provider is reset to SuperUser/LDAPuser, regardless of the previous setting. The next time VPNmanager attempts to connect it may use either SuperUser account or the VPNmanager user's LDAP account.

Advanced/SuperUser Password OFF - only LDAP authentication is allowed. The VSU only attempts to authenticate VPNmanager via the user's LDAP account. A successful connection requires that the VSU authorization provider be set to LDAPuser or SuperUser/LDAPuser (default). When a new configuration is downloaded to the VSU, the VSU authorization provider is reset to LDAPuser, no matter the previous setting. The next time VPNmanager attempts to connect it must use the VPNmanager user's LDAP account.

If VPNmanager has been incorrectly set with VSU/Advanced/SuperUser Password OFF and no LDAP server/user account is configured or available, you must access the VSU console and reset the authorization provider. Before re-attempting to connect, the VPNmanager must set VSU/Advanced/SuperUser Password back to ON, or only a single connection is authenticated, and with SuperUser password left in the OFF position, the VSU only allows LDAP authentication on the next attempt.

Note: The VSU determines what type of authentication it permits, but this is dependent upon the authentication policy last downloaded from VPNmanager (SuperUser Password OFF or ON). Remember that if you set the SuperUser Password to OFF you are no longer able to connect to the VSU using the SuperUser account. The only way to recover SuperUser authentication is to change the setting to back to ON, then do one of the following:

1. Authenticate via your LDAP user account or

2. Go to the VSU console and reset the Configuration/VPNmanager Authorization/Authorization Provider value to SuperUser/LDAPuser, then authenticate by either your LDAPuser account or SuperUser account.

Tunnel Persistence (VPNos 3.x)

This feature consists of the following radio buttons:

• Maintain VPN tunnels on VSUupdate

• Rebuild all VPN tunnels on VSU update

In a multiple VPN structure with tunnel persistence set to *Maintain VPN tunnels*, traffic is interrupted within the modified VPN only. In a multiple VPN structure with tunnel persistence set to *Rebuild All VPN tunnels*, all VPNs related to the modified VPN are interrupted until the routing tables are rebuild by each VSU.

Figure 63, illustrates tunnel persistence between VSUs. The addition of VSU_D to VPN_2 interrupts tunnel persistence in VPN_2 only. Because no modifications have been made in VPN_1 (VSU_A and VSU_B), tunnels remain persistent.

The addition of VSU_D creates VPN_3 (VSU_B and VSU_D). The creation of VPN_3 does not interrupt tunnel persistence on VPN_1. VSU_D is not included in VPN_1 (VSU_A and VSU_B), tunnels remain persistent.

Figure 63. VSU Tunnel Persistence

VPN1 = VSU_A & VSU_B VPN2 = VSU_A , VSU_C , VSU_D VPN3 = VSU_B & VSU_D
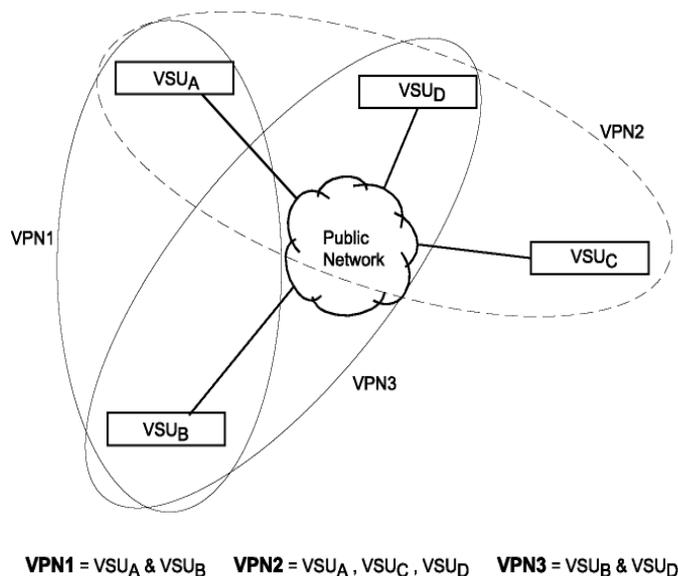
Figure 64, illustrates tunnel persistence between VSUs and remote users (RUser). The addition of VSU_D to VPN_2 (VSU_A, VSU_C, VSU_D, and Remote User) interrupts tunnel persistence in VPN_2, thus breaking the remote connection. Once the VSUs in VPN_2 rebuild their routing tables, the remote connection can be restored by the remote user. Because no modifications have been made in VPN1 (VSU_A and VSU_B), tunnels remain persistent.

The addition of VSU_D creates VPN3 (VSU_B and VSU_D). The creation of VPN3 does not interrupt tunnel persistence on VPN1. VSU_D is not included in VPN1 (VSU_A and VSU_B), tunnels remain persistent.

Figure 64. Remote User Tunnel Persistence

VPN1 = VSU_A & VSU_B **VPN2** = VSU_A , VSU_C , VSU_D , RUser   **VPN3** = VSU_B & VSU_D
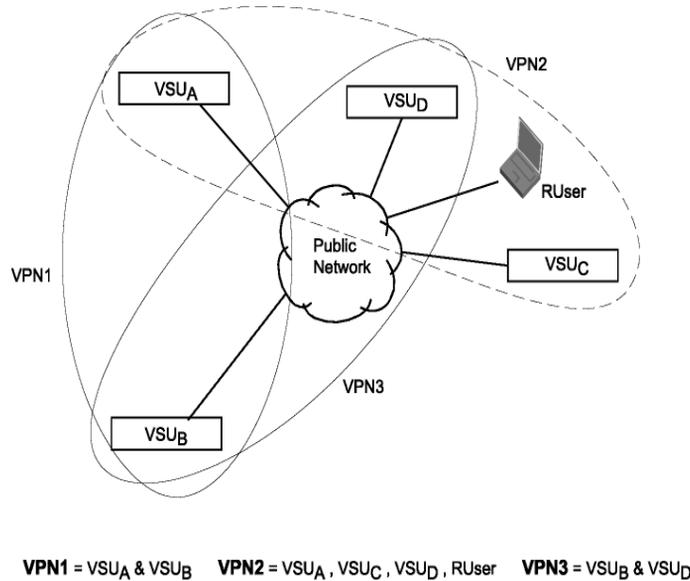
TEP Policy

The Tunnel End Point (TEP) Policy tab provides control of the security policy applied to the traffic that flows between the end points of a tunnel. The default is off, or Do not apply configured VPN policies to TEP traffic.

Figure 65. Tunnel End Point Policy
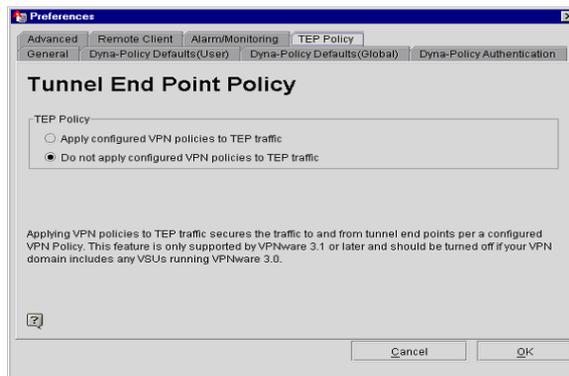
Enabling apply configured VPN policies to TEP traffic encrypts the traffic destined to and from tunnel end points when the following conditions are met:

• Primary IP address of VSUs in your VPN domain must be included in the IP group they are protecting.

• SKIP tunnel mode or IKE is being used (SKIP Transport mode NOT being used).

Failing to meet these conditions, packets be subject to the non-VPN traffic policy (Permit or Deny) selected in the VSU Packet Filtering/Advanced tab.

A typical example of when enabling Apply configured VPN policies to TEP traffic is desired is in the situation of remotely reading an Active Sessions MIB object of a VSU. The information returned here includes the user name or IP address for each session currently active on the selected VSU. Obviously, having this SNMP information pass over the internet in the clear is not desirable.

This feature is not supported in releases of VPNmanager prior to 3.1. Because both tunnel end points must have Apply configured VPN policies to TEP traffic enabled, the VSUs on each end must also be running VPN NOS 3.1 or later.

Servers tab

The Servers tab is used for adding backup *directory servers* to a specific security gateway. There is no practical limit on how many backups you can configure. Backup servers can be added at anytime, and they can be organized so that when one fails, a specific one can be used as a backup.

To install additional servers, see your iPlanet Directory Server documentation for instructions. The following procedure only establishes it as a backup server. The Directory Servers tab is shown in Figure 66.

Figure 66. The Directory Servers tab
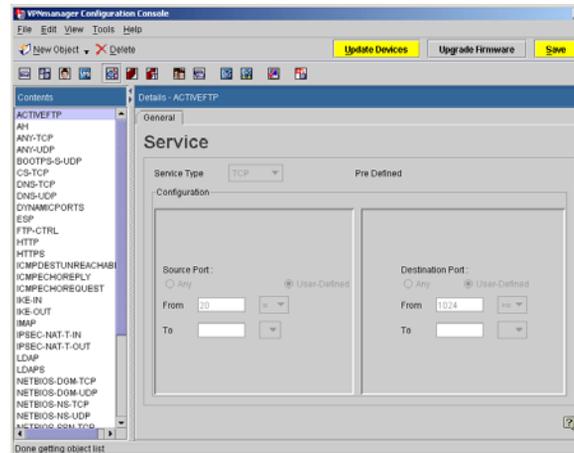
Servers

Servers list presents a list of available directory servers. Three columns appear which include IP address or DNS Name, port, and SSL state.

Move Up/Down arrows are provided to change the position of the highlighted server.

Edit/Delete/Add buttons are provided at the bottom of the pane.

Add servers

Brings up a dialog box to add additional servers. Enter the new server's IP address or DNS Name. The Locate This Server box contains three radio buttons used to place the new server:

- Beginning of List

- End of List (default)

- After Selected Item

To create a backup server:

1. Move to the *Configuration Console* window.

2. From the **Device>Contents** column, select the security gateway that needs to have the backup server.

3. Click the **Directory Servers** tab to bring it to the front.

4. Click **Add** to open the *Add Directory Server* dialog box.

5. Use Table 15 configuring a connection to a server.

Table 15. Add Directory Server Commands

Item	Description
Enter IP Address or DNS Name	Type in the IP address or host name used by the server.
Locate This Server	Use these options to insert the server into a specific position in the *Directory Servers* list.
Port	Type in the port number of the server (default is 389). To verify the number, move to the computer running *iPlanet Directory Server*, then start the iPlanet Console; the number can be seen from the *Console* tab.
Use SSL	Select this check box to protect the communication between the VPNmanager Console and the Directory Server with a *Secure Socket Layer (SSL)*. Read Appendix A, Using SSL with Directory Server, before making this selection.

6. Click **OK** to return to the Directory Servers tab. The new backup server appears in the *Directory Servers* list.

7. When finished, click **Save** to save your work.

Managing the server list

The backup servers shown in the *Servers* list can be edited, have their sequence changed, or even deleted. The list organizes the servers in the sequence in which they must be used, where the one at the top of the list is always used first.

To edit, change the sequence, or delete a backup server:

1. Move to the *Configuration Console* window.

2. From the **Device>Contents** column, select the security gateway that has the backup server that needs to be changed.

3. Click the **Servers** tab to bring it to the front.

4. From the *Servers List*, select a specific secondary end-point.

5. Use Table 16 for performing specific management tasks.

Table 16. Servers list commands

Command	Description
Edit	Use this command to edit the server with the *Add Directory Server* dialog box.
Move Up	Click this button to move the server higher in the list.
Move Down	Click this button to move the server lower in the list.
Delete	Click this button to remove the server from the list.

When finished, click **Save** to save your work.

Resilient Tunnel tab

Tunnels are used to protect VPN traffic that moves through the public networks. The endpoints for tunnels are located in VSUs. *Resilient Tunnels* are used for backing up a specific primary tunnel. Up to three resilient tunnels can be created to backup a specific security gateway. VSUs can report tunnel switching to a common SNMP manager (See "Using SNMP to monitor the device" on page 277.).

Figure 67 illustrates a simple example. San Francisco LAN has two gateways to the WAN. The high-speed route is used by the primary tunnel, and the low-speed route is used by the resilient tunnel. If the circuit in which VSU_B is located goes out of service, traffic automatically switches to VSU_C. Once VSU_B is back in-service, VPN traffic then switches to the primary tunnel. The switching is controlled by VSU_A which is located in the Tokyo LAN.

Figure 67. Primary and Resilient Tunnels

Resilient Tunnels are used for backing-up Primary Tunnels.
Should a Primary Tunnel go out of service, the Resilient Tunnel
will automatically be used for VPN traffic.
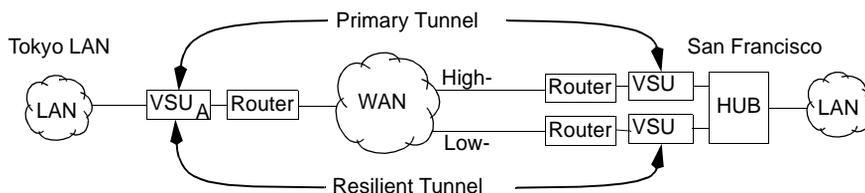
Tunnel Switching

The switching mechanism involves *time* and a packet called a *Heartbeat*. Figure 68 illustrates how tunnels are switched.

Figure 68. Tunnel Switching
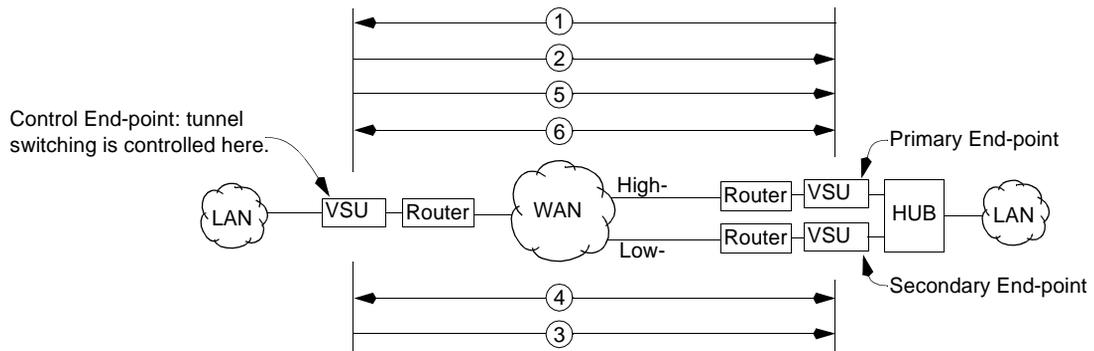
Explanation for Figure 68

1. VSU_A listens to VSU_B's heartbeat. The heartbeat has a configurable period called a *Heartbeat Interval*.

2. If VSU_A realizes a dead heartbeat, it asks VSU_B for a heartbeat.

   The number of times that VSU_A can make a request is configurable, and is called the *Heartbeat Retry Limit*.

3. If the number of requests exceeds the *Heartbeat Retry Limit*, VSU_A then begins to establish a connection with VSU_C.

   Since VSU_C uses a low-speed connection, VSU_A must anticipate a delayed response from VSU_C. That delay is called *Hold-up Time*, and is configurable with VPNmanager Console.

4. After VSU_A establishes a connection with VSU_C, the resilient tunnel is used for VPN traffic.

5. On a *periodic* basis, VSU_A continues to request a heartbeat from VSU_B. The period is called *Dead Primary Poll Interval*.

6. If VSU_A reconnects with VSU_B, VSU_A waits for a specific time before it switches traffic back to VSU_B. The waiting period is called *Hold-down Time*.

Note: If packet filtering is used, be sure the heartbeat packets are not filtered. The security gateway heartbeat listening port = 1643 using UDP protocol.

Creating a resilient tunnel

Resilient tunnels are configured from the Resilient Tunnel tab.

Figure 69. The Resilient Tunnel tab for a security gateway Object
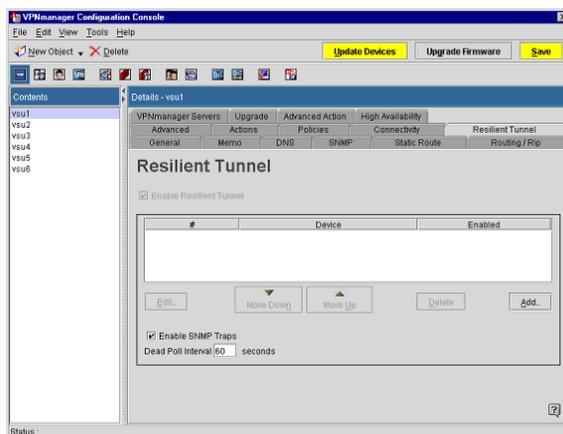
- • Enable SNMP Traps. Check this box if you want SNMP traps enabled for the resilient tunnel.

- • Dead Poll Intervals. The number of seconds between heartbeat poll requests to a dead primary. This is different from the normal Heartbeat Interval because the primary security gateway is believed to be inactive and no response is expected. Therefore, the interval is much longer than a normal heartbeat request interval.

Add resilient tunnel

There are four parameters associated with Resilient Tunnel automatic backup mode. They are:

- • Heartbeat Interval

  The time, in seconds, between heartbeat request attempts made by the remote security gateway to the primary security gateway. Default is 10 seconds.

- Heartbeat Retry Limit

  The number of times a heartbeat request is sent by the remote security gateway before the primary security gateway is declared inactive. Default is 3 tries.

- Hold Up Time

  The time (in seconds) to wait before the remote security gateway attempts to contact the secondary tunnel endpoint security gateway. This allows for the latency of a dialup link, typically much longer than the heartbeat interval. Default is 0.

- Hold Down Time

  Wait time between the remote security gateway determining that the primary endpoint security gateway is able to reconnect, and when the switchover actually occurs. This wait time ensures that the primary security gateway is stable before switching occurs. Default is 20 seconds.

Prerequisites

- Security gateway for the controlling, primary, and secondary end-points must exist. For instructions, see the "Configuring a security gateway" on page 63.

- A VPN Object that uses the controlling and primary security gateway objects must exist. For instructions see the "Creating a new VPN object" on page 156.

To create a resilient tunnel:

1. Move to the *Configuration Console* window. The Device tabs are displayed.

2. From the **Device>Contents** column, select the device that is operating as the primary end-point (see Figure 68).

3. Click the **Resilient Tunnel** tab to bring it to the front.

4. Click **Add** to open the *Add Resilient Tunnel Device* dialog box.

5. From the **Select a Device** list, select the security gateway that is the secondary end-point.

6. Select the **Save as Enabled** check box so Resilient Tunnel services begins as soon as the VSUs are updated.

7. From the **Properties** list, click on **Heartbeat Interval** so the *heartbeat interval* values appears.

   • In the **Heartbeat Interval** drop-down list, select a unit of time.

   • In the **Heartbeat Interval** text box, type in a duration that defines the period of the primary end-point's heartbeat.

8. From the **Properties** list, click on **Heartbeat Retry Limit** so the *heartbeat retry limit* values appears.

   • In the **Times** text box, type in the number of times a heartbeat must be requested by the controlling end-point before it switches traffic to the secondary end-point.

9. If the secondary end-point uses a slower circuit than the primary end-point, the controlling end-point must be aware of the expected delay. That delay is called *Hold-Up Time*.

   From the **Properties** list, click on **Hold Up Time** so the *Hold Up Time* values appears.

   • In the **Hold-Up Time** drop-down list, select a unit of time.

   • In the **Hold-Up Time** text box, type in a duration that the controlling end-point may have to wait for a response from the secondary end-point.

10. From the **Properties** list, click on **Hold-Down Time** so the *hold-down time* values appears.

    • In the **Hold-Down Time** drop-down list, select a unit of time.

    • In the **Hold-Down Time** text box, type in a duration that the controlling end-point must wait before it switches VPN traffic from the secondary end-point to the primary end-point. The wait begins after the controlling end-point reconnects with the primary end-point.

11. Click **OK** to return to the Resilient Tunnel tab. Your new secondary end-point appears in the *Resilient Tunnel* list.

