

IBM SecureWay[®] Firewall for Windows NT[®]



User's Guide

Version 4 Release 1

IBM SecureWay[®] Firewall for Windows NT[®]



User's Guide

Version 4 Release 1

Note

Before using this information and the product it supports, be sure to read the general information under "Appendix. Notices" on page 197.

Fourth Edition (September 1999)

This edition applies to Version 4 Release 1 of the IBM SecureWay Firewall for Windows NT (product number 5697-F59). This edition replaces GC31-8658-02.

Portions Copyright © 1995, 1996 by NEC Corporation. All rights reserved.

Contains security software from RSA Data Security, Inc. Copyright © 1990, 1995 RSA Data Security, Inc. All rights reserved.

This application is powered by JScape Components © 1995, 1996, 1997 JScape Corporation, Scottsdale, Arizona. All rights reserved.

© **Copyright International Business Machines Corporation 1994, 1999. All rights reserved.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About This Book	vii	Other Features	24
Prerequisite Knowledge	vii	Common Fields	25
Enhancements	viii	Entering Text in Input Fields	26
Secure Mail Proxy.	viii		
Socks Protocol Version 5	viii		
HTTP Proxy.	viii		
Remote Access Service.	ix		
Terminology	ix		
How to Contact IBM for Service	ix		
Chapter 1. Introducing the IBM Firewall . . .	1	Chapter 5. Getting Started on the IBM Firewall	27
Firewall Concepts	1	Basic Configuration Steps	27
IBM Firewall Tools	2	Designating Your Network Interface	29
Expert Filters	2	Using the Configuration Client to Define a Security Policy	29
Proxy Servers	3	Network Objects	31
Socks Server	4	Using the Configuration Client to Define Network Objects	32
Authentication	5	Network Object Groups	33
Domain Name Service	6		
Secure Mail Proxy	6		
Network Address Translation	6		
Virtual Private Network	7		
Chapter 2. Planning	9	Chapter 6. Handling Domain Name Service 35	
Planning Checklist	9	Configuring DNS Using the Configuration Client	36
Network Configuration Planning Worksheet	10	Configuring the Secure Name Server.	37
		Configuring the Secure Clients.	40
		Publishing Services to the Public	41
		Installing Microsoft's DNS Server	42
		Sample Configurations	42
		Example 1: DNS Server in a DMZ on the Nonsecure Interface	42
		Example 2: DNS in a DMZ on a Dedicated Interface	44
		Example 3: Using the Firewall as the Secure Nameserver.	45
Chapter 3. Setting Up the Configuration Server and the Configuration Client . . .	13	Chapter 7. Secure Mail Proxy	49
Setting Up the Configuration Server	13	How the Secure Mail Proxy Works	49
Sample Logging Output for the Remote Configuration Server	14	Error Handling	50
		Fan-Out Limit	51
		Overflow Server	51
 		Configuring the Secure Mail Proxy Using the Configuration Client	54
Chapter 4. Using the Configuration Client 15		List or Add a Mail Domain Entry	54
Setting Up the Configuration Client (GUI)	15	Change a Mail Domain Entry	55
Starting the Configuration Client	16	Delete a Mail Domain Entry	55
Remote Configuration.	16	Excluded Mail Domains	55
Enabling Remote Configuration through the Configuration Client	16	Domain Name Hiding	56
How to Log On to the Firewall Using the Configuration Client	17	Proxy Characteristics	56
The Setup Wizard	18	Configuring the Overflow Server	59
The Navigation Tree	20	Configuring the Secure Servers.	60
General Features on the Main Panel	21		
The Alerts Display	22		
The Log Viewer	23		

Configuring the Public Domain	60	Socks Logging	95
Using an SMTP Server Instead of the Secure Mail Proxy	60	Client Considerations for Using the Socks Server	95
Disabling the Secure Mail Proxy	61	Tuning the Socks Server	95
Configuring an SMTP Server	61	Socks-Server Chaining	96
Chapter 8. Controlling Traffic through the Firewall	63	Socks Traffic Monitor	97
Versatile Filters Configuration	63	Chapter 12. Configuring Proxy Servers	99
Using the Configuration Client to Build Connections	63	Introducing HTTP Proxy	99
Building Connections Using Predefined Services	65	Migration	99
Ordering Connections.	67	Installation	99
Connection Control	67	Methods Supported	99
Logging	68	Features	100
Determining the Rule States.	69	Browser Configuration	101
Chapter 9. Examples of Services	71	Configuring Particular Features of HTTP Proxy	102
Planning Considerations	71	Proxy Settings	102
Example of Telnet Proxy	72	Proxy Performance	103
Example of Filtered Telnet	73	Logging	105
Example of Proxy HTTP	73	Timeouts.	105
Example of Socks	74	SNMP	106
Example of Virtual Private Networks Using Static Filter Rules	75	User Authentication	107
Chapter 10. Customizing Traffic Control.	77	SSL Tunneling	107
Using the Configuration Client to Create Rule Templates	77	Log Maintenance	107
Change IP Rule Configuration Entry	82	URL Blocking	110
Delete Rule Configuration Entry	82	FTP	110
Predefined Services	83	Transparent FTP	111
Defining Services	86	Telnet	111
Using the Configuration Client to Create Services	86	Transparent Telnet	112
Chapter 11. Configuring the Socks Server	89	Overriding Timeout Values in FTP and Telnet Proxies	113
Authentication	90	Chapter 13. Administering Users at the Firewall	115
Three Authentication Profiles	90	Adding a User to the IBM Firewall	115
Authentication Methods Supported	91	Types of Users	115
Socks and Filters	92	Types of Databases	116
Protocols Supported by Socks Protocol Version 5 Server.	92	Using the Configuration Client to Add a User	116
Configuring the Socks Server Using the Configuration Client	93	Changing a User's Access	126
Add a New Socks Rule	93	Deleting a User from the IBM Firewall.	126
Modify a Socks Rule	95	Administrator Authority Level by Function	126
Delete a Socks Rule	95	Chapter 14. Creating a Virtual Private Network	127
Activate Connection Rules	95	Manual Tunnels	127
		Tunnel Type	129
		IP Tunnel Configuration and Activation	129
		Tunnel Configuration with Endpoints in the Same Subnet	130

Example of Virtual Private Networks Using Static Filter Rules	130	Configuring Remote Access System	174
Configuring Tunnels Using the Configuration Client	131	Configuring the Firewall	174
Add a Tunnel	131	Using Remote Access Service with the Firewall	174
Modify a Tunnel	134		
Delete a Tunnel	134	Chapter 17. Monitoring the Firewall Logging.	177
Export Tunnel Definition Files	134	Threshold Definitions	177
Import Tunnel Definition Files	135	Alert Messages.	177
Tunnel Activation Status	136	Configuring Log Monitor Using the Configuration Client	178
Setting Up Static Filter Rules for a VPN	136	Add Log Monitor	178
Setting Up Tunnels Using Dynamic Filters	138	Change a Threshold Definition	179
Firewall Interoperability	139	Delete a Threshold Definition	179
How to Use the IBM Firewall with the AIX Operating System	139	Page Notification Support.	179
How to Use the conv_export_file Utility	139	What Carriers and Modems are Supported	180
Creating a Tunnel to a Product Other Than the IBM Firewall or the AIX Operating System.	139	Default Configuration Files Supplied with the Firewall	181
Unique SPI Values	143	Configuring Pager Notification Support	181
		Command Customization	181
Chapter 15. Network Address Translation	145	Carrier Administration	183
The IBM Firewall NAT Implementation	145	Modem Administration	184
Static Mapping.	146	Pager Notification Logging	186
Dynamic Mapping	146	Testing Pager Setup	186
More about Packet Changes Made by IBM Firewall NAT	148	Execute Commands	186
IBM Firewall NAT and Routing	150		
IBM Firewall NAT's Position in the Packet-Processing Sequence	151	Chapter 18. Managing Log and Archive Files	189
IBM Firewall NAT and Filters.	152	Log File Creation Using the Configuration Client	189
IBM Firewall NAT and IPsec Tunnels	153	Add Log Facilities	190
IBM Firewall NAT Log Messages	156	Change Log Facilities	191
Configuring Network Address Translation Using the Configuration Client	156	Delete Log Facilities	191
Add NAT Entry	157	Archiving Logs Using the Configuration Client	192
Many-To-One Registered IP Address	158	Plug-in DLL	192
Translate Secured IP Address	158	Log Management Outputs	193
Exclude Secured Network Address	159	Report Utilities.	193
Map Secured Network Address	160	Running Report Utilities Using the Configuration Client	194
Change NAT Entry	160		
Delete NAT Entry.	161	Appendix. Notices	197
NAT Activation	161	Trademarks	198
NAT Configuration Examples.	162		
Case 1: Many-To-One Configuration	163	Bibliography	199
Case 2: Map Configuration.	165	Information in IBM Publications	199
NAT Configuration Statements	169	Firewall Topics.	199
		Internet and World Wide Web Topics	199
Chapter 16. Remote Access Service.	173	General Security Topics	199
Installing Remote Access Service.	173		

Information in Industry Publications 200
Index 201

**Readers' Comments — We'd Like to Hear
from You 203**

About This Book

This book describes how to configure and administer the IBM® SecureWay® Firewall on a Microsoft Windows NT® system so that you can prevent unwanted or unauthorized communication into or out of your secure network.

This book is intended for network or system security administrators who install, administer, and use the IBM Firewall. Although we describe how to access the firewall using client programs, this is not a user's guide for client programs. To use client programs such as telnet or FTP, see the user's guide for your TCP/IP client programs.

Please refer to the IBM SecureWay Firewall Setup and Installation Instructions, available on the product CDROM, for information on how to set up and install this product.

The first time you start the configuration client, the Setup Wizard appears. You can use it to perform certain fundamental firewall tasks. Online help information is available from the configuration client to guide you through the configuration client interface.

Prerequisite Knowledge

It is important that you have a sound knowledge of TCP/IP addressing, masks, and network administration before you install and configure the IBM Firewall. Because you will set up and configure a firewall that controls the access in and out of your network, you must first understand how the network operates. Especially, you need to understand the basics of IP addresses, fully qualified names, and subnet masks.

A recommended book on TCP/IP that covers netstat, arp, ifconfig, ping, nslookup, DNS, sendmail, routing, and much more is *TCP/IP Network Administration*. See "Bibliography" on page 199 for more details.

A recommended book for those performing administration is the *Windows NT Server 4.0 Administrator's Bible*. See "Bibliography" on page 199 for more details.

Enhancements

The IBM SecureWay Firewall V4 for Windows NT offers several extensions:

- Secure mail proxy enhancements
- Socks Protocol Version 5 enhancements
- An enhanced HTTP Proxy using IBM Web Traffic Express technology
- Remote Access Service

Secure Mail Proxy

The IBM Firewall Secure Mail Proxy has been enhanced to include the following new functions:

- Anti-SPAM algorithms including message blocking from known SPAMers (an exclusion list), verification checks on the validity and replyability of messages (known ways of blocking undesirable messages), configurable limits on the number of recipients per mail messages, configurable limits on the maximum size of a message
- Anti-spoofing support including integration with strong authentication mechanisms
- SNMP trap support and support for the MADMAN MIB
- Message tracking including the ability to seamlessly track messages between the firewall and Domino

Socks Protocol Version 5

Socks Protocol Version 5 has been upgraded to include username-password authentication (UNPW), challenge/response authentication (CRAM), and authentication plug-ins.

Logging has been enhanced to give the user more control in classifying log messages and in specifying logging levels.

HTTP Proxy

The IBM SecureWay Firewall provides a full-featured HTTP proxy implementation based upon the IBM Web Traffic Express (WTE) product. The HTTP proxy efficiently handles browser requests through the IBM Firewall, eliminating the need for a socks server for Web browsing. Users can access useful information on the Internet, without compromising the security of their internal networks and without altering their client environment to implement HTTP proxy.

Remote Access Service

Windows NT Remote Access Service (RAS) provides network connections over dial-up, ISDN, or X.25 media using Point-to-Point Protocol (PPP). NDISWAN is a networking driver which is provided as part of RAS and converts the underlying PPP data to resemble Ethernet LAN data.

Terminology

You can access the IBM Dictionary of Computing at:
<http://www.networking.ibm.com/nsg/nsgmain.htm>.

How to Contact IBM for Service

The IBM Support Center provides you with telephone assistance in problem diagnosis and resolution. You can call the IBM Support Center at any time; you will receive a return call within eight business hours (Monday–Friday, 8:00 a.m.–5:00 p.m., local customer time). The number to call is 1-800-237-5511. Or you can access the following web site:
<http://www.software.ibm.com/security/firewall/support/> and specify the country where service is required.

Outside the United States or Puerto Rico, contact your local IBM representative or your authorized IBM supplier.

Chapter 1. Introducing the IBM Firewall

The IBM SecureWay Firewall for Windows NT[®] is a network security program running on the Windows NT operating system.

In essence, a firewall is a blockade between one or more secure, internal private networks and other (nonsecure) networks or the Internet. The purpose of a firewall is to prevent unwanted or unauthorized communication into or out of the secure network. The firewall has three jobs:

- Enforce your Internet security policies
- Let users in your own network use authorized resources from the outside network without compromising your network's data and other resources
- Keep unauthorized users outside of your network

Firewall Concepts

The any-to-any connectivity of the Internet can introduce many security risks. You need to protect your own private data and also protect access to the machines inside your private network to prevent abusive external use. The first step to achieving this protection is to limit the number of points at which the private network is connected to the Internet. A configuration where the private network is connected to the Internet by just one gateway gives you control over which traffic to allow into and out of the Internet. We call such a gateway a firewall.