12. Click **Save** to save your work.

Managing the resilient tunnel list

The secondary end-points shown in the *Resilient Tunnel List* can be edited, have their sequence changed, or even deleted. The list organizes the secondary end-points in the sequence in which they must be used, where the one at the top of the list is always used first.

To edit, change the sequence, or delete a filtering policy:

1. Move to the *Configuration Console* window.

2. From the **Device>Contents** column, select the security gateway that acts as the primary end-point for a tunnel.

3. Click the **Resilient Tunnel** tab to bring it to the front.

4. From the *Resilient Tunnel List*, select a specific secondary end-point.

5. You can edit, move up, move down or delete.

6. When finished, click **Save** to save your work.

Stopping and starting resilient tunnel services

Resilient tunnel services for a specific *primary end-point* or secondary end-point can be stopped or started at any time.

Primary end-point service

To stop or start resilient tunnel services for a primary end-point:

1. Move to the *Configuration Console* window. Select **Devices.**

2. From the **Device>Contents** column, select the device that acts as the primary end-point for a tunnel.

3. Click the **Resilient Tunnel** tab to bring it to the front.

4. Do one of the following:
   - Select the **Enable Resilient Tunnel** check box to start services.
   - Clear the **Enable Resilient Tunnel** check box to stop services.

5. Click **Save** to save your work.

Secondary end-point service

To stop or start resilient tunnel services for a secondary end-point:

1. Move to the *Configuration Console* window. Select **Devices**.

2. From the **Device>Contents** column, select the security gateway that acts as the secondary end-point for a tunnel.

3. Click the **Resilient Tunnel** tab, to bring it to the front.

4. From the *Resilient Tunnel List*, select a specific secondary end-point.

5. From the **Enabled** column, do one of the following:
   - Select the check box to start services.
   - Clear the check box to stop services.

6. Click **Save** to save your work.

Advanced Action tab

The Device Advanced Action tab provides access to advanced security gateway functions including switching the NOS execution from flash 0 to flash 1 (or back), resetting the security gateway's password, or disabling FIPS on the selected security gateway.

Figure 70. Advanced Action tab
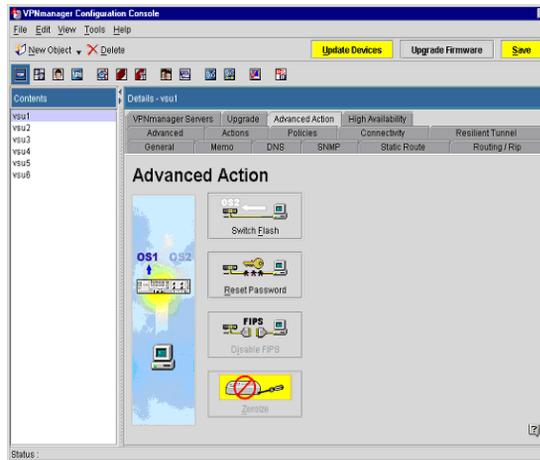
Switch Flash

Switch flash is used to switch the flash chip from which the security gateway is executing its NOS. Normally, a duplicate image of the NOS is loaded into the second flash bank, however, a new or previous NOS image may alternately be loaded when it is desired to switch between the two NOS versions.

The flash from which the security gateway is currently executing its NOS is indicated (Flash 0 or Flash 1). Additional information can be found in the security gateway Data portion of the security gateway General tab.

Reset password

Reset password is used to change the console password on the selected security gateway. An example of when this is used is if you were to forget the security gateway console password, you may change it using this dialog box.

Disable FIPS

This key is used to turn Federal Information Processing Standards (FIPS) mode off. FIPS indicates whether the VSU is running in the normal or FIPS level 2 mode. Avaya recommends that this mode be used only if an organization's policy requires FIPS 140-1 level 2 certification for cryptographic devices.

High Availability tab

This tab provides access to the High Availability (HA) functions for the security gateway including enabling high availability, setting the public and private virtual addresses, adding security gateway members to the HA group, viewing the status of the HA group, converting a passive member to an active member, configuring member VSUs, the VRRP advertisement interval, version number, third party reference points for the public and private interfaces, and minimum connectivity to reference hosts.

This feature provides a fault tolerant infrastructure that minimizes the downtime of the protected network. Fault tolerant infrastructure is achieved by pairing two like VSUs together to form a HA group. The HA group is comprised of a primary or active security gateway and secondary or passive security gateway. Only one instance of the security gateway is

visible in the security gateway contents list. The active security gateway is listed with the passive security gateway visible in the Members pane of the High Availability tab. Because configuration within the HA group is identical, only the primary security gateway of the HA group is displayed.

Figure 71. High Availability
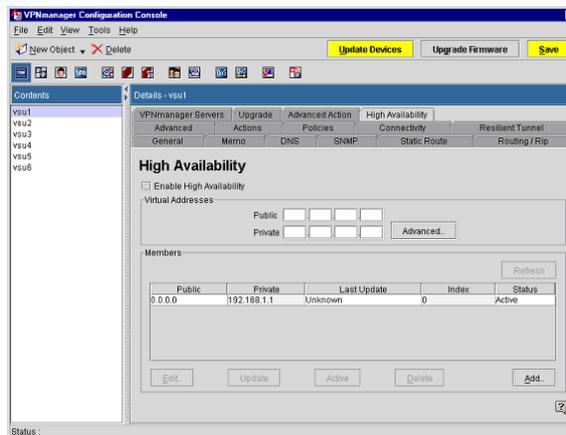
Enabling high availability

This check box is used to enable high availability (HA). Use this check box after you have confirmed the public and private VSUs in the HA group have been configured to deny all non-VPN traffic. Beginning with VPNmanager 3.2, Deny all Non-VPN Traffic is the default selection.

For additional information about configuring the security gateway to deny all non-VPN traffic through the security gateway console, refer to *Preparing the security gateway for Configuration*, of the security gateway User's Guide.

For additional information on how to configure the security gateway to deny all non-VPN traffic through the initial security gateway Quick Setup, refer to the *"Configuring a security gateway" on page 63*.

To configure the security gateway to deny all non-VPN traffic through the VPNmanager:

1. Move to the *Configuration Console* window. Select **Devices**.

2. From the *Device>Contents* column, select the security gateway you want to configure.

3. Click the **Policies** tab to bring it to the front.

4. From the drop-down list, select **Packet Filtering**, then click **GO** to open the *Policy Manager* for *Packet Filtering*.

5. Click **Advanced** to display the Packet Filter Rule Advanced window.

6. Select the **Deny all non VPN traffic** radio button.

7. Click **OK**.

8. Click **Save**.

9. From the upper right-hand of the window, click the close button to return to the *Configuration Console* window.

10. When you want to send the configuration to the security gateway, click **Update Devices**.

Virtual addresses

Once you enabled High Availability by selecting the check box, configure the public and private Virtual Addresses.

The configured Virtual Addresses are shared among all members in the HA group. The public Virtual Address is used as the tunnel end point while the private Virtual Address can be used as the default route for the network behind the security gateway. Configuring the Virtual Addresses in this manner ensures that any member in the HA group has the same configuration and that this configuration does not change.

Advanced parameters

The Advanced Parameters are displayed by clicking the Advanced button.

Once configured, the Advanced Parameters are common to all members in the HA group.

Advertisement interval. The time interval the passive member must detect before it becomes the active member in the HA group. The passive member must detect the elapsed time interval three times before it forces the election to become the active member. The Advertisement Interval range is 1 to 255 seconds.

Group ID. The Group ID allows configuration of a unique identifier for the HA group. By using the Group ID, the HA group avoids conflicts with other VRRP implementations on the network. The values for the Group ID can range from 0 to 255.

Third point of reference hosts. If the network requirements do not permit having the private interface and the public interface plugged into the same network device, configure a Third Point of Reference Hosts (TPRH).

In this network configuration and before a passive member can become active, the passive member must be able to ping the TPRHs on both the private and public interfaces. The TPRH connectivity is configurable from the High Availability Advance dialog box. One TPRH must be configured and up to 8 hosts can be configured for each interface.

Members

The Members table displays all configured members in the HA group. By default, the primary member displays an active status while the secondary and remaining members display a passive status.

The Member table also displays the primary, secondary, last update, current config, and status of each member in the HA group.

- **Refresh** - Displays the current status of each member of the HA group.

- **Public** - Displays the public IP address of the HA group.

- **Private** - Displays the private IP address of the HA group.

- **Last Update** - Displays the date and time stamp of the configuration update of the HA group. VPNmanager handles the different time zones for you. In whatever time zone the update configuration occurred, VPNmanager always displays the time stamp in the local time of the administrator confirming the last update. For example, if an administrator in New York updated the HA at 1:00 pm EST and administrator in Oregon wanted to confirm the last update. The time

stamp displayed in the member table for the administrator in Oregon would display in PST, not EST. All configuration updates are saved in GMT.

- **Index** - Displays the current configuration revision number of all members. This allows the administrator to confirm that all the members have the same configuration. A successful configuration revision revises the index number. The member with the highest index number is eligible to become the active member in the HA group. If the active member has the highest index number and a passive member is revised to also have the highest index number, the active member maintains the active status. The passive member that has been updated to the highest index number does not replace the active member.

- **Status** - Displays the active or passive status of the devices in the HA group.

By selecting the member in the table, the following actions can be performed:

- **Edit** - This action allows the member to be edited.

- **Update** - This action allows the selected member configuration to be updated. If you suspect that a passive member does not have the most current configuration for the HA group, use the Update button to update the passive member's configuration. Using Update revises the configuration on the passive member to match the active member's index number.

- **Active** - This action allows a passive member to become active. A trap is generated when there is a change in status in the HA group. VPNmanager is notified through the trap that a change in status has occurred and updates the Member table accordingly.

- **Delete** - This action allows the member to be deleted from the HA group. VPNmanager notifies the member that it is no longer part of the HA group.

- **Add** - This action allows a new member to be added to the HA group. The minimum configuration of a new member is the public and private IP addresses. By default, the primary IP address is used as the management address when communicating to the member.

Configuring high availability

Creating a High Availability Group

Use the following procedure to create High Availability (HA) groups:

1. Create a new security gateway Object that includes in the HA group. For additional information on creating a new security gateway object, see "Configuring a security gateway" on page 63.

   Note: Because configuration within the HA group is identical, only the primary security gateway of the HA group is displayed.

2. After the security gateway is created, select the security gateway from the **Device>Contents** column.

3. Click the **High Availability tab** to bring it to the front.

4. Click the **Enable High Availability** check box to enable High Availability on the security gateway.

5. Enter the **Virtual Addresses** for the public and the private interfaces. Configuring the virtual addresses in this manner ensures that any member in the HA group has the same configuration.

   Note: Virtual Addresses must be valid routable addresses.

6. Click the **Add** button to add members to the HA group.

7. Enter the **private IP addresses** of the Active security gateway.

   The private IP address may have been entered during the initial creation of the security gateway object. If the private IP address has already been entered, confirm the IP address is correct and move to the next step.

8. Enter the **public and private IP addresses** of the Passive security gateway(s).

9. Click the **Update security gateway** button to update the HA configuration.

Updating a high availability group using Update Device

High Availability groups can be updated using the Update security gateway button in the VPNmanager Configuration Console window.

When using the Update Device, VPNmanager displays the selected security gateway to be updated. If the selected security gateway is a HA member, the Member Update screen displays. By default, all members in the HA group are selected for update.

To update HA VSUs:

1. Move to the *Configuration Console* window. Select **Device**.

2. Select the security gateway to be updated.

3. Click the **High Availability tab** to bring it to the front.

4. Click **Update security gateway** from the *Configuration Console*.

5. The *Member Update* window appears. By default, all members in the HA group that are part of the site-to-site VPN configuration are selected for update.

6. Click **OK** to complete update.

Deleting a high availability group

Use the following procedure to delete High Availability (HA) groups:

1. Click the **High Availability tab** to bring it to the front.

2. Click the **Refresh** button in the Members section of the screen to refresh the status of the HA groups members.

3. From the Members section, select the security gateway to be deleted.

4. Click the **Delete** button to delete the security gateway from the HA group.

5. Click the **Enable High Availability** check box to disable High Availability on the remaining security gateway.

6. Click **Update Devices** from the *Configuration Console*.

Click **OK** to complete update.

Failover

Use the Failover object to configure up to five IP addresses for tunnel endpoint (TEP) for the security gateways. These IP addresses are used for failover locations in the case of VPN or clear traffic failure.

Figure 72. Failover Tab
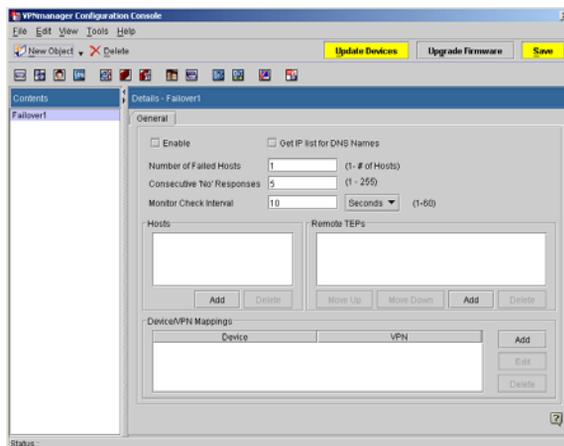
When Failover is configured, a security gateway periodically checks connectivity to designated devices to evaluate the availability of the network path to the central-site resources. These devices can be within the VPN, such as the corporate e-mail server at the central site. These devices can also be outside the VPN, such as a public DNS server.

When a network path fails, the remote security gateway tries to establish a network path through an alternate central-site. If the remote security gateway cannot use that second central-site TEP to establish a network path, the remote security gateway continues through the list of configured TEPs, and tries to establish a usable network path to the central-site resources. If none of the configured tunnels can establish a network path, and the remote security gateway is configured with a public-backup interface, the remote device tries to establish a path through this alternate link.

When the public-backup zone is in use, the security gateway does not perform failover connectivity-checks to the designated hosts. When the idle timer is enabled, and as long as there is traffic, this alternate network link is used. If the configured idle time elapses, the public-backup interface is taken down. The security gateway then tries to reestablish the network connectivity through the primary network path.

Note: If the public-backup interface idle timer is disabled, the security gateway continues to use the alternate network interface.

Network path failure is defined as the configured number of consecutive connectivity checks without a response from the number of hosts that need to fail. The following is an example of a network path failure criteria .

The configuration is as follows:

- The number of consecutive "no" responses is five.

- The idle time between each connectivity check is 10 seconds

- The number of hosts to monitor is three.

- The number of hosts that must fail to respond, out of the hosts configured is two.

Table shows which hosts respond (Y) and which hosts do not respond (N) during the 10-second interval connectivity check.

Table 17. Failover connectivity checks in 10-second intervals

	10	20	30	40	50	60	70	80	90	100	110	120	130
Host													
1	Y	Y	Y	N	N	N	Y	Y	Y	Y	Y	Y	Y
2	Y	Y	Y	Y	N	N	N	N	N	N	N	N	N
3	N	N	N	N	N	Y	Y	Y	N	N	N	N	N

The network path failure criteria are met only when *both* hosts 2 and 3 *concurrently* fail to respond five times (at the 130 second mark) to the connectivity checks. Host 3 failed to respond five consecutive times (between the 10-second interval and the 50-second interval). Host 2 failed to respond five consecutive times (between the 50-second interval and the 90 second interval). But only when host 2 and host 3 both fail to respond to the same five consecutive security checks are the failure criteria met.

To configure failover

1. From the VPNmanager Console main window, select **Failover** as a New Object. The Failover tab appears.

2. From the **Failover>Contents** column select the device to configure for Failover.

3. Select **Enable** to provide an alternate network path to re-establish access to the central-site resources.

4. Select **Get IP List for DN** so that when a DNS query is made, the security gateway keeps all the IP addresses that are returned in the cache. The security gateway attempts to respond to the queries in the same order that the queries were received.

5. If this parameter is not selected and a DNS query is made, the security gateway uses the first IP address of the DNS response that is returned.

6. In the **VPN Name** field, enter the name of the VPN that you want to monitor. This field is optional, if Set VPN Mode is set to dynamic.

7. In the **Number of Failed Hosts** field, enter the number of configured hosts that can fail before network path failover criteria is reached. If multiple hosts are configured and all hosts are critical, enter 1. If any one of the configured hosts failed to respond, network path failover occurs.

8. In the **Consecutive "No"Responses** field, enter the number of consecutive connectivity checks without a response that you want to allow. The default is 10.

9. In the **Monitor Check Interval** field, Enter the number of seconds that you want to allow between connectivity checks to the configured host or hosts. The interval is also used to define the response time of the host. Monitor checks are made at the same time to each host. The default is 10 seconds.

10. In the **Hosts** field, click **Add**, to enter the network host or hosts for which you want to monitor connectivity. You can define up to five DNS names or IP addresses. These hosts can be either within the VPN or

outside the VPN. If the host is within the VPN, the host information is encapsulated in the associated VPN policy. If the host is outside the VPN, the host information is sent in the clear.

11. In the **Remote TEP field**, click **Add**, to enter the tunnel endpoints (TEP) for the central site that the remote VPN device establishes a network connection. If the network path failure criteria is met while the remote security gateway is trying to establish a network connection, the remote VPN tries to alternate TEPs until a network connection is made.

12. Click **Save.**

Policy Manager - My Certificates

If you are creating VPNs that use certificates for authentication and security, use the *Policy Manager* for *My Certificates* to install signed certificates into specific VSU.

After one or more certificates have been installed, see IKE Certificate Usage on page 272 about configuring a target for a signed certificate, and Issuer certificates on page 269 about installing issuer certificates on a target.

About VSU certificates

VSUs use *public-key* certificates based on CCITT Recommendation X.509. Within the framework of the recommendation, each certificate includes a *Rivest, Shamir, and Adleman (RSA) Public-Key Cryptography Standard (PKCS) Number 10* for authentication.

A VSU can store up to nine certificates. One is a *default certificate* which is only used for the SSL connection between the VSU and the VPNmanager Console. The remaining eight certificates are *My Certificates* and are statically stored in the flash memory of the VSU. The default certificate is issued by Avaya Inc..

Note: The default certificate has a six year period of validity, which starts at the factory when it's put into the VSU. Reprogramming the flash is the only way to change the default certificate.

Up to eight certificates can be stored in a VSU. During IKE negotiation, a VSU sends a specified certificate to its target. Those other VSUs and clients are called targets. Likewise, the target that received a certificate must distribute its [unique] certificate to the sender to complete the exchange. The VSUs use the exchange to authenticate each other and to distribute their public keys. These additional certificates can be created then installed into a VSU. Each certificate is assigned a target (see for additional information about making those assignments). A VSU only needs a single certificate to distribute its public-key to multiple VSUs, but additional certificates can be created for establishing secure connections with special targets. The process of getting a certificate for a specific VSU is illustrated in the following figure.

Figure 73. Installing a Signed Certificate into a VSU
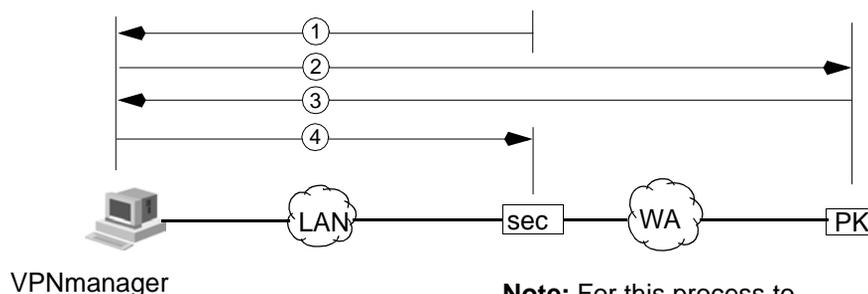
VPNmanager

Note: For this process to work, the security gateway must have already been

Explanation for Figure 73:

1. An administrator uses VPNmanager Console to get a *Certificate Request* from a specific VSU.

2. The administrator sends the *Certificate Request* to a *Public Key Infrastructure (PKI) System*.

3. The *PKI* System sends a *Signed Certificate* to the administrator.

4. The administrator uses VPNmanager Console to install the *Signed Certificate* into the VSU.

Creating and Installing a Signed Certificate

Shown in Figure 74 is the Policy Manager for My Certificates. Use it for generating certificate requests, installing signed certificates in a VSU, and for selecting which certificate the VPNmanager Console must be configured as the target.