To understand how a firewall works, consider this example. Imagine a building where you want to restrict access and to control people who enter in. The building's single lobby is the only entrance point. In this lobby, you have some receptionists to welcome people who enter the building, some security guards to watch over them, some video cameras to record their actions, and some badge readers to authenticate their identity.

This works very well to control entry to a private building. But if a non-authorized person succeeds in getting past the lobby, there is no way to protect the building against any actions from this person. However, if you supervise the movement of this person, you might be able to detect any suspicious behavior.

When you are defining your firewall strategy, you may think it is sufficient to prohibit everything that presents a risk for the organization and allow the rest. However, because of new attack methods, you need to anticipate how to prevent these attacks and, as in the case of the building, you need to monitor

for signs that somehow your defenses have been breached. Generally, it is much more damaging and costly to recover from a break-in than to prevent it in the first place.

IBM Firewall Tools

The IBM Firewall is like a tool box you use to implement different firewall architectures. Once you choose your architecture and your security strategy, you select the necessary IBM Firewall tools. The IBM Firewall configuration client provides a user-friendly graphical user interface for administration. The IBM Firewall provides comprehensive logging of all significant events, such as administration changes and attempts to breach security.

Because the IBM Firewall is, at heart, an IP gateway, it divides the world into two or more networks: one or more nonsecure networks and one or more secure networks. The nonsecure network is, for instance, the Internet. The secure networks are usually your corporate IP networks. Some of the tools that the IBM Firewall offers are:

- Expert filters
- Proxy servers
- Socks servers
- Authentication
- Specific services such as domain name service (DNS) and Secure Mail Proxy
- Network Address Translation
- Virtual Private Networks

Expert Filters

Expert filters are tools that inspect packets at the session level based upon multiple criteria such as time of day, IP address, and subnet. The filter rules work with the IP gateway function so the machine is required to have two or more network interfaces, each in a separate IP network or subnet. One set of interfaces is declared nonsecure and the other set is declared secure. The filter acts between these two sets of interfaces, as illustrated in Figure 1 on page 3.




Figure 1. Firewall with Expert Filtering

Objectives of Expert Filters

Expert filtering provides the basic protection mechanism for the firewall. Filters allow you to determine what traffic passes across the firewall based on IP session details, thereby protecting the secure network from external threats such as scanning for secure servers or IP address spoofing. Think of the filtering facility as the base on which the other tools are constructed.

Proxy Servers

Unlike filtering, which inspects packets passing through, proxy servers are applications that are part of the firewall and perform specific TCP/IP functions on behalf of a network user. The user contacts the proxy server using one of the TCP/IP applications, such as Telnet. The proxy server makes contact with the remote host on behalf of the user, thus controlling access while hiding your network structure from external users. Figure 2 illustrates a proxy Telnet server intercepting a request from an external user.




Figure 2. Firewall with a Proxy Server

The proxy services available are telnet, FTP, HTTP, WAIS, GOPHER, and HTTPS, and Secure Mail Proxy.

The IBM Firewall proxy servers can authenticate users with a variety of authentication methods. Users can access useful information on the Internet, without compromising the security of their internal networks.

Objectives of Proxy Servers

When you connect through a proxy server, the TCP/IP connections are broken at the firewall, so the potential for compromising the secure network is reduced. Users may be required to authenticate themselves, using one of a number of authentication methods.

One major advantage inherent in proxy servers is internal address hiding. All outbound proxy connections use the firewall address. Another major advantage of the proxy server is security. IBM experts have developed these proxy servers to guard against security weaknesses, which might be on the client machine.

Another advantage of the proxy server is that you do not need a special version of the client program on the client machine. Therefore, once you have installed your firewall, every user recorded in the Firewall can have access to the nonsecure network without any additional software installation.

Socks Server

Socks is a standard for circuit-level gateways that provides address hiding but does not require the overhead of a more conventional proxy server. The Socks server is similar to a proxy server in that the session is broken at the firewall. The difference is that socks can support all applications instead of requiring a unique proxy for each application.

The IBM Firewall provides the Socks Protocol Version 5, which enables clients inside the secure network to pass an authentication stage before accessing applications in the nonsecure network. It also provides for authenticated generic proxy and the proxy of some streaming audio and video protocols.

Transparently, the socks client starts a session with the Windows NT socks service on the IBM Firewall host then validates that the source address and user ID are permitted to establish onward connection into the nonsecure network and then creates the second session.

Figure 3 on page 5 illustrates a firewall with a socks server.




Figure 3. Firewall with a Socks Server

Socksified clients (clients, that are Socks-aware) are available with many applications like Netscape Navigator or Microsoft® Internet Explorer, or through TCP/IP software such as Aventail AutoSocks.

Objectives of the Socks Server

For outbound sessions (from a secure client to a nonsecure server) the socks server has the same objectives as a proxy server, that is to break the session at the firewall and provide a secure door where users must prove their identity in order to pass. It has the advantage of simplicity for the user, with little extra administrative work.

Authentication

Authentication means you can use a password or a stronger method to access your network. This is especially useful when you want to log in remotely, such as when you are traveling or working at home. The IBM Firewall lets you choose which method you need for authenticating clients. You can use just a password for access, or you can use more sophisticated methods, such as the Security Dynamics SecurID token.

The authentication method from Security Dynamics includes a user ID and a SecurID token. When you log in remotely, you get your password from the SecurID token. The password changes every 60 seconds and is good for one-time use only. So, even if someone does intercept your password over the open network, the password is not valid for additional use.

The IBM Firewall comes with the Security Dynamics ACE/Server with a two-user license.

You can also customize a user exit to support any other authentication mechanism. The IBM Firewall includes an application programming interface (API) to help you define your own authentication technique.

If you choose to authenticate users with passwords, the rules are robust. The IBM Firewall applies extensive password rules to ensure nontrivial passwords are used.

Domain Name Service

Access to the domain name records for the secure network is of great assistance to intruders, because it gives them a list of hosts to attack. A subverted domain name service server can also provide an access route for an intruder. From the external network, the name server on the firewall only knows itself and never gives out information on the internal IP network. From the internal network, this name server knows the Internet network and is very useful for accessing any machine on the Internet by its name.

Objectives of the DNS Server

Running the DNS server on the firewall has the dual advantage of preventing name resolution requests flowing across the firewall and hiding secure network hosts from the nonsecure world.