Figure 74. The Policy Manager for My Certificates
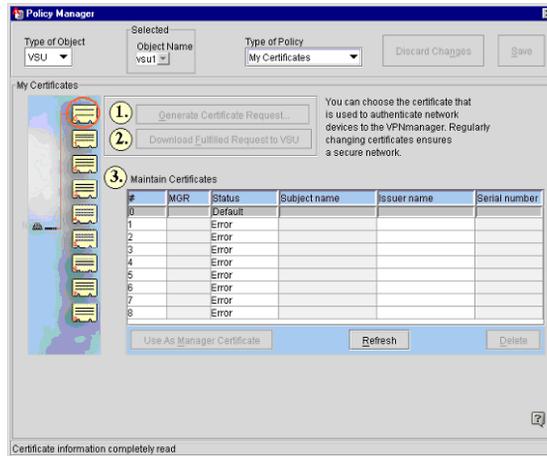
To install a signed certificate into a VSU:

1. From the *Device>Contents* column, select the VSU that needs a Signed Certificate.

2. Click the **Policies** tab to bring it to the front.

3. From the *drop-down list*, select **My Certificates**, then click **GO** to open the *Policy Manager* for *My Certificates*.

4. Click **Generate Certificate Request** to open the *Save as* dialog box.

5. Use the **Look in** drop-down list to navigate to a directory where you want to save the certificate request.

6. In the **File name** text box, type in a name for the *Certificate Request*, then click **Save***.*

The VSU saves a *Certificate Request* into this new file then update the *Maintain Certificates* list with information about the new *Certificate Request*. The status column for the *unsigned* request displays *Request Ready.* The request exists in the Privacy-Enhanced Mail (PEM) using PKCS #10 format.

7. Send the *Certificate Request* to a PKI System.

The PKI System must use the *Distinguishing Encoding Rules (DER)* format for creating the *Signed Certificate*.

The PKI System creates a *Signed Certificate* for the VSU. Figure 75 shows what a certificate in PEM format looks like (its body has been shortened for the example). Currently a VSU accepts the certificate delivery formats of PEM, DER, Base64X509, and PKCS#7.

Figure 75. An Example of a Signed Certificate

Header

-----BEGIN CERTIFICATE-----

nfi897rho987fb+mht>,oi$s25hgj98iJop)kjh

GrDfgyui987jg55dJ99KJY6%$3@@Sd5()~

43dbi0oMl=_+;mhjuuhJ8*&tfeEckiooplkjghf

Footer

8. Cut the signed certificate from whatever file the PKI System sent it in, then paste it to the file you created in Step 6. Include the header and footer.

Note: The alignment of the right side of the certificate must be even (justified), so if the certificate was sent to you in a web page where the last line may run past the right side, just place a carriage return in the appropriate place of the line to even it up. Also, place a carriage return at the end of the footer line.

9. Return to the *Policy Manager* for *My Certificates* for the specific VSU.

10. From the **Maintain Certificates** list, select the item identifying the requested certificate.

11. Click **Download Fulfilled Request to VSU** to open the *Open* dialog box.

12. Use the **Look in** drop-down list for navigating to the location of the signed certificate file. The manager uses DER as the default filename extension, but TXT can be used.

13. Select the signed certificate file, then click **Open** to return to the *Policy Manager* window. After the VSU has received the signed certificate, the *Status* column changes from *Request Ready* to *Cert Accepted*.

Switching certificates used by VPNmanager Console

VPNmanager Console uses the default certificate of the VSU for establishing a secure connection with a VSU. The default certificate can be used until it expires (6 years), or until the VPNmanager Console is made to use a different certificate. The VSU certificate used by VPNmanager Console can be changed anytime.

To switch certificates:

1. From the *Device>Contents* column, select the VSU you want to configure.

2. Click the **Policies** tab to bring it to the front.

3. From the *drop-down list*, select , then click **GO** to open the *Policy Manager* for *My Certificates*.

4. From the **Maintain Certificates** list select the certificate that you want the VPNmanager Console to use.

The default VSU certificate is identified by an asterisk in the *MGR* column. Although a specific certificate may have other targets, as assigned through the *IKE Certificate Usage* tab (), the VPNmanager Console can still use it.

5. Click **Use as Manager Certificate** to make the VPNmanager Console a target of the certificate.

Issuer certificates

Targets use an *Issuer Certificate* to authenticate a *Signed Certificate*. VSU targets can dynamically store up to eight *Issuer Certificates*. Storage on VPNremote Client targets is only limited by the amount of physical memory of the computer. *Issuer Certificates* must be installed on targets before they are needed to authenticate a *Signed Certificate*. This section explains how to retrieve and install *Issuer Certificates* for VSU targets. For information about installing *Issuer Certificates* on VPNremote clients, see the *VPNremote Administrator's Guide.*

About Issuer Certificates

The *Signed Certificates* stored in VSUs are X.509 public-key certificates. They're used for distributing a public-key of the VSU to targets (other VSUs, VPNremote Clients, and IKE compatible clients). Every *Signed Certificate* identifies which *Public Key Infrastructure (PKI) System* has signed it. However, targets must use a method to authenticate every *Signed Certificate* they receive.

Note: An Issuer Certificate may be called a "Signing Certificate" or "Certification Authority (CA) Certificate."

Targets use an *Issuer Certificate* to authenticate a *Signed Certificate*. Therefore, the *Issuer Certificate* must be from the same *PKI System,* as the *Signed Certificate was signed by the issuer's private key.* Figure 76 illustrates how *Issuer Certificates* fit in the scheme of signed certificate exchange.

Figure 76. Issuer Certificates
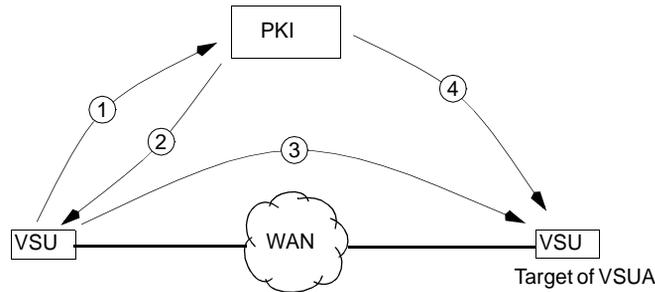
Targets use Issuer Certificates to authenticate
Signed Certificates they receive. The Issuer
Certificate must be from the same PKI System that
created the Signed Certificate. Issuer Certificates
are stored on targets.

Explanation for Figure 76:

1. A *Certificate Request* from VSU_A is sent to a PKI System to be signed.

2. The PKI uses the *Certificate Request* to create a *Signed Certificate* specifically for VSU_A. The *Signed Certificate* is then stored on VSU_A.

3. Every target of VSU_A must have *VSU_A's Signed Certificate*.

4. The target uses an *Issuer Certificate* to authenticate *VSU_A's Signed Certificate*. The *Issuer Certificate* must be from the same PKI which created the *VSU_A's Signed Certificate*.

Installing an issuer certificate

Use the Policy Manager for installing Issuer Certificates in a specific VSU. The VSU then uses the Issuer Certificate to authenticate certificates received from other VSUs.

The process is explained in Figure 76.

To install an Issuer Certificate into a VSU (target):

1. Get an *Issuer Certificate* from a PKI System. Use the same PKI System that created the *Signed Certificate*.

The PKI System must use the *Distinguishing Encoding Rules (DER)* format for creating the *Issuer Certificate*. Figure 77 shows what a certificate looks like (its body has been shortened for the example).

Figure 77. An Example of an Issuer Certificate

Header

-----BEGIN CERTIFICATE-----

nfi897rho987fb+mht>,oi$s25hgj98iJop)kjh

GrDfgyui987jg55dJ99KJY6%$3@@Sd5()~

43dbi0oMl=_+;mhjuuhJ8*&tfeEckiooplkjghf

Footer

2. Cut the issuer certificate from whatever file the PKI system sent it in, then paste it into a text file. The file can have a DER or TXT file name extension.

Note: The alignment of the right side of the certificate must be even (justified), so if the certificate was sent to you in a web page where the last line may run past the right side, just place a carriage return in the appropriate place of the line to even it up. Also, place a carriage return at the end of the footer line.

3. Return to the *Policy Manager* for *Issuer Certificates* for the VSU needing the certificate.

4. Click **Add** to open the *Open* dialog box.

5. Use the **Look in** drop-down list for navigating to the location of the *Issuer Certificate*.

6. Select the *Issuer Certificate*, then click **Open** to return to the *Policy Manager* window.

After the VSU has received the Issuer Certificate, the certificate appears in the *Issuer Certificates* list.

IKE Certificate Usage

If you are creating VPNs which use certificates for authentication and security, use the *Policy Manager* for *IKE Certificate Usage* to configure how *VSU Certificates* must be used. Those certificates were created and installed in VSUs from the *My Certificates* policies (page 264). The *IKE Certificate Usage* policies is the mechanism used for exchanging certificates in a VPN.

About Certificate Usage (Exchange)

Every certificate identifies its owner and contains the owner's public-key. The concept of certificate usage is based on *Owners* and *Targets*. An *owner* sends its certificate to a *target*, who then uses it to encrypt any information it sends to the *owner*. Owners and targets can be a VSU, Remote Client, or any device that can use the *Internet-Key Exchange (IKE)* protocol to exchange certificates.

The roles of owners and targets is purely based on point-of-view. Whenever a *target* needs to receive encrypted traffic from an IKE compatible device, the *target* is viewed as an *owner* because it must send its certificate to the IKE device. The concept of owners and targets is illustrated in Figure 78. It's important to understand that a target must have an owner's certificate before it can send encrypted traffic to the owner.

Figure 78. Certificate exchange between VSUs
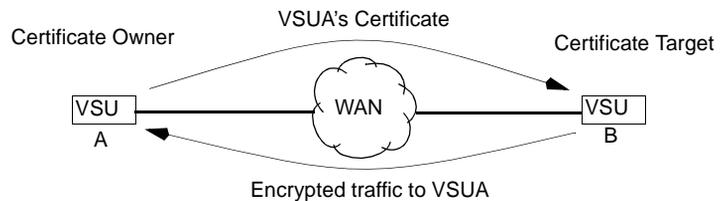
Before VSUA can receive encrypted traffic from VSUB, a certificate owned by VSUA must be sent to VSUB (the target of VSUA's certificate). The roles and process is reversed when VSUB needs to receive encrypted traffic from VSUA.

Assigning a Target for a Certificate

After a certificate is installed in a VSU (as described on), it must be assigned a target.

A *Bundle* is used to define a certificate having a specific target type, address, description, and queue position. The *Policy Manager* for *IKE Certificate Usage* lists all the bundles for a specific VSU.

The *Bundle Numbers* identify which *VSU Certificate* is associated with the bundle. For example, *Bundle Number 3* means that *VSU Certificate* number 3 is associated with the bundle. Up to eight bundles can be created, which directly relates to the number of signed certificates that can be dynamically stored in a VSU. The certificates stored on a specific VSU can be viewed from the *Policy Manager* for *My Certificates* ().

The target of a bundle is usually another VSU, but it can be any IKE compatible device. A target can be configured as an IP address, VPN object, fully qualified domain name, e-mail address, or director server name.

When a VSU recognizes that an target wants to communicate, the VSU uses the *IKE Certificate Usage* list to determine which bundle to send to the target. The search always starts at the top of the list, so it's important to put the most frequently used bundles at the top of the list.

There can be cases when you have to make a general purpose bundle that applies to any type of target. Always place that bundle at the bottom of the *IKE Certificate Usage* list.

- **Add** (IKE Certificate Policy). This screen is used to add a new IKE Certificate Policy to the IKE Usage Certificate list.

- **Bundle**. Combo box listing bundle numbers 1 through 8. 0 is the VPNet factory default bundle.

- **Memo**. Use this area to record notes about this IKE Certificate policy.

- **Target**. Type Identification of the remote tunnel endpoint. Used to determine which certificate to present to the other side. Target Type may be:

  - IP Address

  - VPN

  - FQDN (Fully Qualified Domain Name)

  - email

  - Directory Name

  - Any (target endpoint)

  Depending on the selection made, an appropriate field type appears to capture the respective information for the target type.

- **Locate This IKE Certificate Policy**. Allows you to specify the placement of the IKE Certificate Policy in the IKE Certificate Usage list.

To assign a target for a certificate:

1. From the *Device>Contents* column, select the VSU containing the certificate needing a target.

2. Click the **Policies** tab to bring it to the front.

3. From the *drop-down list*, select **IKE Certificate Usage**, then click **GO** to open the *Policy Manager* for *IKE Certificate Usage*.

4. Click **Add** to open the *Add IKE Certificate Policy.*

5. From the **Number** drop-down list, select which VSU certificate you want to configure.

   *Note: A VSU can dynamically store up to eight certificates. To identify how many certificates exist, click **Cancel** to return to the IKE Certificate Usage window, then from the **Type of Policy** drop-down list, select **My Certificates**.*

6. In the **Description** text box, type in information about the target. If the target is a VSU, typing in its name could be useful.

7. From the *Target Type* drop-down, select the type of target for the certificate.

    •   IP Address. Select to show the *Enter Target Address* text boxes. Type in the address of any IKE compatible device as a target. Typically, this is a VSU.

    •   VPN. Select to show the *Select Target VPN* list. VPN objects that have been created appears in the list. Select a specific VPN to be a target for the certificate. This only applies to Avaya Inc. VSUs of Version 3.0 and higher.

    •   FQDN. Select to show the *Enter Target Information* text box. Type in the *Fully Qualified Domain Name (FQDN)* to identify the target by its absolute name. For example, a target having the name **xyz** and a root of **vpnet.com**, has an absolute name of **xyz.vpnet.com**. The DNS Server that is used is configured from the "DNS tab" on page 69.

    •   e-mail. Select this item to show the *Enter Target Information* text box. Type in an e-mail address to identify the target by an e-mail address.

    •   Directory Name. Select this item to show the *Enter Target Information* text box. Type in an e-mail address to identify the target by an e-mail address.

    •   Any. Select this item for general purpose situations. For example, if you do not have enough certificates to configure.

8. From the *Locate this IKE Certificate Policy* options, select a queue position for the bundle.

    VSUs use bundles on an as-needed basis.

Chapter 10   Monitoring your network

This chapter describes the real-time monitoring facilities that the VPNmanager application provides. This includes the following

- Using SNMP to monitor the device

- Syslog Services

- Using Monitor

- Monitoring alarms

- Report Wizard

Using SNMP to monitor the device

The VPNmanager uses the SNMP protocol to monitor the security gateway. The security gateway includes a SNMP agent that supports MIB-II and a proprietary MIB. This agent is read-only and cannot be used to configure the security gateway. The agent can also send traps to a list of trap targets.

You configure the SNMP properties from the *Device Object SNMP* tab. Use this tab to configure the SNMP target devices or SNMP destination devices to which all security gateways report their status and alarm information.

SNMPv1, SNMPv2c, or, beginning with VPNos 4.2, SNMPv3 can be selected. You configure the trap and monitor strings and trap targets for SNMPv1 and SNMPv2c. You configure the trap targets and the SNMP user for SNMPv3. Since SNMPv1 and SNMPv2c send data in the clear, you can disable access to sensitive data including Filter Statistics, VPN Active Session, and the Event Log.

To configure SNMPv3, an Admin User with the required SNMPv3 privacy and authentication settings must be created. The same admin user can be used in the SNMP settings of different security gateways. This version of the VPNmanager does not have the capability to monitor security gateways using SNMPv3. These devices can be monitored using third party monitoring tools or MIB browsers that support SNMPv3.

If you select None, SNMP is disabled on the security gateway.

If you check *poll this security gateway* for either SNMPv1 or SNMPv2c, the VPNmanager actively monitors the device to see if the security gateway can be reached. If the device cannot be reached from the VPNmanager, an alarm is logged in the Alarm Console. For VPNmanager 4.2 and 4.3, if SNMPv3 is configured, the ability to poll the security gateway is disabled.

The traps that are generated by the security gateway are sent to the list of trap targets that are configured. The version of the trap that is sent is the same as the version of the SNMP Agent, that is, if the security gateway is configured for SNMPv1, a SNMPv1 trap is sent. A maximum of five trap targets can be specified and one of these can be the Directory Server. In large enterprises, the security gateways might also report to a network monitoring application, such as HP Open View.

Figure 79. The SNMP Tab for a security gateway Object
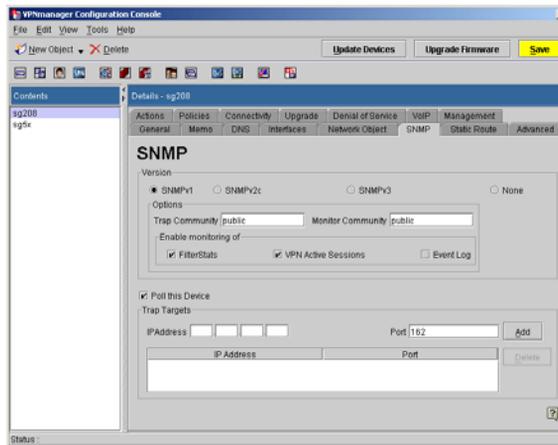
To add SNMP trap targets

To add an SNMP Trap Target for security gateway's at version VPNos 4.2 or later, do the following.

Note: To configure SNMPv3, see "Adding Admin Users for SNMPv3" on page 280.

1. From the *Contents* column, select the security gateway you want to configure.

2. Click the **SNMP** tab to bring it to the front.

3. Type in the SNMP Target IP Address

4. Enter the trap port. The default is 162

5. Click **Add** to add this setting to the list of Trap Targets.

6. When finished, click **Save**.

To add an SNMP Trap Target for security gateway's running versions prior to VPNos 4.2, do the following:

1. From the *Contents* column, select the security gateway you want to configure.

2. Click the **SNMP** tab to bring it to the front.

3. In the **Trap Community** text box, type in a unique community name.

4. Click **Add** to open the *Add SNMP Trap Target* dialog box.

5. In the SNMP Trap Target text boxes, type in the SNMP Target IP address.

6. Click **Close** to return to the *SNMP* tab, or **Apply** to add another address.

7. When finished, click **Save**.

8. When you want to send the configuration to one or more security gateways, click **Update Devices**.

To delete SNMP trap targets

1. From the *Contents* column, select the security gateway you want to configure.

2. Click the **SNMP** tab to bring it to the front.

3. From the **Trap Target** list, select the target you want to delete.

4. Click **Delete** to remove the target.

5. Click **Save**.

Adding Admin Users for SNMPv3

Configuring SNMP for a security gateway

1. In the *SNMP* tab choose the version as SNMPv3.

2. A list of Admin users who were configured for SNMPv3 are displayed in the list. Select one of the admin users.

3. Click **Save**.

Syslog Services

Security gateways have a syslog messaging facility for logging system error messages. The messages can be automatically sent to a destination running a Syslog server.

Use Policy Manager to configure and enable Syslog services, then move to your computer's command prompt and type in a start command.

Figure 80. The Policy Manager for Syslog Services
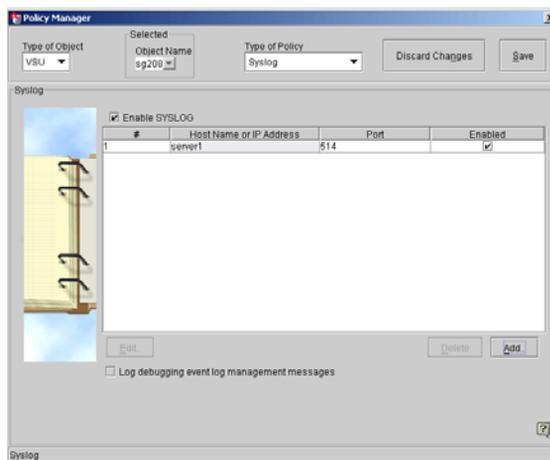
Enable SYSLOG When this box is checked, Syslog reporting to the target hosts in the list occurs.