Secure Mail Proxy

Mail is one of the primary reasons why an organization would want to access the Internet. Secure Mail Proxy is an efficient IBM mail gateway designed to control mail traversal through your network. It allows mail to flow securely inside and outside of your network.

The Secure Mail Proxy function does not store mail on the gateway. The firewall gateway public domain name is substituted in place of the private domain names on outgoing mail so that mail appears to be coming from the firewall's address instead of the user's address. Secure Mail Proxy supports Simple Mail Transfer Protocol (SMTP) and Multipurpose Internet Mail Extensions (MIME).

Network Address Translation

Addresses within a private network can be assigned out of a very large address space (typically a 10.0.0.0 class A address space). These addresses are private and are not exposed to the Internet. Therefore these addresses can be reused by another IP network. With NAT, a single registered IP address is used to hide many private network addresses. NAT converts the unregistered addresses into valid registered Internet addresses. In the inbound direction, NAT converts the registered Internet address back to the unregistered addresses.

The advantage of NAT is that it transparently allows a network that uses private or illegal addresses to communicate with hosts on the Internet,

effectively allowing the private network to have a large address space. Furthermore, by using NAT, addresses in the private network are hidden from the external world providing an additional level of security.

Figure 4 illustrates basic NAT operation in an IBM Firewall environment.

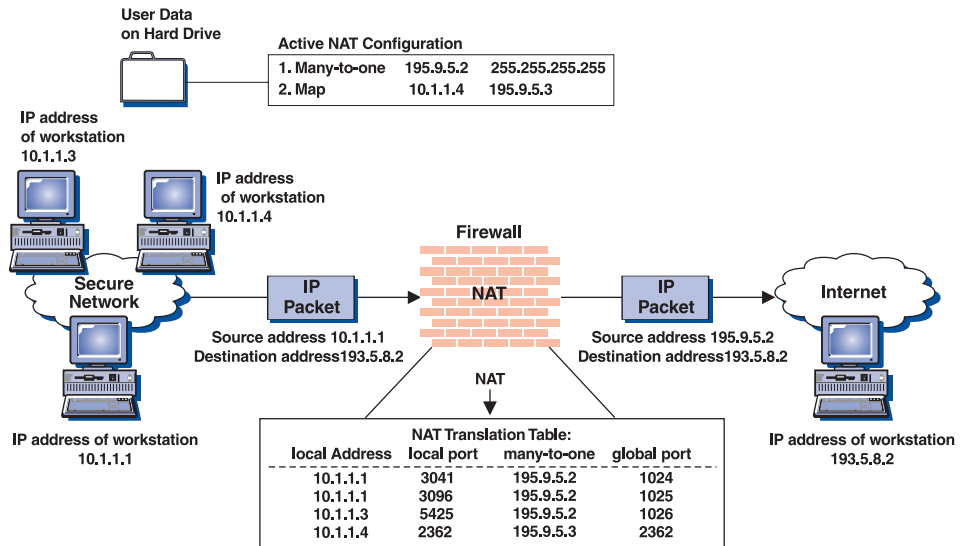


Figure 4. Network Address Translation

Virtual Private Network

The IBM SecureWay Firewall V4R1 provides support for manually configured VPN tunnels. A Virtual Private Network (VPN) is an extension of an enterprise's private intranet across a backbone network, which typically will be a public backbone such as the Internet. A VPN allows you to create a secure connection to protect your data while it is in transit over the backbone. The VPN tunnel uses the open IPSec security standards to protect your data from modification or disclosure while it is travelling between firewalls. Your data will flow within a VPN tunnel, which can provide data origin authentication, confidentiality, and integrity checking on every packet. IPSec protocols can keep your data private, hiding it from any eavesdroppers on the public network. Packet filtering in the firewall can be used in conjunction with IPSec technologies to further protect your intranets from unwanted intrusions.

VPNs can be created by manually configuring VPN tunnels between pairs of IBM Firewalls or between an IBM Firewall and any other device (client, router, server, or firewall) that supports the latest open IPSec standards.

Encryption support can include 3DES, DES, and CDMF. Authentication support includes HMAC-MD5 and HMAC-SHA. Filters for your VPN tunnels can be customized, or you can choose a default set of filter rules.

Figure 5 illustrates a secure IP tunnel and a VPN.




Figure 5. Tunnel, All IP Traffic between Two Secure Networks. FW₁ and FW₂ represent nonsecure interface IP address and mask. SN₁ and SN₂ represent any host in the secure network. The shaded area of the picture represents a VPN.

Objectives of the Virtual Private Network

The virtual private network allows you to obscure the real data being sent between two private networks and also allows you to be assured of the identity of the session partners and the authenticity of the messages.

Chapter 2. Planning

Before you configure the IBM Firewall, use the checklist and the planning worksheets to help you understand your network configuration.

Planning Checklist

1. Define your objective. Do you want to:
 - Access the Internet (telnet, anonymous FTP, etc.)?
 - Partition parts of your internal network?
 - Provide *external* access to your network?
2. For security reasons, the Firewall should not be a Primary (PDC) or Backup Domain Controller (BDC). The Firewall should be a standalone machine or reside in an NT domain as a standalone server.
3. Evaluate the topology of your network at the IP subnetwork level.
 - Is one secure and one nonsecure interface a correct configuration?
 - Are your addresses able to support subnet masks in rules?
4. Decide how you will use DNS. Refer to "Chapter 6. Handling Domain Name Service" on page 35.
5. Decide how you will use Secure Mail Proxy. Refer to "Chapter 7. Secure Mail Proxy" on page 49.
6. If you want to use Socks, ensure socksified clients, such as the Netscape Navigator or the Microsoft Internet Explorer are installed on the client host. For information on using Socks, see "Chapter 11. Configuring the Socks Server" on page 89.
7. What type of authentication is required?
 - If you are going to use the Security Dynamics ACE/Server to authenticate users, we suggest that you install the ACE/Server server code on a non-IBM Firewall host inside the secure network.
For information about installing and using a Security Dynamics ACE/Server and the SecurID card, see the information that is provided by Security Dynamics Technologies Inc.
 - If you use your own authentication method, see the chapter on "Providing Your Own Authentication Methods" in the *IBM SecureWay Firewall Reference*.
 - Because NETBIOS will be disabled by the hardening process, you must configure the Windows client code that implements the ability to search trusted Windows NT domains for authentication purposes, to use TCP instead of NETBIOS. The trusted Windows NT servers must have TCP/IP host names and addresses and have TCP/IP connectivity between them and the firewall. The firewall administrator needs to

create connections between the firewall and the trusted NT servers in order to permit traffic to flow between the two.

Set up this connection using the following predefined services:

- a. Domain Controller Authentication - which allows the use of the Domain Controller for user authentication
- b. NetBT Name Services broadcasts - which allows NetBIOS over TCP/IP Name Services broadcasts

And use the NT configuration utilities to define the trust relationships.

8. If you use filtering, start with simple filter rules and make them highly restrictive. Become familiar with ports and protocols used by services you need.
9. Decide on a method for archiving log files. Archiving is an ideal candidate for a scheduled job in the Windows NT Scheduler service. See "Chapter 18. Managing Log and Archive Files" on page 189.
10. If you are going to set up a virtual private network, decide on your tunnel end points and methods of encryption and authentication. See "Manual Tunnels" on page 127
11. If you are using a version of the Firewall that requires key recovery to be installed for IPSec processing for export/import encryption regulations, then you need to obtain policy management configuration files. For more information, see your IBM Marketing representative.

Network Configuration Planning Worksheet

Fill in the following information as part of the planning for your IBM Firewall configuration.

Host name of firewall _____

Secure network interface(s) (connected to internal secure network)

IP address _____ Subnet Mask _____

IP address _____ Subnet Mask _____

IP address _____ Subnet Mask _____

IP address _____ Subnet Mask _____

Nonsecure network interface(s) (connected to untrusted nonsecure network)

IP address _____ Subnet Mask _____

IP address _____ Subnet Mask _____

IP address _____ Subnet Mask _____

IP address _____ Subnet Mask _____

Name of router _____

Address of router _____

Secure domain name _____

IP address of secure domain name server (DNS) _____

IP address of nonsecure domain name server(s) (DNS) _____

Secure Mail Server _____

Public Domain Name _____

IP address of the configuration client _____

IP address of the remote client(s) _____

Root directory of your Windows NT Firewall _____

(We refer to this as ROOTDIR throughout the documentation)

(We assume that Windows NT is installed in c:\winnt)

Chapter 3. Setting Up the Configuration Server and the Configuration Client

This chapter tells you how to set up the configuration server and the configuration client, which is the graphical user interface (GUI) for the IBM Firewall.

Setting Up the Configuration Server

The configuration server is the configuration client's interface to the Firewall. The configuration server processes requests from the configuration client. It runs on the Firewall machine and can handle requests from configuration clients that are on either local or remote machines. Once you have set it up, consider it part of the Firewall machine.

The configuration server's port number is specified in the Windows NT services file located in the directory where you installed the Windows operating system: `c:\winnt\system32\drivers\etc\services`. The port number defaults to 1014, but you can change this, for added security, by stopping the configuration server service, modifying the services file, and restarting the configuration server service.

The configuration server is initially set up to only accept requests from configuration clients on the local machine. Initial requests are not encrypted. To change these options, use the following command from the command line:

```
fwcfgsrv cmd=change
```

localonly=

Indicates if the Firewall can only be administered from a local machine.

localonly=yes

The configuration can occur only on the local machine; this is the default.

localonly=no

The configuration can occur from any machine.

encryption

Indicates if the configuration server expects incoming data to be encrypted through secure sockets layer (SSL) or not.

If you change the encryption option or the `sslfile`, you must stop and restart the configuration server service.

encryption=none

No encryption will occur; this is the default.

encryption=ssl

SSL encryption will occur.

sslfile=

Indicates the name of the SSL keyfile to be used with SSL encryption; the default is `ROOTDIR\config\fwkey.kyr`. *ROOTDIR* is the directory that you have selected during the installation process as the target location for the IBM Firewall. For information on how to create the keyfile, see the *IBM SecureWay Firewall Reference*.

If a configuration client cannot connect to the Firewall machine, and is on a different machine, use `fwcfgsrv cmd=list` to check that `localonly=no` is set. Also, the language used by the client and the server must match. Finally, ensure that the configuration server service is running by bringing up the services panel and checking its status. To do this, go to the control panel, double-click the Services icon to check the status of each service. If it is not running, the service should be restarted.

Sample Logging Output for the Remote Configuration Server

The following is a sample of the logging output for the remote configuration server:

```
Feb 03 13:52:15 1998 mr16n18: ICA9005i: Starting remote configuration server.
```

```
Feb 03 13:52:21 1998 mr16n18: ICA2024i: User administrator successfully authenticated using NT authentication from secure network:127.0.0.
```

```
Feb 03 13:52:21 1998 mr16n18: ICA2169i: User administrator successfully authenticated for Remote Administration Server using NT from secure network:127.0.0.1.
```

Chapter 4. Using the Configuration Client

Use the configuration client, which is a graphical user interface, to configure and administer the IBM Firewall.

You must use the configuration client shipped with this IBM Firewall.

When you first install the IBM Firewall, it is initially set up to only accept requests from the configuration client on the local machine. However, you can install the configuration client on another machine and administer the Firewall remotely. See “Setting Up the Configuration Server” on page 13 for information on how to do this.

Setting Up the Configuration Client (GUI)

When you install the IBM Firewall, the configuration client is automatically installed. The configuration client can also be separately installed on any Windows NT or 9X machine without the Firewall, which enables you to perform remote administration. To start the configuration client, select it from the start program menu. When the configuration client is started, you must first log on to the Firewall using a Windows NT administrator account.

Only Windows NT administrators and firewall administrators that have the appropriate administration authentication can use the configuration client to log on to the Firewall.

After the Firewall is installed, all Windows NT administrators are designated as primary firewall administrators. Use the configuration client to log on to the configuration server using a primary firewall administrator and define the additional firewall administrator usernames, if necessary. See “Chapter 13. Administering Users at the Firewall” on page 115 for information on how to define firewall administrators using the configuration client.

To set the logon timeout value for faster or slower machines, make the following change by clicking the IBM Firewall Configuration Client icon, then click **Properties**. Modify Properties by using the **Shortcut** tab. Change the parameter `timeout` to `20`, where 20 equals the number of seconds to wait for a connection to occur.

To increase the level of debug information in the Java console, run `ibmfw.bat` in `ROOTDIR\cfgcli\gui` instead of using the configuration client icon. Note however, that enabling console logging can degrade performance.

Starting the Configuration Client

To set the configuration client to start in the language for your specific locale, click the IBM Firewall Configuration Client icon, then click **Properties**. Modify Properties by using the **Shortcut** tab. By default, the locale of the host machine is used. The IBM Firewall will support these locales:

- en_US - US English
- ja_JP - Japanese PC
- ko_KR - Korean
- zh_CN - Simplified Chinese EUC
- zh_TW - Traditional Chinese (Big 5)
- fr_FR - French
- it_IT - Italian
- pt_BR - Brazilian Portuguese
- es_ES - Spanish PC
- de_DE - German

A mouse is required to use the configuration client.

A **Help** button is located near the top of the configuration client main panel. Click **Help** for information on any function.

Remote Configuration

The IBM SecureWay Firewall for Windows NT V4R1 configuration client can be used to remotely administer an IBM SecureWay Firewall for AIX V4R1 and the reverse is true. However, the IBM SecureWay Firewall V4R1 configuration clients cannot be used to administer previous versions of the IBM Firewall.