- **#**. The number here is the order by rank of the target host to which Syslog data is sent.

- **Host Name or IP Address**. The domain name or IP address of the target logging archive machine.

- **Type**. UDP.

- **Port**. The port number of the Syslog host.

Add Syslog Policy

The Add Syslog Policy screen allows you to designate the host to which syslog messages are sent by the selected security gateway or all devices. It also enables syslog messages to be sent to the VPNmanager through a designated UDP port.

- **Hosts to receive log messages**. Enter the name or the IP Address of the target machine you are designating to receive syslog data.

- **Send event log message via**. Enter the port service (UDP) and number through which syslog messages are reported.

To run Syslog services:

1. From the **Device>Contents** column, select the security gateway you want to configure.

2. Click the **Policies** tab to bring it to the front. Select **Syslog**, then click **GO** to open the *Policy Manager* for *Syslog*.

3. Select the **Enable SYSLOG** check box so the security gateway will run Syslog services.

4. Click **Add** to open the *Add Syslog Policy* dialog box.

5. Use the *Hosts to receive log messages* options to configure the address of the Syslog Server.

   • To use a DNS name of the server, select the **Host Name** radio button, then type in a name.

   • To use the IP address of the server, select the **IP Address** radio button, then type in an address.

6. If you want the security gateway to send syslog messages to VPNmanager Console, configure the *Send event log message via* option.

   • To send the messages to a UDP Port, select the **UDP Port**, then type a number into the **Port Number** text box. The default number is 514.

7. Click **OK** to return to the *Policy Manager* for *Syslog* window.

8. Click the **Log debugging event log management messages** checkbox to log the messages.

9. Click **Save**.

10. From the upper right-hand of the window, click the close button to return to the *Configuration Console* window.

11. Move to the command prompt for your MS Windows computer.

12. Type in the following command line to create a directory for the syslog file, its size limit, protocol used, port number.

(Directory) \\Program Files\Avaya\VPNmanger\Console\Syslog\
..\jre\bin\java SyslogServer "-Lc:\ProgramFiles\AvayaVPN\Syslog" [-S*size*] [-P*port*] [-N*number*]

- If you want the size of the log file to be limited to a specific size, type in a specific size in kilobytes, otherwise the 8000 KB (8 MB) default size will be used.

- If the default UDP port numbers were not used in Step 6, type in the number used. The default values is 514 for UDP.

13. When you want to send the configuration to the security gateway, click **Update Devices** from the VPNmanager Console.

Using Monitor

When the Monitor button on the Main Console screen is selected, the Monitoring Wizard is launched. This wizard facilitates quick and easy construction of the desired presentation format and VPN information you wish to monitor. Once this setup is completed, the data and its presentation type is displayed on your VPNmanager console screen and is dynamically updated at your specified intervals. A hardcopy can be printed on demand.

Enterprise MIB

Monitoring is accomplished by selecting specific MIB objects from MIB-II and the VPNet Enterprise MIB within the VSUs. These individual items are assembled into preconfigured report groups for convenience. Individual parameters are also available for creating custom monitoring groups.

Monitoring wizard

The Monitoring wizard is designed to help you quickly set up the VPN objects and parameter groups you wish to monitor, and the format most appropriate for displaying the information produced.

The first Monitoring wizard dialog allows you to perform a high-level selection of the domain and VPN(s), then to choose specific network devices within the VPN. You can also select a monitoring group, which is a predefined suite of VPN parameters to monitor.

Device List For VPN Domain. This drop-down menu allows you to select a specific domain, or all domains to monitor.

Select Device(s). A list of all available network objects available for monitoring. You can select a single device, or select all devices displayed.

Select Monitoring Group. This window displays a list of all possible pre configured groups you may wish to monitor. These groups are constructed from one or more logically related items from MIB-II and the VPNet Enterprise MIB. The groups include:

- **Log Group** (provides details about attack events including time, attack type and a description).

- **System Group** (shows security gateway CPU Utilization)

- **ActiveSessions Group** (shows various details about the session in progress on the selected security gateway)

- **AddressTable** (atTable from MIB-II)

- **IPRouteTable** (ipRouteTable from MIB-II)

- **FilterStats** (provides detailed reporting on filtering statistics for the current security gateway)

- **FilterRules** (details about filter rules in effect and the traffic through them)

The following tables detail the individual enterprise MIB items in each of the monitoring groups.

Table 18. Log Group Parameters

Parameter	Description
Log Index	An integer identifying this row in the Log table.
Time	sysUpTime value when this attack occurred.
Attack Type	Indicates the reason that the packet was registered in the attack log. Six identifier types are reported: • 1 = SKIP header error (packet was not IPSec AH or IPSec ESP). • 3 = SKIP Algorithm mismatch. The parameters of the VPN that this packet belongs to does not match the VPN parameters in the SKIP header. • 4 = SKIP Authentication error. The authentication key in the offending packet was not correct. This type of attack results in an authFailure trap which shows up on the VPNmanager as Invalid Authentication Signature. • 6 = SKIP Encryption Header error. The packet's ESP trailer wasn't correct. • 7 = Remote client has exceeded the configured number on login attempts.
Packet Header (Hex)	The first 48 bytes of the packet header.

Table 19. System Group Parameters

Parameter	Description
CPU Utilization	A number from 1 to 100 representing CPU utilization in this security gateway.

Table 20. ActiveSessions Group Parameters

Parameter	Description
ActiveSessionsName	A VPNremote client name or a security gateway name as defined in VPNmanager.
Length	Length of this session in seconds.
Original IP	VPNremote client's originating IP address or remote security gateway IP address.
Xlated IP	VPNremote client's assigned address from the Client IP Address pool if configured. If the Client IP Address pool is not configured or this session is from a security gateway, then this attribute is empty.
Description	Textual description of this VPN indicating what key management is being used and what encryption, authentication and compression algorithms are being used. For example, IKE, 3DES, MD5, Compression.
Pkts In	Number of packets sent to this security gateway from the VPNremote Client or remote security gateway identified by Name during this session
Pkts Out	Number of packets sent from this security gateway to the VPNremote Client or remote security gateway identified by Name during this session.
Bytes In	Number of bytes sent to this security gateway from the VPNremote Client or remote security gateway identified by Name during this session
Bytes Out	Number of bytes sent from this security gateway to the VPNremote Client or remote security gateway identified by Name during this session.

Table 21. Address Table

Parameter	Description
Address Table Index	The interface on which this entry's equivalence is effective. The interface identified by a particular value of this index is the same interface as identified by the same value of ifIndex.
Physical Address	The media-dependent physical address.
Network Address	The Network Address (e.g., the IP address) corresponding to the media-dependent 'physical' address.

Table 22. ipRouteTable

Parameter	Description
Destination	The destination IP address of this route. An entry with a value of 0.0.0.0 is considered a default route. Multiple routes to a single destination can appear in the table, but access to such multiple entries is dependent on the table- access mechanisms defined by the network management protocol in use.
IP RouteTable Interface Index	The index value which uniquely identifies the local interface through which the next hop of this route should be reached. The interface identified by a particular value of this index is the same interface as identified by the same value of ifIndex.
Metric 1	The primary routing metric for this route. The semantics of this metric are determined by the routing-protocol specified in the route's ipRouteProto value. If this metric is not used, its value should be set to -1.
Metric 2	An alternate routing metric for this route. The semantics of this metric are determined by the routing-protocol specified in the route's ipRouteProto value. If this metric is not used, its value should be set to -1.
Metric 3	An alternate routing metric for this route. The semantics of this metric are determined by the routing-protocol specified in the route's ipRouteProto value. If this metric is not used, its value should be set to -1.
Metric 4	An alternate routing metric for this route. The semantics of this metric are determined by the routing-protocol specified in the route's ipRouteProto value. If this metric is not used, its value should be set to -1.
Next Hop	The IP address of the next hop of this route. (In the case of a route bound to an interface which is realized via a broadcast media, the value of this field is the agent's IP address on that interface.)

Table 22. ipRouteTable (Continued)

Parameter	Description
Route Type	The type of route. Note that the values direct(3) and indirect(4) refer to the notion of direct and indirect routing in the IP architecture.
	Setting this object to the value invalid(2) has the effect of invalidating the corresponding entry in the ipRouteTable object. That is, it effectively disassociates the destination identified with said entry from the route identified with said entry. It is an implementation-specific matter as to whether the agent removes an invalidated entry from the table. Accordingly, management stations must be prepared to receive tabular information from agents that corresponds to entries not currently in use. Proper interpretation of such entries requires examination of the relevant ipRouteType object.
	Enumerated values:
	1. other 2. invalid 3. direct 4. indirect
Route Proto	The routing mechanism via which this route was learned. Inclusion of values for gateway routing protocols is not intended to imply that hosts should support those protocols.
	Enumerated values:
	1. other 2. local 3. netmgmt 4. icmp 5. egp 6. ggp 7. hello 8. rip 9. is-is 10. es-is 11. ciscoIgrp 12. bbnSpfIgp 13. ospf 14. bgp
Route Age	The number of seconds since this route was last updated or otherwise determined to be correct. Note that no semantics of 'too old' can be implied except through knowledge of the routing protocol by which the route was learned.

Table 22. ipRouteTable (Continued)

Parameter	Description
Route Mask	Indicate the mask to be logical-ANDed with the destination address before being compared to the value in the ipRouteDest field. For those systems that do not support arbitrary subnet masks, an agent constructs the value of the ipRouteMask by determining whether the value of the correspondent ipRouteDest field belong to a class-A, B, or C network, and then using one of: mask network 255.0.0.0 class-A 255.255.0.0 class-B 255.255.255.0 class-C If the value of the ipRouteDest is 0.0.0.0 (a default route), then the mask value is also 0.0.0.0. It should be noted that all IP routing subsystems implicitly use this mechanism.
Metric 5	An alternate routing metric for this route. The semantics of this metric are determined by the routing-protocol specified in the route's ipRouteProto value. If this metric is not used, its value should be set to -1.
Route Info	A reference to MIB definitions specific to the particular routing protocol which is responsible for this route, as determined by the value specified in the route's ipRouteProto value. If this information is not present, its value should be set to the OBJECT IDENTIFIER { 0 0 }, which is a syntactically valid object identifier, and any conformant implementation of ASN.1 and BER must be able to generate and recognize this value.

Table 23. FilterStats

Parameter	Description
FilterStatsName	Interface name to which the filtering stats apply.
Pass In	Number of inbound packets allowed to pass through this interface.
Pass Out	Number of outbound packets allowed to pass through this interface.
Block In	Number of inbound packets not allowed to pass through this interface.
Block Out	Number of outbound packets not allowed to pass through this interface.
No Match In	Number of inbound packets that did not match any rule. This count includes all non-rule-matching packets, regardless of whether the packets were ultimately passed or blocked per the default rule.
No Match Out	Number of outbound packets that did not match any rule. This count includes all non-rule-matching packets, regardless of whether the packets were ultimately passed or blocked per the default rule.
Pass Log In	Number of inbound packets that were allowed to pass which have been logged. When a filtering rule is declared using the 'log option' (different from 'log action'), and the rule action is declared to be 'pass', a log entry is generated for each packet that matches the rule.
Pass Log Out	Number of outbound packets that were allowed to pass which have been logged. When a filtering rule is declared using the 'log option' (different from 'log action'), and the rule action is declared to be 'pass', a log entry is generated for each packet that matches the rule.
Block Log In	Number of inbound packets not allowed to pass which have been logged. When a filtering rule is declared using the 'log option' (different from 'log action'), and the rule action is declared to be 'block', a log entry is generated for each packet that matches the rule.
Block Log Out	Number of outbound packets not allowed to pass which have been logged. When a filtering rule is declared using the 'log option' (different from 'log action'), and the rule action is declared to be 'block', a log entry is generated for each packet that matches the rule.

Table 23. FilterStats (Continued)

Parameter	Description
No Match Log In	Number of inbound packets on a given interface that did not match any filtering rule and were subsequently logged, regardless of whether or not the packet was ultimately passed or blocked per the interface's default rule.
No Match Log Out	Number of outbound packets on a given interface that did not match any filtering rule and were subsequently logged, regardless of whether or not the packet was ultimately passed or blocked per the interface's default rule.
Packets Logged In	Total number of inbound packets that should have been logged. This number includes packets that matched filtering rules declared using either the 'log option' or the 'log action'.
Packets Logged Out	Total number of outbound packets that should have been logged. This number includes packets that matched filtering rules declared using either the 'log option' or the 'log action'.
Skip Log In	Number of inbound packets that should have been logged, but the log buffer was full. Log records are stored in a fixed size non-circular buffer. When the buffer is full, no new log records are written until the buffer is drained via either the security gateway console or the VPNmanager.
Skip Log Out	Number of outbound packets that should have been logged, but the log buffer was full. Log records are stored in a fixed size non-circular buffer. When the buffer is full, no new log records are written until the buffer is drained via either the security gateway console or the VPNmanager.
Return In	Number of inbound packets that matched a rule requiring that a TCP-Reset or ICMP packet be sent in response.
Return Out	Number of outbound packets that matched a rule requiring that a TCP-Reset or ICMP packet be sent in response.
Account In	Number of inbound packets that matched a filtering rule with a declared action of 'count'.
Account Out	Number of outbound packets that matched a filtering rule with a declared action of 'count'.

Table 23. FilterStats (Continued)

Parameter	Description
Bad Frag Alloc In	Number of failed attempts to allocate a Fragment table entry for inbound packets. This occurs when a filter rule is declared using the 'keep frag' option. A packets matching this rule cause a Fragment table entry to be allocated. If the table is full, the allocation fails.
Bad Frag Alloc Out	Number of failed attempts to allocate a Fragment table entry for outbound packets. This occurs when a filter rule is declared using the 'keep frag' option. A packets matching this rule cause a Fragment table entry to be allocated. If the table is full, the allocation fails.
New Frag Alloc In	Number of successful attempts to allocate a Fragment table entry for inbound packets. This occurs when a filter rule is declared using the 'keep frag' option. A packets matching this rule cause a Fragment table entry to be allocated. This value does not reflect the size of the table, only the number of entry allocations which succeeded.
New Frag Alloc Out	Number of successful attempts to allocate a Fragment table entry for outbound packets. This occurs when a filter rule is declared using the 'keep frag' option. A packets matching this rule cause a Fragment table entry to be allocated. This value does not reflect the size of the table, only the number of entry allocations which succeeded.
Unneeded Frag Alloc In	Number of successful, but unnecessary attempts to allocate Fragment table entries for inbound packets. When a filter rule is declared using the 'keep frag' option, matching packets cause a Fragment table entry to be allocated. This allocation takes place before the determination is made that the packet is indeed a fragment. If the packet is later determined NOT to be a fragment, the table entry is de-allocated and this counter is incremented.

Table 23. FilterStats (Continued)

Parameter	Description
Unneeded Frag Alloc Out	Number of successful, but unnecessary attempts to allocate Fragment table entries for outbound packets. When a filter rule is declared using the 'keep frag' option, matching packets cause a Fragment table entry to be allocated. This allocation takes place before the determination is made that the packet is indeed a fragment. If the packet is later determined NOT to be a fragment, the table entry is de-allocated and this counter is incremented.
Bad State Alloc In	Number of failed attempts to allocated State table entries for inbound packets. This occurs when a filter rule is declared using the 'keep state' option. Packets that match the rule cause a State table entry to be allocated. This allows expected return packets to bypass other filtering rules that might normally block them. Allocation fails if the State table is full and a new entry cannot be allocated.
Bad State Alloc Out	Number of failed attempts to allocated State table entries for outbound packets. This occurs when a filter rule is declared using the 'keep state' option. Packets that match the rule cause a State table entry to be allocated. This allows expected return packets to bypass other filtering rules that might normally block them. Allocation fails if the State table is full and a new entry cannot be allocated.
Keep State Alloc In	Number of successful attempts to allocated State table entries for inbound packets. This occurs when a filter rule is declared using the 'keep state' option. Packets that match the rule cause a State table entry to be allocated. This allows expected return packets to bypass other filtering rules that might normally block them.
Keep State Alloc Out	Number of successful attempts to allocated State table entries for outbound packets. This occurs when a filter rule is declared using the 'keep state' option. Packets that match the rule cause a State table entry to be allocated. This allows expected return packets to bypass other filtering rules that might normally block them.

Table 23. FilterStats (Continued)

Parameter	Description
Cache Hit In	Number of cache hits for inbound packets on this interface. Each inbound packet is examined to see if a packet with identical characteristics exists in the outbound cache for this interface. If a match is found, the resulting rule match applied to the previous packet is applied to the current one, bypassing the rest of the filtering mechanism.
Cache Hit Out	Number of cache hits for inbound packets on this interface. Each outbound packet is examined to see if a packet with identical characteristics exists in the outbound cache for this interface. If a match is found, the resulting rule match applied to the previous packet is applied to the current one, bypassing the rest of the filtering mechanism.
Good Pullup In	Number of failed pullup operations occurring for inbound packets. These occur when a packet is fragmented across multiple internal memory buffers and there is insufficient information available to properly process the packet. Successive memory buffers are read until there is enough information to process the packet.
Good Pullup Out	Number of failed pullup operations occurring for outbound packets. These occur when a packet is fragmented across multiple internal memory buffers and there is insufficient information available to properly process the packet. Successive memory buffers are read until there is enough information to process the packet.
Bad Pullup In	Number of failed pullup operations occurring for inbound packets. These occur when a packet is fragmented across multiple internal memory buffers and there is insufficient information available to properly process the packet. Successive memory buffers are read until there is enough information to process the packet.
Bad Pullup Out	Number of failed pullup operations occurring for outbound packets. These occur when a packet is fragmented across multiple internal memory buffers and there is insufficient information available to properly process the packet. Successive memory buffers are read until there is enough information to process the packet.

Table 23. FilterStats (Continued)

Parameter	Description
No Match Pass In	Number of inbound packets for a given interface which did not match any filtering rule and were ultimately allowed to pass per the interface's default rule.
No Match Pass Out	Number of outbound packets for a given interface which did not match any filtering rule and were ultimately allowed to pass per the interface's default rule.
No Match Block In	Number of inbound packets for a given interface which did not match any filtering rule and were ultimately blocked per the interface's default rule.
No Match Block Out	Number of outbound packets for a given interface which did not match any filtering rule and were ultimately blocked per the interface's default rule.

Table 24. Filter Rules

Parameter	Description
Rule	Filtering Rule description. Shows the rule parameters as they would appear when typed in or displayed at the security gateway console.
Rule Match	Number of packets which matched this filtering rule.
Rule Byte Count	Total byte count for packets which match this filtering rule.

Table 25. Stats Group

Parameter/ Group	Description
Active Ports	The number of active ports on this security gateway.
Traffic Rate Table Group	See Table 26 Traffic Rate Table Group.
Overview Statistics Table Group	See Table 27 Overview Statistics Table Group.
Ethernet Statistics Table Group	See Table 28 Ethernet Statistics Table Group.

Table 26. Traffic Rate Table Group

Parameter	Description
Traffic Port Description	A description of each port.
Traffic Port Index	The index of this port. Indices are: Private = 0 or 2, Public = 1 or 3 (2 and 3 appear only for the security gateway-100).
Traffic Port Interface Index	The ifIndex value from the MIB-II if Table.
Summary Interval	A time interval used to average this traffic rate.
Packets From Port	The average rate (in packets per second) at which packets have been transmitted from this port over the last <Summary Interval> seconds.
Packets To Port	The average rate (in packets per second) at which packets have been received on this port over the last <Summary Interval> seconds.
KBits From Port	The average rate (in KBits per second) at which packets have been transmitted from this port over the last <Summary Interval> seconds.
KBits To Port	The average rate (in KBits per second) at which packets have been received on this port over the last <Summary Interval> seconds.

Table 27. Overview Statistics Table Group

Parameter	Description
Overview Port Description	A description of each port.
LAN Frames Received	The number of LAN frames received on this port.
LAN Frames Xmitted	The number of LAN frames transmitted from this port.
LAN Frames Discard	The total number of LAN frames discarded on this port because of errors.
Ethernet Header Errors	The number LAN frames discarded on this port because of Ethernet header errors.
Non-IP Packets Received	The number of non-IP packets received on this port.
VPN Packets Received	The number of VPN packets received on this port.
Non-VPN IP Packets Received	The number of non-VPN IP packets received on this port.
Non-VPN Packets Blocked	The number of non-VPN packets blocked on this port.
Tunnel Config Errors	The number of packets dropped on this port because of tunnel configuration problems.
IP Header Length Errors	The number of packets dropped on this port because of an invalid IP header length.
Address Map Discards	The number of packets dropped because of IP Address Map errors.

Table 28. Ethernet Statistics Table Group

Parameter	Description
EtherStat Port Description	A description of each port.
Total Frames Received	Total number of frames received on this port.
Total Frames Xmitted	Total number of frames transmitted from this port.

Table 28. Ethernet Statistics Table Group (Continued)

Parameter	Description
Total Frames Filtered	Total number of frames discarded on this port because the destination MAC address of the frame was determined by the bridge logic to be attached to the same network segment as this port.
Total Frames Discarded	Total number of frames discarded on this port because of some error.
Local Frames Received	Total number of frames received on this port destined for the this unit's IP address.
Local Frames Xmitted	Total number of frames transmitted from this port with the unit's MAC address as the source MAC address.
Available Xmit Buffers	Total available VPNos transmit buffers on this port.
No-Receive-Buffer Errors	The number of packets dropped on this port because no VPNos receive buffers were available.
Missed Frames	The number of frames dropped on this port because the Ethernet chip had no receive buffers available.
CRC Errors	The number of packets dropped on this port because of CRC errors.
Frame Errors	The number of packets dropped on this port because of frame errors.
Overflow Errors	The number of packets dropped on this port because of overflow errors.
No-Xmit-Buffer Errors	The number of packets not transmitted on this port because no VPNos transmit buffers were available.
Lost Carrier Errors	The number of packets not transmitted on this port because of lost carrier errors.
Xmit Collisions	The number of packets not transmitted on this port because time collisions.
Time Underflow Errors	The number of packets not transmitted on this port because of time underflow errors.
Timeout Errors	The number of packets not transmitted on this port because of time-out errors.

Table 28. Ethernet Statistics Table Group (Continued)

Parameter	Description
Retry Overflow Errors	The number of packets not transmitted on this port because of retry overflow errors.
Miscellaneous Errors	The number of packets dropped on this port because of other miscellaneous errors.

Define Custom

The Define Custom screen allows you to define a custom monitoring group that only collects the data you specify. You select the desired MIB parameters from the *Available Data* column, then moving them into the *Current Data in Group* column. All of the available MIB-II and VPNet Enterprise parameters in the Monitoring Groups are available.

- **Name**. Enter the name you wish to call your custom group.

- **Current Data in Group**. This is a list of the individual enterprise MIB parameters that compose your new Group.

- **Available Data**. This is a list of all possible enterprise MIB parameters you may select to monitor. Use the Move Left arrow to transfer the highlighted parameter into the Current Data in Group column.

Monitoring wizard (Presentation)

The Monitoring presentation screen is used to select the display type for the monitored data. The update frequency is also indicated here.

Presentation

There are four types of presentations:

- Bar graph

- Line graph

- Pie chart

- Table

Some types of data cannot be displayed in all four presentation styles. For example, only the System Group can be presented as a bar graph.

Only the available presentation types appear in this field for the group previously selected (the table is the most common format for most of the groups).

Update Time. Update time defines how often your presentation is updated (security gateway MIB is re-read). You can choose minutes or seconds.

Display. The display area offers two selections for how your security gateway groups are presented, either one window per security gateway, or a single window in which the desired security gateway is selected from a drop-down menu.

Monitoring alarms

On the main VPNmanager window, the Alarm pane displays alarm information arriving from the security gateways in the VPN. When an alarm arrives, a rotating red beacon light activates. Conditions causing alarms include a security gateway device not reachable and several security attacks, such as a manager authentication failure, a key failure, or a CCD failure. The alarm console can also be used as a general trap target gather SNMP trap information from other network devices.

Two buttons appear at the bottom of the pane: Properties, and Delete.

By default, all device alarms are displayed, however, alarms from a specific security gateway can also be shown. All alarm information is stored locally on the VPNmanager Console.

This window provides detailed information about the alarm including a time stamp, the security gateway generating the alarm, alarm definition, first and last occurrence. This window appears even if it does not contain any content. The most recent entry is at the top of the list.

- **Properties**. The Alarm Properties screen displays a list of specific alarm types and their corresponding disposition action: ignore or take action. Refer to Table 29 for Alarm Type descriptions. The default is Take action on Alarm.

- **Delete**. A Delete button appears at the bottom of the window. The highlighted alarm(s) is deleted when the Delete button is clicked.

Figure 81. VPNmanager, Alarm Pane, Properties
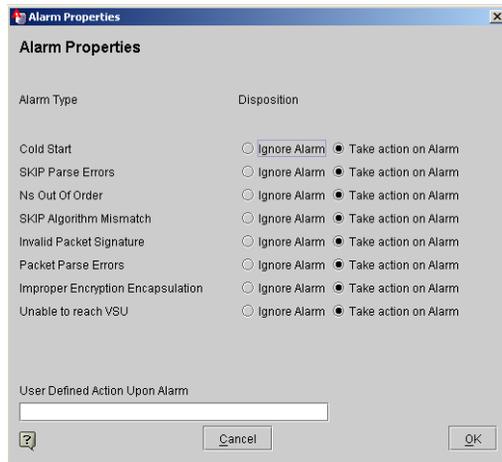
Alarm Types

Table 29. Alarm Descriptions

Alarm Type	Description
Cold Start	Indicates a security gateway was restarted via a power cycle, the security gateway console, or VPNmanager.
SKIP Parse Error	Indicates a packet that had an incorrect SKIP header was received. This alarm could result from a cryptographic attack.
Ns Out of Order	Indicates a packet that contained an expired key was received. This alarm could result from a cryptographic attack.
SKIP Algorithm Mismatch	Indicates that a packet for which one of the three algorithms (compression, encryption, or authentication) used to secure it did not match the VPN configuration within the security gateway where it was received. This alarm could result from a cryptographic attack.
Invalid Packet Signature	Indicates a packet that failed authentication was received. May indicate that the packet's source is using invalid encryption keys. This alarm could result from a cryptographic attack.
Packet Parse Error	Indicates the receipt of a packet that could not be properly decrypted. Usually due to an incorrect IP packer header.
Improper Encryption Encapsulation	Indicates a packet that could not be properly decrypted was received. This alarm could result from a cryptographic attack.
Unable to Reach device	Indicates no response was received from a security gateway to a VPNmanager polling request for management data.

You can select to either ignore the alarm or take action on the alarm. If "Take Action on Alarm" is checked, the User Defined Action Upon Alarm is executed.

User Defined Action Upon Alarm. Enter the name of the application to be launched when any alarm is generated. The action can be any executable file (for example, an application that pages the system administrator).

Report Wizard

The Report wizard is used to generate summaries of configuration details and a variety of reports about the VPN, its components, and how they are performing. This is especially useful in the configuration debugging process, and as an audit trail to document the overall VPN configuration. (For accounting, see SYSLOG).

The first Report wizard screen allows you to specify the objects you wish to include in the report. The available objects include:

- IP Group

- User

- User Group

- Security gateway (device)

- VPN

To create a report using the report wizard:

1. Move to the Main Console.

2. Click **Repor**t to start the *Report Wizard*.

3. In the Report Contents portion of the screen, select the object types to be included in the report.

   The Select All and Deselect All buttons are provided for convenience.

4. Click **Next**.

5. In the **Show Report Title** text box, type the report title.

   Report format details including date and time, report title, author, page numbering, and the type font and font size.

   The available font types are: Arial, Times Roman, and Helvetica. The available font sizes range from 8 points to 72 points.

6. Click **Next**.

   Depending on the objects selected in the initial screen, each object is displayed as part of the report wizard.

7. Select the desired object groups to be included in the report.

    Note: The Summary button presents a single-screen overview of all the currently set report selections and options. Advanced users may wish to jump to this screen immediately.

8. Click **Next**.

9. Select additional information for the object group to be included in the report.

10. Click **Next**.

11. Click **Finished** when all report information has been selected.

    You then have a choice of the output file type, HTML or PDF. The output file may be viewed on the screen, then sent to a printer if hardcopy is desired. Be sure you have an Adobe Acrobat reader to view the PDF file, or a web browser to view the HTML file.

Generating the report

When you are satisfied with the report selections made, click on the Finished button to generate the report. The report window appears after a short pause. If a hardcopy is desired, you may save the report as a PDF or html file, then print from Acrobat or a browser (respectively).

Figure 82. Report Sample
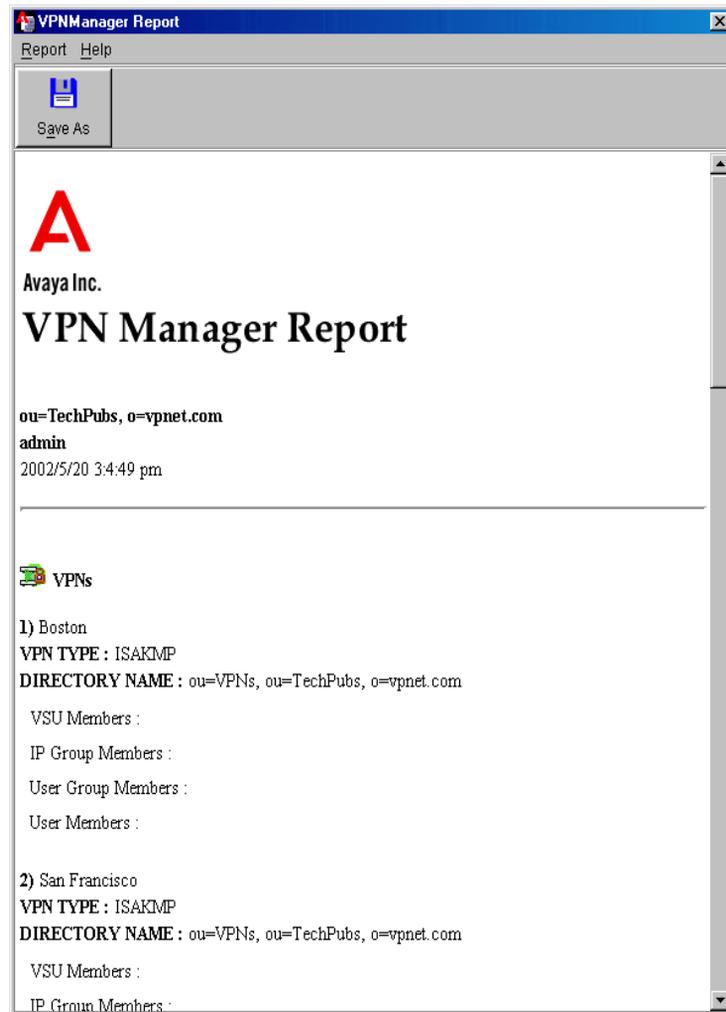
Chapter 11   Device management

From the VPNmanager Console, you can manage and check that status of the security gateways This chapter describes:

- Using the Management tab to change administrataive passwords and set up SSH and Telnet to connect to a security gateway
- Using the Connectivity tab to ping the security gateway
- Using the Device Actions tab to reboot the device, set the device time and import a device configuration
- Importing and exporting VPN configurations to a device
- Exporting RADIUS

Using the Management tab

The Device>Management tab is used to set up the SSH/Telnet feature and to change the administrator's password for the security gateway.

Setting Up SSH and Telnet

Beginning with VPNos 4.31, SSH (Secure Shell) and Telnet can be used to access the security gateway's CLI. When you use SSH to transfer data, the entire log in session, including transmission of the password, is encrypted. If you use Telnet to communicate with the security gateway, data transfer is not encrypted.

You can turn on both SSH and Telnet, and you can specify the port to use and the allowed IP addresses that can access the security gateway. The default is the following:

- SSH is enabled for *Any* network objects on the *private* zone, all other zones are disabled. Only the root and the monitor users can use SSH to access the security gateway.

- Telnet is disabled on all zones.

Use the *Device>Management* tab to change the defaults and to configure or change the security gateway SSH/Telnet feature.

When you log in to the security gateway using either SSH or Telnet, the security gateway's CLI interface is displayed. You can then use the CLI commands to troubleshoot the security gateway. To use CLI commands see the VPNos Configuration Guide.

Note: To restrict access to hosts or networks, Firewall rules limit access from specific zones. See Firewall rules template chapter

To set up SSH or Telnet

1. Move to the *Configuration Console* window.

2. From the Icon tool bar, click **Devices** to list all security gateways in the *Contents* column.

3. From the Contents column, select the security gateway to configure for SSH or Telnet connection.

4. Click the **Management** tab, to bring it to the front.The SSH/Telnet page is displayed.

5. By default SSH is enabled and the port 22 is configured, and Telnet is disabled. Make the appropriate changes to enable or disable either or both of these and to change the port if required

6. In the **Allowed** area, select **Zones** to set which zones can be used. The **SSH/Telnet Zones Configuration** dialog is displayed, and the zones that are configured as listed.

   For SSH, by default, the private zone is allowed.

   For Telnet, you must select a zone as all zones are disabled by default.

7. Move the zones from **Blocked** to **Allowed**. Click **OK**.

8. Select **Networks**, to configure the IP address to use to access the security gateway

   • To add an IP address, click **Add**, enter the address and click **OK**.

   • To add network objects, from **Available** list, select the network object and click **Move Left** to the **Allowed** column. Click **OK.**

   For SSH, by default **Any** is allowed.

9. Click **Save** and then click **Update Devices** to send the configuration change to the security gateway.

Changing device administrator's passwords

(VPNos 4.31)The following security gateway administrators configure and monitor the security gateway.

• *Super user* is the VPNmanager centralized management administrator. The VPNmanager super user has full read and write privileges to configure and monitor security gateways. The super user name and the password are entered from the VPNmanager console and are authenticated before VPNmanager is used to make configuration changes on the security gateway. For centralized management, the security gateway must have the *Permit Centralized Management* feature enabled. See the VPNos Configuration Guide for details.

• *Root* is the login name for the security gateway administrator. The root administrator has full privileges to configure and maintain a specific security gateway network and user configuration.

• *Monitor* is the login name for an administrator who can view the Inspect properties and monitor sub functions of the security gateway's interface software. The monitor user has read-only permissions.

These administrator's cannot be deleted but their passwords can be changed. Go to the *Device>Management* tab to change the passwords.

To reset the passwords

1. Move to the *Configuration Console* window.

2. From the Icon tool bar, click **Devices** to list all security gateways in the *Contents* column.

3. From the **Contents** column, select the security gateway that requires the administrator passwords reset.

4. Click the **Management** tab, to bring it to the front.

5. Select **Reset Password** to see the configuration reset buttons.You can reset the super user, root user or monitor user password.

6. Click **Reset** for the administrator user that should be changed.The **Reset Password** dialog is displayed.

7. Enter the new password. The password must be a minimum of six characters.

8. Click **OK**. The new password is automatically reset on the security gateway.

Using the Connectivity tab

The Device>Connectivity tab provides basic communications testing. Ping between the VPNmanager workstation and a security gateway, or the VPNmanager and an address or DNS server.

Figure 83. The Connectivity tab for a security gateway Object
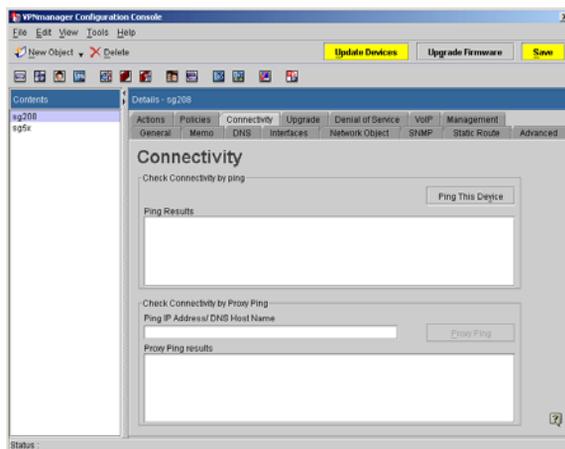
Two methods for testing the connectivity of a security gateway are:

• Ping between the VPNmanager workstation and a security gateway

• Proxy ping, which has been initiated by the VPNmanager, from a security gateway to any node.

A ping between the VPNmanager workstation and a security gateway is useful for verifying that the security gateway is powered on and operational, and that an IP network connection from the VPNmanager workstation to the security gateway exists.

The *Ping This Device* button initiates a clear text (non-VPN traffic) ping from the security gateway to the VPNmanager workstation.

Check connectivity by ping

To execute this ping:

• Select a security gateway from the Contents list, then click on the **Ping This security gateway** button.

• The ping results are displayed in the Ping Results window.

The Ping Results window indicates that connectivity to the security gateway's IP address is being checked.

A result of "<IP address of security gateway> is alive" indicates a reply was received from the IP address of this security gateway.

A result of "security gateway unreachable" indicates no reply was received.

To directly ping a specific security gateway:

1. Move to the *Configuration Console* window.

2. From the **Contents** column, select the security gateway that you want to ping.

3. Click the **Connectivity** tab to bring it to the front.

4. Click **Ping This security gateway** to start the ping.

   Information about the ping appears in the *Ping Results* text box.

Check Connectivity by Proxy Ping

Ping this Address/DNS name: Enter the IP address or DNS name.

Results are displayed in the Proxy Ping Results window.

To proxy ping a specific security gateway:

1. Move to the *Configuration Console* window.

2. From the **Contents** column, select the security gateway that you want to ping.

3. Click the **Connectivity** tab to bring it to the front.

4. In the **Ping This Address/DNS Server**, type in an address or host name of the proxy.

5. Click **Proxy Ping** to start the ping.

   Information about the ping appears in the *Ping Results* text box.

Using the Device Actions tab

The Device Actions tab is used to perform basic functions on the security gateway. Basic functions include Refresh Configuration, Reset Device Time, Reboot Device, Re-setup security gateway, Import Configuration, and Ethernet Speed.

Note: The Import Configuration and Ethernet Speed features are visible and only on some models.

Figure 84. The Actions tab for a security gateway Object
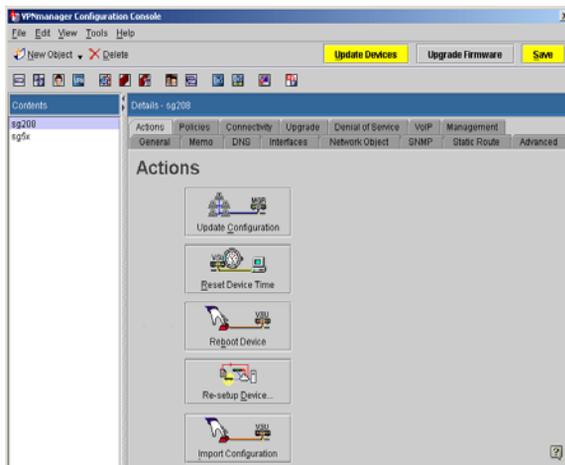
Update Configuration

When changes are made to a Device Object, use the *Update Configuration* button to send the changes from the server to a specific security gateway.

Reset Device Time

Click Reset Time to synchronize the security gateway and VPNmanager workstation to Greenwich Mean Time.

Reboot Device

To restart a security gateway at any time, click Reboot. (A Cold Start alarm is logged by VPNmanager and any other trap targets specified.) Note that any existing VPN connections are dropped and are re-established following the security gateway reboot sequence.

Reboot should normally not be necessary except when the fundamental configuration changes (such as changing the security gateway's IP address) are made.

The time for the reboot process to complete varies with each security gateway series. The VSU-1200/7500 series taking up to approximately two minutes during which VPN connections through this security gateway are down. For this reason, security gateway reboots should be performed during scheduled maintenance whenever possible.