Enabling Remote Configuration through the Configuration Client

To enable remote configuration through the configuration client, make sure the administrator that is going to log on has the following attributes defined on the Firewall machine:

- If the administrator is on the secure side of the network and using a secure interface on the Firewall machine, then he or she must be defined with the appropriate authentication method for secure administration. (It cannot be set to deny all). This applies to logging on to the Firewall locally as well.
- Similarly, if the administrator is on the nonsecure side and using a nonsecure interface on the Firewall machine, then he or she must be defined with the appropriate authentication method for nonsecure administration. (It cannot be set to deny all).

All of the user attributes can be set through the Modify User dialog box in the configuration client or by using the command `fwuser`. All firewall administrators will have all of the above fields set appropriately after installation of the Firewall. Refer to “Chapter 13. Administering Users at the Firewall” on page 115 for more information.

How to Log On to the Firewall Using the Configuration Client

To log on to the IBM Firewall using the configuration client (on the local or remote machine):

- The user must be a firewall administrator
- The firewall administrator must have an authentication scheme defined. See “User Authentication Methods” on page 120.
- The user must have the authority to perform specific configuration functions

1. For Logon Type, select Local if you are on the same machine as the firewall. Local is the default. Select Remote if you want to remotely access another Firewall. Remote requires that you enter a host name.
2. If you selected Remote logon, you need to enter the host name or the IP address of the firewall machine you want to log on to.
3. Select either SSL or none depending upon which encryption is used for the Firewall. For the Client, the default for Local is None and the default for Remote is SSL.

If you wish to use SSL, you must first create a certificate on the Firewall. Refer to the *Using the Make Key File Utility* chapter in the *IBM SecureWay Firewall Reference*.

4. Enter a user name of a firewall administrator or a Windows NT administrator.
5. Enter the port number on which the server is listening. The default is 1014.
6. For Mode, select Host if you want to configure a Windows NT firewall machine that you are logging on to. With host administration, the administrator can locally or remotely update one Firewall at a time.
7. After you log on, you will see authentication messages and you might be prompted to enter a password if that is the authentication method setup for your user name. If you are prompted for a password, enter your password in the User Response field and either press Enter or click Submit. If you enter an incorrect password, you get a message. Click Close and restart the logon process. If you are not prompted for a password, your user authentication method might be permit all. In this case you will immediately get the IBM Firewall configuration client panel.

- After you have successfully been authenticated, the Setup Wizard will automatically appear. See “The Setup Wizard” for more information. With subsequent logons, you will see the main configuration panel.




Figure 6. Configuration Client Logon Panel

The Setup Wizard

The setup wizard aids you with the initial configuration of the Firewall. It is especially helpful if you do not have an extensive knowledge of the firewall configuration because it enables you to have a basic firewall configuration up and running quickly after installation.

The setup wizard appears automatically after you log onto the Firewall for the first time, as shown in Figure 7 on page 19.

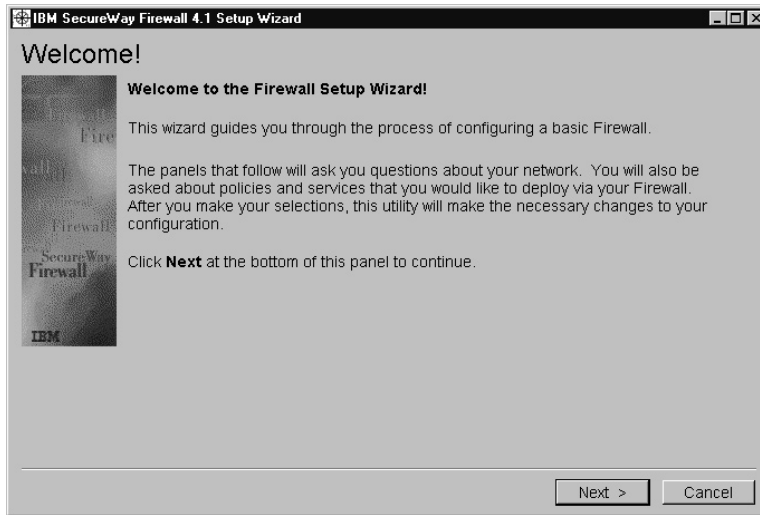


Figure 7. Setup Wizard

Thereafter, the setup wizard is available under the **Help** menu item on the GUI. The setup wizard is optional; you are not required to use it to configure the Firewall.

The setup wizard guides you through the following fundamental tasks:

- Basic security policies
- System administration tasks having to do with interfaces, DNS, mail, and log setup
- Setup to allow secure users to access nonsecure networks through the Web, Telnet, or FTP
- Setting authentication methods for the default user based upon services that they enabled with the wizard
- Creating an alert log
- Setting up some basic log monitor thresholds

The setup wizard can be helpful for getting started on a variety of firewall installations. However, depending upon your circumstances, the wizard may not be recommended. The wizard is not recommended for:

- Migrating a configuration from a previous version of the Firewall
- Setting up a demilitarized zone (DMZ) that involves designating two or more network interfaces as secure
- Setups that require more than one security policy for the secure networks

The Navigation Tree

The configuration client has a collapsible tree-style navigation aid along the left side, as shown in Figure 8 on page 21.

If a node or function has items under it, a file folder icon appears at the left of the node. To see the subfunctions you can expand the view by double-clicking on the icon. Double-clicking on the icon again collapses the view of this node back to the original view.

Any function that you click is considered selected and is highlighted. You can expand and collapse the nodes without any change to the window view on the right. When the expanded tree exceeds the vertical space available, a scroll bar appears at the right of the navigation tree. A horizontal scroll bar appears if any of the function names do not fit into the navigation tree.