Re-setup Device

Allows a complete re-setup of the security gateway. This is normally done when the security gateway created did not exist in the network, or when the security gateway has been replaced with a new unit.

Import Device Configuration

You can use the Import Device Configuration feature in VPNmanager to import configuration data from security gateways running VPNos 4.31, for use in VPNmanager.

While it is feasible to configure a small number of security gateways using the VPNos Web interface or the security gateway's CLI, it quickly becomes impractical for larger installations. When switching to VPNmanager for centralized management of devices which have already been configured, the Import Device Configuration feature allows the devices' existing configuration data to be easily migrated to VPNmanager.

When a device configuration is imported into VPNmanager, only the device-level configuration settings are imported. The domain-level settings, e.g. VPNs, Firewall templates, Users and Failover are not

imported. The configuration settings that are imported apply only to the specified device.

Note: If VPNmanager already has any configuration data for a particular device, the retrieved data overwrites the existing data for that device.

In VPNmanager 3.4, the Import Device Configuration feature supports importing of the following configuration settings:

- Interfaces
- Static Routes
- Network Objects
- Services
- VoIP
- DNS
- NAT
- NAT Traversal
- DoS
- SSH/Telnet
- Management Access

To import configuration data for a device:

1. Select "Devices" on the Configuration window in VPNmanager.

2. Select the device from which configuration data will be imported. (If the device entry does not yet exist in VPNmanager, simply create a new device, specifying its IP address and selecting "Set Up Later" in the Device Setup Wizard.)

3. Select the device Actions tab.

4. Click the Import Configuration button.

Ethernet Speed

The Ethernet Speed button only appears when a VSU10000 is the selected device.

Ethernet Speed button allows the VPNmanager to configure the Ethernet speed on a per port basis.

When the Ethernet Speed button is selected, there is a short delay in presenting the Ethernet Speed dialog box. This delay is due to VPNmanager trying to contact the security gateway to retrieve the current port speed settings. When the VPNmanager has retrieved the current speed settings, the Ethernet Speed dialog box displays the public port settings by default. The current private port settings are displayed at the top of the Ethernet Speed dialog box.

Port. Select the public or private port to configure the port speed of the selected security gateway.

Set Speed. Configure the Ethernet speed by selecting one of the following speed options:

Note: When selecting the port speed, be sure to select a speed that is supported by the host PC. If the host PC does not support the selected speed, the VPNmanager looses connectivity to the security gateway.

Auto Negotiate. Auto negotiation allows the security gateway's Ethernet port and host PC to automatically select the correct port speed and duplex mode to be used between the two ports.

1000 Mbps, Full Duplex. This option allows the VPNmanager to configure the security gateway's Ethernet port speed to 1000 Mbps in full duplex mode. In full duplex mode, the Ethernet port is capable of sending and receiving packets simultaneously over the network at 1000 Mbps.

1000 Mbps, Half Duplex. This option allows the VPNmanager to configure the security gateway's Ethernet port speed to 1000 Mbps in half duplex mode. In half duplex mode, the Ethernet port is capable of either sending or receiving packets over the network at 1000 Mbps.

100 Mbps, Full Duplex. This option allows the VPNmanager to configure the security gateway's Ethernet port speed to 100 Mbps in full duplex mode. In full duplex mode, the Ethernet port is capable of sending and receiving packets simultaneously over the network at 100 Mbps.

100 Mbps, Half Duplex. This option allows the VPNmanager to configure the security gateway's Ethernet port speed to 100 Mbps in half duplex mode. In half duplex mode, the Ethernet port is capable of sending or receiving packets over the network at 100 Mbps.

10 Mbps, Full Duplex. This option allows the VPNmanager to configure the security gateway's Ethernet port speed to 10 Mbps in full duplex mode. In full duplex mode, the Ethernet port is capable of sending and receiving packets simultaneously over the network at 10 Mbps.

10 Mbps, Half Duplex. This option allows the VPNmanager to configure the security gateway's Ethernet port speed to 10 Mbps in half duplex mode. In half duplex mode, the Ethernet port is capable of sending or receiving packets over the network at 10 Mbps.

Redundancy

This button only appears when a VSU-1200/7500 is the selected device.

This screen appears when the Redundancy button on the security gateway Action tab is clicked. It is used to set up specific redundancy attributes when two VSU-1200/7500s are being used to backup each other.

This function also allows you to check the status of the redundant systems in the VSU-1200/7500, and allows you to manually switch over the active Ethernet ports from the primary to secondary ports or vice versa. This switch-over function can be performed for both ports on a single card or for an individual active port.

Network Interface Status

Card 1, Card 2 Shows the current status of the public and private Ethernet ports located on the primary and secondary Ethernet interface cards. The port names are shown next to three icons indicating the current port status. The first box indicates whether the port is on, off, or defective. The second box indicates if the port is connected and at what speed its operating (100 or 10 megabits per second), and the last box indicates Full or Half duplex.

Fan/Power Status Indicates the power supply fan status. The Fan/Power Status section shows the current state of the redundant cooling fans and power supply modules. If a fan or power supply modules fails, a FAILED status is displayed indicating which component failed. Refer to the

VSU-1200 User Guide for instructions on how to replace the failed component.

IPSec Engine Status The IPSec Engine Status section shows the current state of the VSU-1200's two packet processor engines (PPE). If either PPE fails, a FAILED status is displayed indicating which PPE failed. Both PPEs must be functional for the VSU-1200 to operate correcting. The PPEs and Ethernet cards are enclosed in a tamper-evident case and can only be serviced by an authorized technician. Contact your customer service representative for instructions on getting the VSU-1200 repaired.

Switching

To individually switch the active public or private ports, select which active ports to switch from and which passive ports to switch to, then click the Switch Ports button. Note that the active public and private ports can only be switched to passive ports of the same type. A public port cannot be switched to a private port or vice versa.

Importing and exporting VPN configurations to a device

A secure, inter-company extranet can be created by exporting a VPN configuration to a file that is then imported by other VPNmanager installations. Select Import VPN when you receive your exported VPN file and have it copied to a local directory. You will need the password from the exporting administrator.

Export VPN

Creating an extranet is a cooperative effort between system administrators running independent copies of VPNmanager and involves the same steps as creating any other VPN: create the VSUs, then the Groups and Clients, and finally the VPN.

The names chosen for VPN components must be synchronized within each corporation's VPNmanager. This requires close coordination between the system administrators during the VPN component creation process and can be achieved by performing the following procedure:

- The administrators at each corporation agree that all VPN components will be created by one of the administrators (the "exporting" administrator) and that the exporting administrator will create and deliver an export VPN configuration file to the other administrators (the "importing" administrators).

- The exporting administrator then creates security gateways, groups, users, and VPNs required, with the exception of the security gateways under management control of importing administrators.

  The VPN name must be unique to both the exporting and importing administrators' VPNmanager databases.

- When creating an "alien Group," which is a group that includes IP address/mask pairs residing within an importing administrator's network, the exporting administrator associates each alien Group with an extranet device.

  In the Group configuration, the IP address of the importing administrator's security gateway must be specified if any tunnel mode VPNs include this security gateway.

- After creating the VPN, the exporting administrator exports the VPN configuration file and delivers it, along with the password used to protect the file, to the importing administrators.

- The importing administrators import the VPN configuration file using the supplied password.

- Finally, the importing administrators edit the alien Group, modifying the security gateway association appropriately.

The Export VPN screen appears allowing you to select the VPN to be exported.

Once you have entered the password, click OK. The new VPN file decodes and is entered into the VPNmanager server and the new VPN objects appear.

If any pair in the "Current IP Network/Mask Pairs" list represents a network under your management control, associate the Group with the appropriate security gateway by modifying the "Associate this Group with security gateway" picklist.

For Groups with network/mask pairs that are not under your management control, leave the "Associate this Group with security gateway" picklist as an extranet device and confirm that the "Extranet IP Address" entry field contains the correct IP address, especially if any tunnel mode VPNs include this security gateway.

Repeat this step for all Groups in the imported VPN.

NOTE: For any Certificate Based IKE extranet VPNs, verify that the proper certificates are installed on all devices.

Exporting RADIUS

The Export RADIUS function is used to export VPN information to an existing RADIUS database. This is primarily for backwards compatibility, but also useful if you wish to convert your existing VPN (using local security gateway-based user authentication) into a dynamic VPN for future scalability. It is, however, expected that LDAP will be the preferred method of building dynamic VPNs.

In this procedure, your existing client configuration information is migrated to the RADIUS database through a RADIUS-compatible export file. The Export RADIUS pane appears with a list of all users you wish to include in the export. When you click OK, VPNmanager creates a text file.

The saved text file consists of entries that must be added to the RADIUS server "users" file.

The Users file variable parameters are:

- <Client_name> – The name of the Client as entered in VPNmanager. Case and spelling are significant. This parameter is written by VPNmanager.

- <authentication password> – The response required from the Client to the authentication challenge sent through the security gateway by the RADIUS server. Case and spelling are significant. This field must be entered by the system administrator.

- <VPN-specific algorithm and key information> – Information specific to the VPNs for which the Client is a member. There may be one or more of these entries. These parameters are written by VPNmanager.

NOTE: The export RADIUS Users file created by VPNmanager contains no entries in the authentication password field. Consequently, after creating the file, you must edit it to add the authentication password field to each Client. Additionally, the security of cryptographic keys used to secure VPNs are not compromised during the VPNmanager-to-RADIUS transfer. All VPN keys are encrypted with Triple DES encryption (56-bit DES encryption for the DES only version of VPNmanager).

This completes the process for configuring RADIUS support. If any Clients are rekeyed, they must be re-exported to the RADIUS server to reflect the new key.

·Note: Telnet sends traffic, including the login password in the clear. Remember to disable telnet after you use it.

.

Chapter 12 Upgrading firmware and licenses

You can upgrade the VPNos firmware and license from the VPNmanager and set encryption strength and remote access for VSU100s.

Centralized firmware management

The VPNmanager centralized firmware management allows you to upgrade the firmware for one or many security gateways at one time. You can quickly verify the firmware release for any security gateway or VSU model. VPNmanager validates that the firmware image is correct before upgrading the device. The available firmware images are stored in the policy server.

Before upgrading the firmware using the centralized firmware management feature, you must download the latest firmware from Avaya Inc.

The security gateway firmware download is password-protected. Contact technical support at vpnsupport@avaya.com to request a password prior to beginning the download.

Read the latest security gateway product readme file, before beginning the upgrade. For the latest version of the file for all security gateways, go the VPN and Security page from the Avaya Support Technical Database Web site, at http://support.avaya.com, and select the security gateway type to be downloaded, follow the links to the Readme file.

Following are a few definitions that you should be familiar with prior to using the centralized firmware management feature:

- Device Inventory
  The device inventory is displayed when the Upgrade Firmware button is selected. The device inventory lists the name of the available devices to be upgraded, type of the device available for upgrade, current firmware version, and the available versions of firmware for the specific device.

- Firmware Library

  The firmware library list the devices and the available firmware versions for that specific device. The firmware library is a repository that is stored and maintained on the policy server. The various versions of firmware for the different devices are stored in the firmware library.

  Firmware versions can be added to the firmware library. Click the Add button to Browse to the firmware location and add to the firmware library repository on the policy server. Previous or older versions of the firmware can be deleted from the firmware library repository on the policy server.

- Upgrade Options
  The upgrade options are:

    - Skip devices that are up-to-date
      This option is the default setting. The devices that up-to-date will not display in the upgrade list. If a device should be downgraded, this option must be unchecked to view all devices in the upgrade list.

    - Prompt for reboot
      This option is not the default setting. All devices selected in the upgrade list to be upgraded will reboot when the upgrade is completed. All devices must be rebooted in order for the upgrade to take effect.

- Upgrade Devices

  The upgrade devices button activates the upgrade wizard. Use the upgrade wizard to walk you through the steps to upgrade using the centralized firmware management feature.

  Note: The upgrade devices wizard dose not allow downgrading of devices.

To upgrade the firmware using centralized firewall management:

1. From the configuration console, click the **Upgrade Firmware** button.

2. The Device Inventory dialog appears.

3. Select the **Upgrade Devices** button to begin the upgrade devices wizard.

4. Select the device(s) to be upgraded from the Available Devices column.

5. Click the **Move Left** button to move the selected devices into the Device(s) to Upgrade column.

6. Click **Next** to review pending device(s) upgrade.

7. Click **Upgrade** to complete the device(s) upgrade.

Device - Upgrade tab

The Upgrade tab provides access to security gateway upgrade facilities including firmware upgrades and optional feature activation. For devices with firmware version 4.2 or later, license files can be uploaded from the Upgrade tab.

Figure 85. Device Upgrade tab
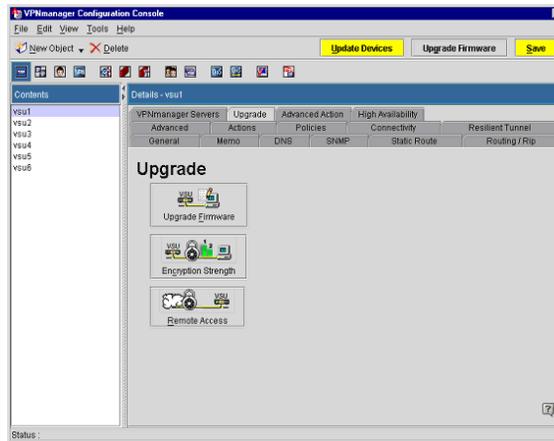
Upgrading a security gateway's firmware

Use the *Upgrade Firmware* button for upgrading the firmware of a specific security gateway. Before upgrading firmware from the VPNmanager, you must download the latest firmware from Avaya Inc.

The security gateway firmware download is password-protected. Contact technical support at vpnsupport@avaya.com to request a password prior to beginning the download.

Read the latest security gateway product readme file, before beginning the upgrade. For the latest version of the file for all security gateways, go the VPN and Security page from the Avaya Support Technical Database Web site, at http://support.avaya.com, and select the security gateway type to be downloaded, follow the links to the Readme file.

Note: Because the upgrade procedure removes the security gateway from service, firmware upgrades should be a scheduled maintenance activity.

To upgrade a security gateway's firmware:

1. Once you have received your password, go to the Avaya Support Technical Database Web page at http://support.avaya.com, click **VPN and Security** and select the appropriate security gateway type to download.

2. Click **Software Downloads** and follow the links. Click the *security gateway type* link to begin the download process.

3. Select **Save this file to disk**. Click **OK**.

4. Browse to the directory where the VPNos download files should be saved. Click **Save**.

5. Navigate to the directory where the VPNos file was saved.

6. Double-click the firmware zip file to begin extracting the VPNos image. The Password screen appears.

7. Enter the password from technical support.

8. Go to the VPNmanager Console, then move to the *Configuration Console* window.

9. Click **View>Device** to list all the security gateway in the *Contents* column.

10. From the **Contents** column, select the security gateway to upgrade.

11. Click the **Upgrade** tab, to bring it to the front.

12. Click **Upgrade Firmware**; the *Open* dialog box appears.

13. Navigate to the directory where the VPNos firmware image was saved.

14. Select the **update.bin** file.

15. Click **Open** to install the update.bin file.

16. When installation is complete, a message box appears asking if you want to reboot the security gateway.

    • If the subdirectory has an **upstage2.bin** file, click **NO**. Do not reboot the security gateway. You need to install the **upstage2.bin** file. Follow the instructions, starting from Step 9; in step 14, select the **upstage2.bin** file.

    • If the security gateway subdirectory *does not* have an **upstage2.bin** file, click **YES**. If you answered **YES** to rebooting the security gateway, your upgrade is complete.

17. Click **OK** to return to the VPNmanager Console.

18. The task summary is displayed.

19. Close the task summary window and check the security gateway status. The security gateway status should be success.

20. If you have not communicated with the target security gateway, the **security gateway logon screen** appears: enter your login credentials to complete the download.

21. When the download is finished, click **Reboot Device** to reboot the security gateway.

    Note: A security gateway takes at least two minutes to reboot.

License

Beginning with VPNos 4.2, you can obtain additional licenses to increase the number of remote users and site-to-site VPN connections that are allowed during a secure session.

When you purchase additional licenses, you receive a file with the encrypted information. This file is created based on the serial number of the security gateway and the number of licenses that are available on that security gateway. This file cannot be applied to another security gateway.

Use the License button to upload the licenses from the VPNmanager Console.

Once you have received the license file from your sales representative, upload the license file to the security gateway as follows:

1. Save the license file to a directory on the computer.

2. From the security gateway object Upgrade tab, click **License.**

3. Navigate to the directory where the license was saved and select the license file. Click **Open.**

4. Choose the security gateway for which the license needs to be updated. Click **OK**. The license is uploaded to the security gateway and the status of the upload is shown.

Encryption Strength

This button launches the Encryption Strength screen through which DES or 3DES encryption can be activated. When initially launched, the security gateway is polled for the current status of this feature, which is displayed on the first line (DES or 3DES). Click on the radio button for the desired encryption method. Click OK after you have selected the encryption.

Note: You are required to enter a valid license registration number to activate this option. Should you miskey the number, a "Registration number invalid" message appears. Other error messages may also appear if the security gateway is not reachable.

Remote Access (VSU-100 Only)

While the VSU-100 is designed for site-to-site operation, theVSU100R is provided with additional functionality to support dial-in VPNremote clients. To use this facility, Remote Access must be enabled.

Note: You are required to enter a valid license registration number to activate this option. Should you miskey the number, a "Registration number invalid" message appears. Other error messages may also appear if the VSU is not reachable.

Appendix A Using SSL with Directory Server

As an added benefit, all communications with the Directory Server can be secured by *SSL (Secure Sockets Layer)*.

In order to run SSL, a *Public Key Infrastructure (PKI)* is used for creating a *signed certificate* and an *issuer's certificate*. Both signed certificates are then installed on the server. The issuer's certificate is then installed in the VPNmanager Console and devices belonging to the VPN domain. The PKI can be owned and operated by a third party called a Certification Authority, or it can be owned and run by your organization. After the certificates are installed, the VPNmanager Console is started, and during login SSL services are started.

Figure 86. Installing Certificates for Running SSL
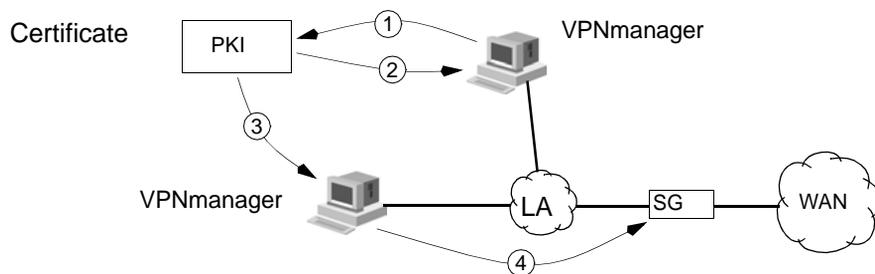
Explanation for Figure 86:

1. An administrator uses Directory Server to send a *Certificate Request* to a PKI.

2. The PKI responds with a *Signed Certificate*.

   The Issuer's Certificate is sometimes called a *Certificate Authority (CA) Certificate*, and can be freely obtained from anyone running a PKI.

3. An Issuer's Certificate is installed in the VPNmanager Console.

4. The administrator uses VPNmanager to install an *Issuer's Certificate* into the devices.

When to Configure Your VPNmanager for SSL

You can configure your VPNmanager to use SSL at anytime, however, it's recommended that it be done before being put into service.

Starting the iPlanet Directory console

The following two procedures describe how to start the *iPlanet server console*. One procedure is for servers running on Windows NT or Windows 2000, and the other is for Solaris OS computers. Start iPlanet Directory Server when you need to generate certificate requests and install certificates into the iPlanet Directory Server.

Starting iPlanet Directory Server from Windows NT or Windows 2000:

1. From the Windows Task bar, click **Start**>**Programs**>**iPlanet Server Products**, then select **iPlanet Console** to open the *iPlanet Console* dialog box.

2. Enter the VPNmanager DN and password that was configured during installation.

3. Click **OK** to open the *iPlanet Console*.

Starting iPlanet Directory Server from Solaris Sparc:

1. From the File Manager, move to the directory where iPlanet was installed. The default directory is //usr/iPlanet/servers.

2. Use the command *startconsole* to open the *iPlanet Console* dialog box.

3. Enter the VPNManager DN string and password that was configured during installation.

4. Click **OK** to open the *iPlanet Console*.

Installing a certificate into the server

Installing a signed certificate into iPlanet Directory Server involves six procedures:

- One: Generating a Certificate Request

- Two: Sending the Certificate Request

- Three: Installing the Certificate

- Four: Confirming the Install

- Five: Activating SSL

One: Generating a certificate request

To create a certificate request:

1. Start the *iPlanet Directory Console*.

2. From the *Navigation Tree*, select the directory server for your VPN.

3. Click **Open** to run the *Server Interface* window.

4. From the **Tasks** tab, click **Certificate Setup Wizard** to start the *Certificate Setup Wizard*.

5. Click **Next**.

6. From the **Select a Token** drop-down list, select **Internal (software)**.

7. Select the **No** radio button.

   Note: If using FORTEZZA, see the iPlanet Directory Server Help for instructions.

8. Click **Next**.

9. If a certificate database does not exist, do the following.

   - Click **Next**.

   - Enter a password, having at least one numeral and eight characters, then confirm it.

   - Click **Next** to continue.

10. Click **Next**.

---

11. If prompted, select a token from the list of legal tokens you can use, enter the password created in Step 9, then click **Next**.

12. Select the **New Certificate** radio button.

13. In the **CA Administrator's** address text box, type in your Certificate Authority's email address or URL.

14. Click **Next**.

15. In the **Your Name** text box, type in your *User ID*.

16. In the **Telephone** text box, type in your telephone number.

17. In the **Server Host Name** text box, type in the fully qualified domain name of the computer running the server, for example, *host_name.domain_name.com*.

18. In the **E-mail Address** text box, type in your e-mail address that will be used by the CA.

19. In the **Organization** text box, type in the legal name of your organization.

20. In the **Organizational Unit** text box, type your organization's department name.

21. In the **Locality** text box, type in the city name where your organization is located.

22. In the **State or Province** text box, type in the full name of the state or province where your organization is located.

23. In the **Country** drop-down list, select the country where your organization is located.

24. Click **Next** and enter your password.

25. Click **Next** and your certificate request will appear in the dialog box.

26. Click **Cancel** to exit the wizard. Your work will not be destroyed.

Now you can send your request to a Certificate Authority.

Two: Sending the certificate request

If you are running the server on a Solaris OS computer, the certificate request Sendmail.

If you are using Windows NT, the certificate information is automatically saved to a file in the *C:\temp* directory. Use the following procedure to send the request to a CA.

To send a certificate request to a certification authority:

1. Use a common e-mail program to create a new message.

2. Start Notepad.

3. Use Notepad to open the certificate request, *mailtemp.1*, that is located in the *C:\temp* directory.

4. From Notepad, copy *everything* below the **Reply-to Address**, then paste it to the body of your new message.

5. From Notepad, copy the **To Address**, then paste into the **To** text box of your new message.

6. From Notepad, copy the **Subject**, then paste into the **Subject** text box of your new message.

7. Send the message to the CA.

When the CA responds with a certificate, save the information to a text file. Backup the file to a safe location. You are now ready to install the certificate.

Three: Installing the certificate

To install a signed certificate into iPlanet Directory Server:

1. Start the *iPlanet Directory Console*.

2. From the *Navigation Tree*, select the directory server for your VPN.

3. Click **Open** to run the *Server Interface* window.

4. From the **Tasks** tab, click **Certificate Setup Wizard** to start the *Certificate Setup Wizard*.

5. Click **Next**.

6. From the **Select Token** drop-down list, choose **Internal (software)**.

7. From the **Is the server certificate already requested and ready to install** options, select **Yes**.

8. Click **Next**.

9. Click **Next** again.

10. In the **Certificate For** options, select **This Server**.

11. In the **Password** text box, type in the password created in Step 9 of page 333.

12. Click **Next**.

13. Do one of the following.

    • In the **The certificate is located in this file** text box, type in the path to the certificate file.

    • Copy the certificate from the certificate file, then paste it into the **The certificate is located in the following text field** text box. Only copy the header, footer, and body of the certificate as shown by Figure 87.

Figure 87. An Example of a Signed Certificate

Header

-----BEGIN CERTIFICATE-----

nfi897rho987fb+mht>,oi$s25hgj98iJop)kjh

GrDfgyui987jg55dJ99KJY6%$3@@Sd5()~

43dbi0oMl=_+;mhjuuhJ8*&tfeEckiooplkjghf

Footer

14. Click **Next** to install the signed certificate.

15. Click **Add** to add the signed certificate to the iPlanet Directory Server.

16. Click **Done** to complete the certificate installation.

Now that the certificate is installed, you need to configure the sever to trust the Certificate Authority.

Four: Confirming the install

To confirm the install:

1. Start the *iPlanet Directory Console*.

2. From the *Navigation Tree*, select the directory server for your VPN.

3. Click **Open** to run the *Server Interface* window.

4. From the **Tasks** tab, click **Manage Certificates** to open the *Certificate Management* dialog box.

5. Scroll through the list to find the certificate named *Server-Cert* that you installed.

If you found your certificates, your server is ready for SSL activation.

Five: Activating SSL

To activate SSL on the server

1. Start the *iPlanet Directory Console*.

2. From the *Navigation Tree*, select the directory server for your VPN.

3. Click **Open** to run the *Server Interface* window.

4. Click the **Configuration** tab to bring it to the front.

5. From the right-hand pane, click the **Encryption** tab to bring it to the front.

6. Select the **Enable SSL** check box.

7. Select the Check box next to the cipher family you want to use.

8. From the **Token** drop-down list, select the token you want the server to use.

9. From the **Certificates** drop-down list, select the certificate named *Server-Cert* that you want to use.

10. Click **Save**.

After activating SSL, you can install the Issuer's Certificate into the VPNmanager Console.

Installing the Issuer's Certificate in the VPNmanager Console

Installing an Issuer's Certificate into VPNmanager Console is done from the command line. The same Issuer's Certificate that was installed in the server can be used here. Since the console can run on Windows NT or Solaris OS, the following two sub sections cover the procedures.

Note: After a certificate is installed, it cannot be seen in the Issuer Certificates list of the Policy Manager for Issuer Certificates.These certificates are specifically used for running SSL services, not for anything else.

Windows NT and Windows 2000 Computers

To install a certificate in VPNmanager Console:

1. Copy the certificate to the *C:\Program Files\Avaya\VPNmanager\ Console* directory.

2. From the task bar, click **Start** >**Run** to open the *Run* dialog box.

3. In the **Open** text box, type the following command line to install the certificate. The *filename* is a name of the certificate file, and *aliasname* is the alias you choose for the certificate file.

   C:\Program Files\AvayaVPNmanager\Console importcert ***aliasname filename***

   The DOS window will appear containing a message confirming the install.

To view all the installed issuer's certificates:

1. From the Task bar, click **Start** then select **Run** to open the *Run* dialog box.

2. In the **Open** text box, type the following command line to view all installed certificates.

   C:\Program Files\AvayaVPNmanager\ Console listcert

   The DOS window will appear listing all the certificates.

To delete an installed issuer's certificates:

1. From the Windows NT Taskbar, click **Start** then select **Run** to open the *Run* dialog box.

2. In the **Open** text box, type the following command line to view all installed certificates, where *aliasname* is the alias you gave the certificate when it was installed.

    C:\Program Files\AvayaVPNmanager\ Console deletecert **aliasname**

Solaris OS Computers

To install a certificate in VPNmanager Console:

1. Copy the certificate to the opt/Avaya/VPNmanager/Console directory.

2. Open a Console window.

3. Move to the opt/Avaya/VPNmanager/Console directory.

4. Type in the following command to install the certificate. The *filename* is a name of the certificate file, and *aliasname* is the alias you choose for the certificate file.

    sh importcert.bat *aliasname filename*

To view all the installed issuer's certificates:

1. Open a Console window.

2. Move to the opt/Avaya/VPNmanager/Console directory.

3. Type in the following command to list all installed issuer certificates.

    sh listcert.bat

To delete an installed issuer's certificates:

1. Open a Console window.

2. Move to the opt/Avaya/VPNmanager/Console directory.

3. Type the following command line to view all installed certificates, where *aliasname* is the alias you gave the certificate when it was installed.

    sh deletecert.bat **aliasname**

Installing the Issuer's Certificate into a security gateway

To create a Device object, refer to Chapter 3, Setting up the network. Once the Device object has been created, perform the following procedure.

To install the issuer's certificate into a security gateway:

1. From the **Tools** menu, select **Policy Manager** to open the *Policy Manager Window.*

2. From the **Object Name** list, select the Device object that you just created.

3. From the **Type of Policy** list, select **Issuer Certificates** to open the *Policy Manager for Issuer Certificates.*

4. From the **Issuer Certificates** list, select a row where the new issuer certificate will be installed.

5. Click **Add** to open the *Open* dialog box.

6. Use the **Look in** list for navigating to the location of the *Issuer Certificate.*

7. Select the *Issuer Certificate*, then click **OK** to return to the *Policy Manager* window.

   After the device has received the Issuer Certificate, the certificate appears in the *Issuer Certificates* list.

8. Close the window.

Repeat Step 1 through Step 7 for each device that needs to have an Issuer's Certificate installed.

Note: The certificates and procedures involved in this appendix are not related to creating a certificate based VPN. They are only for securing the communications between the VPNmanager Console, Directory Server, and the device. For information about certificate based VPNs, see Chapter 7, Configuring VPN objects.

Appendix B Firewall rules template

General

The security gateway contains a powerful multi-layer inspection engine to provide extensive filtering capabilities, essential for a full-time connection to the Internet. You can configure your own rules, but, as a convenience in setting up the Firewall on the security gateway, predefined general firewall rules (templates) can be selected to protect the public, private, semi-private, DMZ, and maintenance zones.

These predefined firewall rules are grouped into security levels of high, medium, and low. One firewall security level is applied to the security gateway, and the rules for each zone are enforced according to the type of zone being protected. How the template rules are applied to a zone are described in this appendix.

The Firewall engine uses a rule-based method of packet filtering, where the priority of the rule is determined by its position in the list (first is highest priority).

Note: The *common services* referred to in this appendix include all of the following:

- Ping
- FTP control, Passive Data FTP
- SSH, TELNET
- HTTP, HTTPS
- POPS, IMAP, SMTP, and NNTP

High Security. Selecting high security enforces a set of rules that try to protect the security gateway itself and the internal network zones. For high security the following policy is defined:

- Private networks and management networks are considered internal networks, and can initiate connections to access common services on the Internet.

- Except for access to the DMZ zone, traffic initiated from the Internet is denied.

- VPN outgoing and incoming traffic is allowed.

- DMZ common services can be accessed from all interfaces. The DMZ network cannot initiate any traffic.

- The semi-private zone is not considered completely trusted. Access from semi-private to private zones is allowed only if it is VPN traffic. All other incoming traffic is blocked.

Medium Security. Selecting medium security enforces the same security policy as high security for all zones except the semi-private zone. The semi-private zone with medium security is trusted the same as the private zone. That is, the same security policy that is enforced on the private zone is enforced on the semi-private zone. In medium security, semi-private zone can also access all the resources in the private zone.

Low Security. Selecting low security enforces the same security policy as specified for medium and the access from the internal network to the Internet is not limited to only the common services. Access to all TCP and UDP services are allowed.

The details about rules and what type of traffic is allowed and denied for each level and zone are in the tables that follow.

Public zone firewall templates

The public network interface provides connection to the Internet and the security gateway functions as the firewall/VPN gateway.

Usually the public interface has the strongest firewall policy. Few incoming packets are allowed and outgoing packets are allowed only for commonly used services.

The public high security rules are enforced for both incoming and outgoing packets as follows.

Incoming traffic to the public zone allowed include:

- VPN packets from private, DMZ, Management or Semi-private zones

- ICMP unreachable packets

- Publicly accessible DMZ services allowed include ping, FTP, SSH, Telnet, HTTP, HTTPS, POP3, IMAP, SMTP, NNTP and DNS.

All other incoming traffic is blocked.

Outgoing traffic from the public zone allowed include:

- Outgoing VPN traffic

- ICMP unreachable

- Ping from any IP to any

- DNS from any IP to any

- Common services originating from all internal networks, private, DMZ, management and semi-private.

All other outgoing traffic is blocked.

The medium security policy for the public zone is the same as that of the high security policy.

The low security policy allows all the traffic allowed for medium security. In addition, all TCP, UDP packets from all networks are allowed to go out.

Table 30. Public high and medium security firewall rules

Rule Name	Action	Source	Destination	Service	Direction	Zone	Keep State	Description
InBoundPublicAccess	Permit	Any	PublicIP	IKE-IN IKE-AVAYA-IN IPSEC-NAT-T-IN AH/ESP ICMPDEST UNREACHABLE	In	Public	no	Permit incoming VPN traffic and ICMP unreachable packet

Table 30. Public high and medium security firewall rules (Continued)

Rule Name	Action	Source	Destination	Service	Direction	Zone	Keep State	Description
InBoundPublictoDMZAccess	Permit	Any	DMZNet	ICMPECHO REQUEST SSH/TELNET FTP-CTRL PASSIVEFTP HTTP/HTTPS DNS-TCP/DNS-UDP NETBIOS-NS-TCP/UDP NETBIOS-DGM-TCP/UDP NETBIOS-SSN-TCP/UDP POP3/IMAP/SMTP NNTP	In	Public	Yes	Permit incoming traffic to DMZ network
InBoundPublicBlockAll	Deny	Any	Any	ANY	In	Public	No	Deny the rest of traffic
OutBoundPublicAccess	Permit	PublicIP	Any	IKE-OUT IKE-AVAYA-OUT IPSEC-NAT-T-OUT AH/ESP ICMPDEST UNREACHABLE	Out	Public	no	Permit outgoing VPN traffic

Table 30. Public high and medium security firewall rules (Continued)

Rule Name	Action	Source	Destination	Service	Direction	Zone	Keep State	Description
OutBoundPublicGeneralAccess	Permit	Any	Any	ICMPECHO REQUEST SSH/ TELNET FTP-CTRL PASSIVEFTP HTTP/ HTTPS DNS-TCP/ DNS-UDP NETBIOS-NS-TCP/ UDP NETBIOS-DGM-TCP/ UDP NETBIOS-SSN-TCP/ UDP POP3/ IMAP/SMTP NNTP	Out	Public	Yes	Permit traffic with the services to go out. The traffic can come from any network.
OutboundPublicActiveFTPActive	Permit	DMZNet	Any	ActiveFTP	Out	Public	Yes	Permit active FTP data connection from FTP server on DMZNet to any FTP client on INATERNET
OutboundPublicNATedFTPActiveFTPActive	Permit	PublicIP	Any	DYNAMICPORTS	Out	Public	Yes	Permit NAT'ed active FTP data connection from FTP server on DMZNet to any FTP client on INATERNET
OutBoundPublicBlockAll	Deny	Any	Any	Any	Out	Public	No	Deny the rest of traffic

Rule Name	Action	Source	Destination	Service	Direction	Zone	Keep State	Description
InBoundPublicAccess	Permit	Any	PublicIP	IKE_IN IPSEC_NAT_T_IN AH/ESP ICMPDestUnreach	In	Public	no	Permit incoming VPN traffic and ICMP unreachable packet
InBoundPublictoDMZAccess	Permit	Any	DMZNet	ICMPEchoReq(PING) FTP-Ctrl/PassiveFTP SSH/TELNET HTTP/HTTPS DNS-TCP/DNS-UDP POP3/IMAP/SMTP NNTP	In	Public	Yes	Permit incoming traffic to DMZ network
InBoundPublicBlockAll	Deny	Any	Any	Any	In	Public	No	Deny the rest of traffic
OutBoundPublicAccess	Permit	PublicIP	Any	IKE_OUT IPSEC_NAT_T_OUT AH/ESP ICMPDestUnreach	Out	Public	no	Permit outgoing VPN traffic
OutBoundPublickPingAccess	Permit	DNZNet PrivateNet SemiPrivateNat ManagementNet	Any	ICMPEchoRequest	Out	Public	Yes	Permit outgoing ping access.
OutBoundPublicDNSAccess	Permit	PublicIP DMZNet PrivateNet SemiPrivateNet ManagementNet	Any	DNS-TCP DNS-UDP	Out	Public	Yes	Permit outgoing DNS access.
OutBoundPublicGeneralAccess	Permit	Any	Any	ICMPEchoReq(PING) FTP-Ctrl/PassiveFTP SSH/TELNET HTTP/HTTPS DNS-TCP/DNS-UDP POP3/IMAP/SMTP	Out	Public	Yes	Permit traffic with the services to go out. The traffic can come from any network.
OutBoundPublicBlockAll	Deny	Any	Any	Any	Out	Public	No	Deny the rest of traffic

Table 31. Public low security firewall rules

Rule Name	Action	Source	Destination	Service	Direc-tion	Interface	Keep State
InBoundPublicAccess	Permit	Any	PublicIP	IKE_IN IPSEC_NAT_T_IN AH/ESP ICMPDestUnreach	In	Public	no
InBoundPublictoDMZAccess	Permit	Any	DMZNet	HTTP/HTTPS POP3/IMAP/SMTP	In	Public	Yes
InBoundPublicBlockAll	Deny	Any	Any	Any	In	Public	No
OutBoundPublicAccess	Permit	PublicIP	Any	IKE_OUT IPSEC_NAT_T_OUT AH/ESP ICMPDestUnreach	Out	Public	no
OutBoundPublicPingAccess	Permit	PublicIP DMZNet PrivateNet SemiPrivateNet ManagementNet	Any	ICMPEchoRequest	Out	Public	Yes
OutBoundPublicGeneralAccess	Permit	Any	Any	ICMPEchoRequest(PING) ALL TCP ALL UDP	Out	Public	Yes
OutBoundPublicBlockAll	Deny	PublicIP DMZNet PrivateNet SemiPrivateNet ManagementNet	Any	Any	Out	Public	No

Private zone firewall templates

The private network interface provides connection to the private/corporate LAN. Private zones are considered trusted networks and because of this most traffic is allowed.

The private high security rules are enforced for both incoming and outgoing packets as follows.

Any incoming traffic from the private zone is allowed except traffic that is destined to the management zone.

For outgoing traffic to the private zone, traffic initiated from DMZ is strictly denied. All other traffic is allowed.

The private medium security rules and the low security rules are the same as the private high security rules.

Table 32. Private high security firewall rules

Rule Name	Action	Source	Destination	Service	Direction	Zone	Keep State	Description
InBoundPrivateToMgmtDenyAccess	Deny	Any	ManagementNet	Any	In	Private	No	Traffic to ManagementNet is denied.
InBoundPrivatePermitAll	Permit	Any	Any	Any	In	Private	Yes	Permit VI/VMGR and VP, clear traffic to PUBLIC
OutBoundPrivateDMZSemiPriDenyAccess	Deny	DMZNet	Any	Any	Out	Private	No	Deny traffic from DMZNet and SemiPrivateNet
OutBoundPrivatePermitAll	Permit	Any	Any	Any	Out	Private	Yes	Permit incoming VPN

Table 33. Private medium security firewall rules

Rule Name	Action	Source	Destination	Service	Direction	Zone	Keep State	Description
InBoundPrivateDenyAccess	Deny	Any	ManagementNet	Any	In	Private	No	Traffic to ManagementNet is denied.
InBoundPrivatePermitAll	Permit	Any	Any	Any	In	Private	Yes	Permit WI/VMGR and VPN, clear traffic to PUBLIC
OutBoundPrivateDenyAccess	Deny	DMZNet	Any	Any	Out	Private	No	Deny traffic from and SemiPrivate Net
OutBoundPrivatePermitAll	Permit	Any	Any	Any	Out	Private	Yes	Permit incoming VPN

Table 34. Private low security firewall rules

Rule Name	Action	Source	Destination	Service	Direction	Zone	Keep State	Description
InBoundPrivateDenyAccess	Deny	Any	ManagementNet	Any	In	Private	No	Traffic to ManagementNet is denied.
InBoundPrivatePermitAll	Permit	Any	Any	Any	In	Private	Yes	Permit WI/VMGR and VPN, clear traffic to PUBLIC
OutBoundPrivateDenyAccess	Deny	DMZNet	Any	Any	Out	Private	No	Deny traffic from and SemiPrivateNet
OutBoundPrivateDenyAll	Permit	Any	Any	Any	Out	Private	Yes	Permit incoming VPN

Semi-private zone firewall templates

A semi-private network interface provides connection to a network whose equipment can be made physically secure, but whose medium is vulnerable to attack (such as a Wireless network used within a corporation's Private network infrastructure).