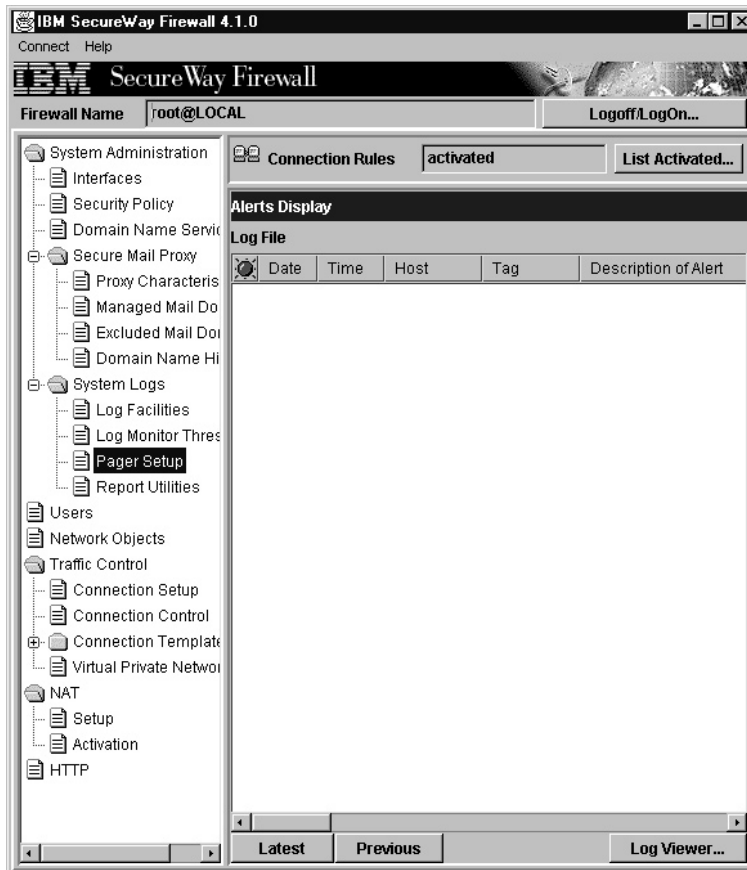


Figure 8. Configuration Client Navigation Tree

General Features on the Main Panel

A **Help** button is located near the top of the configuration client main panel, as shown in Figure 8. Click **Help** to see what to do to get your IBM Firewall up and running. Some of the product publications, the Setup Wizard, and the Readme file are also accessible from this button.

Other buttons that you will encounter on the main panel are:

Logoff/LogOn

A **Logoff/LogOn** button is located in the upper right corner of the configuration client. It is a reconnect button. You can restart the logon sequence to connect to a different Firewall or to log on as a different administrator.

To log off, click Logoff, click Cancel on the logon panel, and the application.

Latest A **Latest** button is located at the bottom of the configuration client main panel. Click **Latest** to see the most recent alerts, if you have defined an alerts log.

Previous

A **Previous** button is located at the bottom of the configuration client main panel. Click **Previous** to see earlier alerts, if you have defined an alerts log.

Log Viewer

A **Log Viewer** button is located in the lower right corner of the configuration client. It allows you to browse firewall logs.

The Alerts Display

You can view alert records generated by the system log monitor in the lower right section of the main configuration client window, as shown in Figure 9 on page 23.

The alert records displayed are obtained from the file identified by the first alert log facility defined in `ROOTDIR\config\syslog.conf`. If no alert log facility is defined, you will see a blank display. See “Add Log Facilities” on page 190 for help in defining an alert log facility.

The panel shows you the name of the alerts file and the line numbers currently displayed from that file. You can click **Latest** to see the most recent alerts. Clicking **Previous** allows you to see earlier alerts.

Each line displayed shows the date and time of the alert, the host name of the firewall on which the alert occurred, the alert message tag, and the text of the alert message. The tag is an indication of the type of the alert.

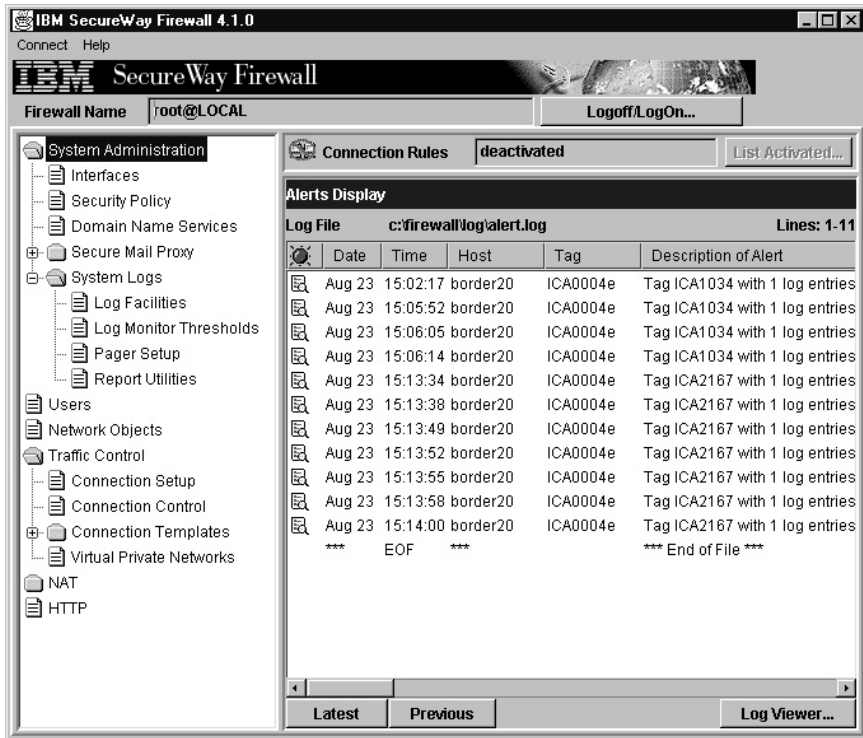


Figure 9. The Alerts Display

The Log Viewer

Clicking **Log Viewer** brings up a log viewer window, as shown in Figure 10 on page 24. The log viewer allows you to view firewall log records. You can specify a log file and a record count (default is 25).

The default log is the file identified by the first firewall log facility defined in `ROOTDIR\config\syslog.conf`. You can select a different target log file from the file name field's pull-down menu or you can type in the name of a file to be viewed.

To request a specific start line, click **Start at Line:**, after typing the line number in the field next to it. To request the last so many lines, click **Bottom**, which takes you to the bottom of the file. **Next** advances you to the next set of lines in the file. **Previous** takes you back to the previous set of lines in the file. **Top** takes you to the top of the file. In the **Expand Firewall Log text** field, you can expand firewall logs to readable text by clicking **Yes**.

See “Log File Creation Using the Configuration Client” on page 189 and “Chapter 17. Monitoring the Firewall Logging” on page 177 for more information about log files, facilities, monitoring and alerts.