Because wireless connections cannot be easily controlled, strict firewall policy should be enforced on the semi-private interface to limit the access from the semi-private zone to VPN traffic. Clear traffic to Private and Management zones is not allowed. Common services to DMZ are allowed and clear traffic to Public is allowed.

The semi-private high security rules are enforced for both incoming and outgoing packets as follows.

Incoming traffic to the semi-private zone allowed includes:

- VPN traffic. The VPN tunnel endpoints could be semi-private IP or Public IP.

- Ping, DNS

- ICMP unreachable packets

The following clear traffic is allowed

- The source is semi-private and the destination is DMZ servers, with the following common services: PING, FTP control, Passive Data FTP, SSH, Telnet, HTTP, HTTPs, POP3, IMAP, SMTP, and NNTP.

- The destination is Public and the services are FTP, SSH, Telnet, HTTP, HTTPS, POP3, IMAP, or ICMPechorequest.

All other incoming traffic is blocked.

Outgoing traffic to the semi-private zone that is allowed includes

- Any allowed traffic from other zones

- VPN traffic

Table 35. Semi-private high security firewall rules

Rule Name	Action	Source	Destination	Service	Direction	Zone	Keep State	Keep State
InBoundSemiPrivateVPNAccess	Permit	Any	SemiPrivateIP PublicIP	IKE_IN IPSEC_NAT_T_IN AH/ESP ICMPDestUnreach	In	SemiPrivate	No	Permit incoming VPN and ICMP unreachable
InBoundSemiPrivatePingAccess	Permit	Any	SemiPrivateIP PublicIP	ICMPEchoReq(PING)	In	SemiPrivate	Yes	Permit incoming PING
InBoundSemiPrivatetoDMZAccess	Permit	Any	DMZNet	ICMPEchoReq(PING) FTP-Ctrl/PassiveFTP SSH/TELNET HTTP/HTTPS DNS-TCP/DNS-UDP POP3/IMAP/SMTP NNTP	In	SemiPrivate	Yes	Permit incoming services to DMZNet
InBoundSemiPrivateDenyAccess	Deny	Any	DMZNet PrivateNet ManagementNet SemiPrivateIP	Any	In	SemiPrivate	No	Deny traffic to PrivateNet, ManagementNet and DMZNet
InBoundSemiPrivatetoPublicAccess	Permit	Any	Any	ICMPEchoReq(PING) FTP-Ctrl/PassiveFTP SSH/TELNET HTTP/HTTPS DNS-TCP/DNS-UDP POP3/IMAP/SMTP NNTP	In	SemiPrivate	Yes	Permit clear traffic to Public network/ VPN traffic with Public IP as tunnel endpoint
InBoundSemiPrivateBlockAll	Deny	Any	Any	Any	In	SemiPrivate	No	Deny the rest of traffic
OutBoundSemiPrivateVPNAccess	Permit	SemiPrivateIP PublicIP	Any	IKE_OUT IPSEC_NAT_T_OUT AH ESP ICMPDestUnreach	Out	SemiPrivate	No	Permit outgoing VPN traffic.
OutBoundSemiPrivatePermitAll	Permit	Any	Any	Any	Out	SemiPrivate	Yes	Permit everything with Keep state. (For any traffic initiated from Private/ ManagementNET)

Table 36. Semi-private medium security firewall rules

Rule Name	Action	Source	Destination	Service	Direction	Zone	Keep State	Description
InBoundSemiPrivateDenyAccess	Deny	Any	ManagementNet	Any	In	SemiPrivate	No	Traffic to ManagementNet is denied.
InBoundSemiPrivateVPNAccess	Permit	Any	SemiPrivateIP PublicIP	IKE_IN IPSEC_NAT_T_IN AH/ESP ICMPDestUnreach	In	SemiPrivate	no	Permit incoming VPN traffic and ICMP unreachable packet
InBoundSemiPrivatePermitAll	Permit	Any	Any	Any	In	SemiPrivate	Yes	Permit WI/VMGR and VPN, clear traffic to PUBLIC
OutBoundSemiPrivateDenyAccess	Deny	DMZNet	Any	Any	Out	SemiPrivate	No	Deny traffic from DMZNet
OutBoundSemiPrivateVPNAccess	Permit	SemiPrivateIP PublicIP	Any	IKE_OUT IPSEC_NAT_T_OUT AH/ESP ICMPDestUnreach	Out	SemiPrivate	no	Permit outgoing VPN traffic
OutBoundSemiPrivateDenyAll	Permit	Any	Any	Any	Out	SemiPrivate	Yes	Permit incoming VPN

Table 37. Semi-private low security firewall rules

Rule Name	Action	Source	Destination	Service	Direction	Zone	Keep State	Description
InBoundSemiPrivateDenyAccess	Deny	Any	ManagementNet	Any	In	Semi Private	No	Traffic to ManagementNet is denied.
InBoundSemiPrivateVPNAccess	Permit	Any	SemiPrivateIP PublicIP	IKE_IN IPSEC_NAT_T_IN AH/ESP ICMPDestUnreach	In	Semi Private	no	Permit incoming VPN traffic and ICMP unreachable packet
InBoundSemiPrivatePermitAll	Permit	Any	Any	Any	In	Semi Private	Yes	Permit WI/VMGR and VPN, clear traffic to PUBLIC

Table 37. Semi-private low security firewall rules (Continued)

OutBoundSemiPrivateDenyAccess	Deny	DMZNet	Any	Any	Out	Semi Private	No	Deny traffic from DMZNet
OutBoundSemiPrivateVPNAccess	Permit	SemiPrivateIP PublicIP	Any	IKE_OUT IPSEC_NAT_T_OUT AH/ESP ICMPDestUnreach	Out	Semi Private	no	Permit outgoing VPN traffic
OutBoundSemiPrivateDenyAll	Permit	Any	Any	Any	Out	Semi Private	Yes	Permit incoming VPN

DMZ zone firewall templates

The Demilitarized Zone (DMZ) network interface is typically used to allow Internet users access to some corporate services without compromising the private network where sensitive information is stored. For all the services setup in the DMZ, access is allowed from any network, including Public, Private, Management and Semi-private. Because the DMZ is not a trusted network, all outgoing traffic is blocked.

The same security rules are enforced for high security, medium security, and low security. The DMZ high security rules are enforced for both incoming and outgoing packets as follows.

Incoming traffic from the DMZ zone is denied.

Outgoing traffic to the DMZ zone allowed includes

• Packets from the following networks: private, management, semi-private, and the destination is the servers with the common services.

Table 38. DMZ high and medium security firewall rules

Rule Name	Action	Source	Destination	Service	Direction	Zone	Keep State	Description
InBoundDMZActiveFTPAccess	Permit	DMZNet	Any	ActiveFTP	In	DMZ	Yes	Permit active FTP data connection from FTP server on DMZNet to any FTP client on INATERNET(this works for both NAT/Non NAT setup)
InBoundDMZBlockAll	Deny	Any	Any	Any	In	DMZ	No	Deny the rest of traffic
OutBoundDMZAccess	Permit	Any	DMZNet	ICMPECHOREQUEST SSH/TELNET FTP-CTRL PASSIVEFTP HTTP/HTTPS DNS-TCP/DNS-UDP NETBIOS-NS-TCP/UDP NETBIOS-DGM-TCP/UDP NETBIOS-SSN-TCP/UDP POP3/IMAP/SMTP NNTP	Out	DMZ	Yes	Permit outgoing traffic with common services
OutBoundDMZBlockAll	Deny	Any	Any	Any	Out	DMZ	No	Deny the rest of the traffic

Table 39. DMZ low security firewall rules

Rule Name	Action	Source	Destination	Service	Direction	Zone	Keep State	Description
InBoundDMZBlockAll	Deny	Any	Any	Any	In	DMZ	No	Deny the rest of traffic
OutBoundDMZAccess	Permit	Any	DMZNet	ICMPEchoReq(PING) FTP-Ctrl/PassiveFTP SSH/TELNET HTTP/HTTPS DNS-TCP/DNS-UDP POP3/IMAP/SMTP NNTP	Out	DMZ	Yes	Permit outgoing traffic with the services
OutBoundDMZBlockAll	Deny	Any	Any	Any	Out	DMZ	No	Deny the rest of the traffic

Management zone security

Management interface connection can be configured to simplify network deployments to eliminate enterprise network dependencies on switches or routers.

The Management zone is a trusted network similar to the Private zone. Outgoing traffic is allowed, but incoming traffic is restricted. Only traffic initiated by the security gateway is allowed.

High, medium and low security rules are the same.

Incoming

All traffic is allowed to come in from the management network.

Outgoing

Only packets from the Management IP to the Management zone are allowed.

Table 40. Management high, medium, and low security firewall rules

Rule Name	Action	Source	Desti-nation	Service	Direct-ion	Zone	Keep State
InBoundManagementInterface PermitAccess	Permit	Any	ManagementIP	Any	In	Management	No
InBoundManagementPermitAll	Permit	Any	Any	Any	In	Management	Yes
OutBoundManagementInterfa ceAccess	Permit	Managem entIP	Any	Any	Out	Management	No
OutBoundManagementBlockA ll	Deny	Any	Any	Any	Out	Management	No

Glossary

A

Aggressive mode
An IKE mechanism used in the first phase of establishing a security association. Aggressive mode accomplishes the same authentication negotiating goal between clients as Main mode but faster (three packets versus six).

AH/ESP
In an IPSec packet, the Authentication Header (AH) and Encapsulation Security Payload (ESP) header. IKE VPNs authenticate IP packets using either an ESP header as defined in draft-ietf-ipsec-esp-v2-03.txt, or AH as defined in IETF draft-ietf-ipsec-auth-header-04.txt.

Alarms
When a security gateway in the VPN reports an alarm condition, details about the alarm including type, timestamp, and the originating security gateway can be found in the VPNmanager main screen Alarm pane.

Authentication

Generic
The process of ensuring that the data received is the same data that was sent from the source.

Local
Local Authentication is used in non-dynamic VPNs (VPNs not using RADIUS or a directory server (LDAP) as the authentication database). Here, the user is authenticated from the database stored in the security gateway's flash memory.

RADIUS
RADIUS Authentication uses an external RADIUS server and database for user authentication.

LDAP
LDAP Authentication uses the designated directory server database for user authentication.

B

Brute Force Attack
A hack attack that attempts to recover a cryptographic key by trying all reasonable possibilities.

C

CCD
Client Configuration Download. The protocol used to download the VPN session parameter configuration file from the security gateway to the remote client as part of a successful authentication when the security gateway is configured for Local Authentication.

Certificate Authority
A trusted company or organization that serves as a repository of digital certificates. Once a CA accepts your public key (with some other proof of identity), others can then request verification of your public key.

Certificates
Issuer
Issuer Certificates also reside in the security gateway and are used to authenticate the other side. For example, if the Directory Server presents a certificate for an SSL session, the security gateway must have an Issuer Certificate that can verify the VPNmanager's certificate is valid. Devices wishing to use IKE must be validated with an Issuer Certificate. All Issuer certificates are public.
My Certificates
My Certificates is a list of nine (0 through 8) certificates that exist inside the security gateway and are used to identify the security gateway to an opposite endpoint. Requires generation of a public/private key pair where the private key never leaves the security gateway.
Signing
Similar to the security gateways Issuer Certificates necessary to verify the VPNmanager Certificate, the Signing Certificates are for the VPNmanager Console to verify the security gateway Certificate.

Certificate Revocation List (CRL), checking
Certificate Revocation List checking looks to a directory server (maintained by CAs) to validate a new certificate by searching a list of no longer valid digital certificates.

D

DCI
Direct Configuration Interface is a Avaya Inc. proprietary protocol developed to facilitate passing setup and configuration data between the VPNmanager console and the security gateway. DCI traffic can pass in the clear if the LAN on which they both reside is behind a firewall, or over SSL if not.

DES

Data Encryption Standard (DES) is a block-cipher algorithm created by IBM used to rapidly encrypt large amounts of data at one time. The technique uses a 56-bit key and operates on blocks of 64 bits. SeeTriple DES.

Diffie-Hellman

A popular mechanism used to define the mathematical parameters used during IKE negotiations. Group 1 specifies use of a 768 bit modulus, Group 2 a 1024 bit modulus (Group 2 is "more secure").

Digital Certificate

An electronic document used to establish a company's identity by verifying its public key. Digital Certificates are issued by a certificate authority.

Domain Name Service (DNS)

The network service that converts text-based names into numeric IP addresses and vice-versa.

Domains, VPN

A VPN Domain is a collection of Virtual Private Network devices that compose a Virtual Private Network.

Dynamic VPNs

Dynamic VPNs are VPNs that can be readily scaled as dictated by business demands. As the remote client user population grows, the authentication and session configuration information for each new user must necessarily also grow. By maintaining this information not in the security gateway's flash memory but on a dedicated network host device, the number of users becomes unlimited. Two techniques of achieving this functionality normally used are LDAP or RADIUS.

Dyna Policy

An Avaya VPN term relating to a dynamic configuration download of VPN session security parameters to the remote client computer upon connection to a security gateway. This technique assures maximum security in a VPN session.

E

Encapsulation

The process of placing the contents of one packet into that of payload of another packet.

Extranet security gateway

It is possible to create a Group associated with a security gateway that is not managed by your company's VPNmanager. This happens when creating "extranets," or VPNs between partner corporations. In an extranet, each corporate network uses VPN components that are managed separately by each company's system administrator.

F

Firewall

A network device acting as a filter to restrict access to private network resources from the public. Filtering typically is based on the types of packets exchanged between two devices on the network.

H

Heartbeat

A special VPN packet broadcast by a primary security gateway used to facilitate the resilient tunnel function.

IKE (Internet Key Exchange)

A key-management protocol, IKE defines procedures and packet formats to establish, negotiate, modify and delete Security Associations (SAs) and defines payloads for exchanging key generation and authentication data. These formats provide a consistent framework for transferring key and authentication data which is independent of the key generation technique, encryption algorithm and authentication mechanism. Now combined with Oakley to form IKE.

IP Groups

IP Groups are a convenient means of managing your VPN resources. IP Groups are collections of IP network mask pairs associated with security gateways, hosts, and workstations located behind the security gateway.

IPSec

The network cryptographic protocols for protecting IP packets.

ISAKMP

The key-management protocol used in conjunction with IPSec.

Issuer Certificates

See Certificates, Issuer

L

LAN

Local Area Network

LDAP

Lightweight Directory Access Protocol is a simplified version of the standard X.500 distributed directory model standard. LDAP specifies how a client accesses a directory server. LDAP has emerged as a favored protocol since it also handles key management with key and certificate storage.

Lifetime, Key

Payload key lifetime defines the extent to which a single set of cryptographic keys is used when applying VPN services to IP packets. Key lifetimes can be defined by either the amount of data acted on by this single set of cryptographic keys or the amount of time these keys are used before a key change. The more often a key is changed, the "more secure" the system, although performance may be affected by frequent key changes.

LZS

Lempel-Ziv-Stac, a compression algorithm.

M

Mask Pairs

A network address and network mask. Two 4-byte pairs. For example, 1.1.1.0 and 255.255.255.0.

MIB - Enterprise

The enterprise-specific Management Information Base in the Avaya Inc. security gateways. The Enterprise MIB information allows the administrator to obtain basic monitoring information such as the network table, packet counter, and general information regarding the security gateway using third party software.

MIB-II (Non-Enterprise)

The non-enterprise specific Management Information Base in the Avaya Inc. security gateways. The MIB-II allows the administrator to obtain basic monitoring information such as device ethernet information, routing and ARP tables, SNMP traps, packet statistics, and other general information regarding the security gateway using third party software.

Migration

A utility by which an existing VPNmanager database is converted into an LDAP database for compatibility with VPNmanager 3.0 or later.

My Certificates

See Certificates, My Certificates

N

NAT

Network Address Translation (NAT) is a mechanism that allows private (non-routable) networks to connect to public (routable) networks.

Not My security gateway

If you are creating an extranet, choose "Not My security gateway" as the Group's associated security gateway. Doing this enables the "IP Address of Extranet security gateway" entry field. Enter the IP address of the your partner company's security gateway. This is required if any VPNs serviced by a VSU-1100, VSU-1010 or VSU-10 are in tunnel mode.

O

Oakley

A key exchange protocol used in IPSec as part of the Internet Key Exchange protocol.

P

Packet Filter

Hardware or software mechanism used in firewalls to discards packets based on the contents of the packet headers.

Perfect Forward Secrecy

Perfect Forward Secrecy defines a parameter of ISAKMP in which disclosure of long-term secret keying material does not compromise the secrecy of the exchanged keys from previous communications. Enabling Perfect Forward Secrecy is "more secure". See the IETF draft-ietf-ipsec-oakley-02.txt for more information on Perfect Forward Secrecy.

PKI

Public Key Infrastructure is the organization of certificate issuers and certificate management processes.

Preshared Secret

Preshared Secret is the simplest key management method used to construct a VPN. Authentication key exchanges between security gateways in the VPN are based on a single pre-shared secret known to all security gateways.

Public Key Certificate

A special block of data used to identify the owner of a particular public key. It describes the value of a public key, the key's owner, and the digital signature of the issuing authority.

R

RADIUS

Remote Authentication Dial In User Service is a client/server remote user authentication protocol in widespread use.

Resilient Tunnel

A mechanism of providing automatic backup of a secure tunnel between two endpoints. In practical application, a primary security gateway sends a "heartbeat" packets to a secondary security gateway every few seconds (configurable). Should the primary security gateway fail, the secondary security gateway will stop receiving the heartbeat packets. When this happens, the secondary security gateway switches over and takes on the role of primary security gateway.

S

SA

Security Association is an IPSec agreement between to communicating devices on which authentication and encryption algorithms (including key lifetimes) are used.

Session Key

A cryptographic key that has a finite life expectancy, typically for a single session.

Signing Certificates

See Certificates, Signing

SKIP

Simple Key-Management for Internet Protocol – SKIP differs from ISAKMP in the area of negotiation. In SKIP, all of the security parameters are identified within each SKIP secured packet in the form of a SKIP header. The cryptographic algorithms defining the VPN services in a SKIP VPN are predefined, instead of negotiated dynamically as in ISAKMP.

Smart Card

A special type of credit-card like authentication device (assigned to an individual user) that offers a greater degree of private network access security.

Split Tunneling

Split tunneling allows the remote client to simultaneously maintain both a VPN (secure) connection and a clear connection. This function is active by default, however, disabling Split Tunneling turns it off allowing only secure VPN traffic from the remote client's computer. Control of Split Tunneling is normally set when the Dyna-Policy configuration download to the remote client's computer occurs.

SSL

Secure Sockets Layer is a protocol that provides authentication for servers and browsers as well as secure communications between a web server and browser. Used by the VPNmanager Console to communication with the security gateways and the Directory Server.

Syslog

Syslog enables each security gateway in the VPN to provide logging data to a specified destination for historical purposes.

T

Triple DES

A cryptographic algorithm based on DES that encrypts a block of data three times with different keys.

U

User Groups

User Groups are logical groups in which individual VPN user members reside. User Groups have a single-level hierarchy. Users can belong to more than one User Group.

V

VPN

Virtual Private Network. A VPN allows the sending of sensitive, secured data through an unsecure network like the Internet by using dynamically created connections between member of the VPN.

Index