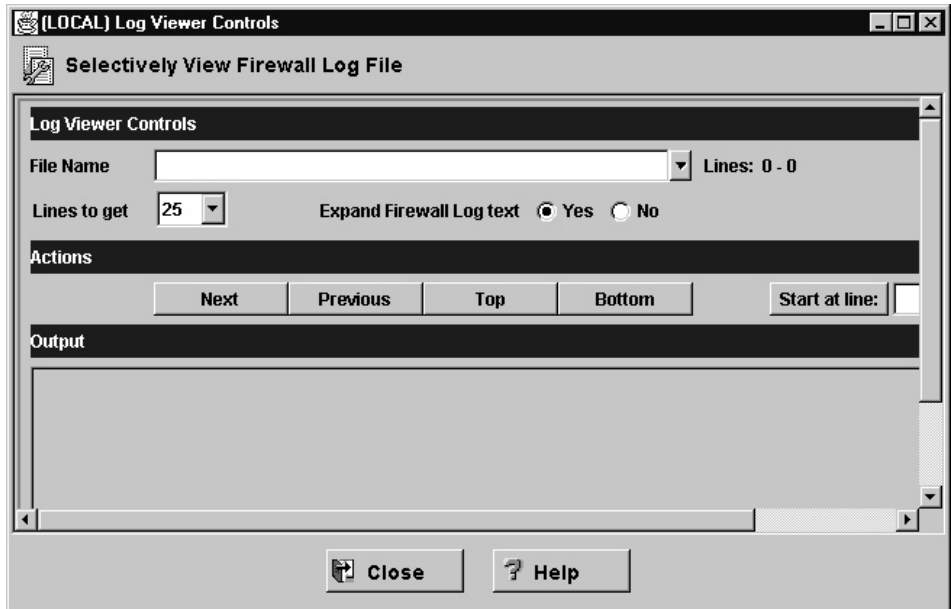


Figure 10. Log Viewer

Other Features

A **Search** field is located near the top left corner of some of the panels. You can enter a search string and click **Find** or press Enter.

Other buttons that you will see on many of the configuration client dialog boxes are:

Apply Click **Apply** to populate the field on the previous panel with your current selection or to save changes you have made on a panel. The **Apply** button will not cause the window to disappear.

Bottom

Click **Bottom** to go to the bottom of a panel.

Cancel

Click **Cancel** to close the window without saving any changes.

Close

Click **Close** to eliminate the window from your display.

Copy

The **Copy** button saves time when adding new items to the list. After

selecting an item on the list, click **Copy** to create an item that is similar to the selected item. When the new item is opened, it will copy field values from the selected item on the list. You will then be able to modify field values as needed for the new item.

Delete Click **Delete** to delete a selected item from the list.

Move Down

Select an item in the list and click **Move Down** to lower the item's relative position in the list. Each click will cause the item to move down one position.

Move Up

Select an item in the list and click **Move Up** to raise the item's relative position in a list. Each click will cause the item to move up one position.

OK Click **OK** to save changes and close the window.

Open After selecting an item on the list, click **Open** to view or modify that item. To add a new item, click **NEW** item on the list and click **Open**.

Refresh

Click **Refresh** to reaccess the data from the firewall and redisplay the data on the panel.

Remove

Click **Remove** to eliminate a selected item from a list. This action will only remove the item from the list. This action will have no effect on other places where the item is defined.

Select Click **Select** to access a list of candidate items that are valid for this function.

Top Click **Top** to go to the top of a panel.

Common Fields

Common fields that you will see on many of the configuration client dialog boxes are:

Output

As the command that you have initiated proceeds, progress information will appear here.

Name Provide a name for this item. This item name must be unique for this particular function in the firewall. The name should NOT contain a pipe symbol(|), a single quote (or apostrophe) character('), or a double quote(") character. Use of these characters can result in unreliable data.

Description

This field is optional and is provided in case you want to provide a comment or additional information about this item.

Entering Text in Input Fields

All user-generated input is restricted to an invariant character set, a character set such as the syntactic character set, whose code point assignments do not change from code page to code page. Basically this means that user input is restricted to ASCII letters and numbers. For example, passwords, user IDs, user-defined object names, descriptions, and so forth are in this restricted set of characters even if messages and so forth are in a language other than English.

The syntactic character set is the set of graphic characters registered in the IBM registry with a GCSGID of 00640:

- All upper and lower case characters of the English alphabet
- 10 digits
- + < = > % & * " ' () , - . / \ : ; ~ ?

Chapter 5. Getting Started on the IBM Firewall

This chapter gives you the basic configuration steps you need to get your IBM Firewall set up. It explains how to define a secure interface, how to determine your security policy, and how to define network objects.

Basic Configuration Steps

For a basic IBM Firewall setup:

1. Plan for your IBM Firewall setup. Decide in advance which functions of the firewall you want to use and how you want to use them. These sections are helpful:
 - “Chapter 1. Introducing the IBM Firewall” on page 1
 - “Chapter 2. Planning” on page 9
 - “Planning Considerations” on page 71
2. Tell the Firewall which of its interfaces are connected to secure networks. You must have a secure interface and a nonsecure interface for your firewall to work properly. From the configuration client navigation tree, open the System Administration folder and click **Interfaces** to see a list of the network interfaces on your firewall. To change the security status of an interface, select an interface and click **Change**. See “Designating Your Network Interface” on page 29 for more information.

If you are going to connect to the Internet, contact your Internet Service Provider (ISP) to obtain a registered IP address for the Firewall nonsecure interface.

3. Set up your general security policy by accessing the **Security Policy** dialog in the System Administration folder. For typical Firewall configurations it is recommended, at a minimum, that you enable the following policies:
 - Permit DNS queries
 - Deny broadcast message to nonsecure interface
 - Deny Socks to nonsecure adapters

See “Using the Configuration Client to Define a Security Policy” on page 29 for more information.

4. Set up your domain name service and mail service. Very little communication will take place efficiently if you do not provide DNS resolution. Access these functions from the System Administration folder on the configuration client navigation tree. First read “Chapter 6. Handling Domain Name Service” on page 35.

5. Define key elements of your network(s) to the firewall using the **Network Objects** function in the configuration client navigation tree. Network Objects are used in setting up connections. Define the following key elements as network objects:
 - Secure Interface of the Firewall
 - Nonsecure Interface of the Firewall
 - Secure Network
 - Each subnet on your secure network
 - A host network object for your Security Dynamics servers and your Windows NT domain servers, if appropriate

See “Network Objects” on page 31 for more information.

6. Enable services on the Firewall. These are the methods by which users in the secure network can access the nonsecure network (such as socks or proxy). Which services get implemented depend on decisions you made at the planning stage. Implementing a service often requires setting up some connection configurations to allow certain types of traffic. For example, if you want to allow your secure users to surf the web on the Internet by using HTTP Proxy, not only do you need to configure the HTTP Proxy daemon on the Firewall, but you also need to set up connections to allow HTTP traffic. See “Chapter 9. Examples of Services” on page 71 for information on how to set up connections that support certain services.
7. Set up firewall users. If you are going to require authentication for functions like outbound Web access or for firewall administrators, you need to define these users to the Firewall. See “Chapter 13. Administering Users at the Firewall” on page 115 for more information.
8. Because the hardening process disables NETBIOS, if you want to use Windows NT domain passwords for authentication, you must configure the Windows client code that implements the ability to search trusted Windows NT domains for authentication purposes, to use TCP instead of NETBIOS. The trusted Windows NT servers must have TCP/IP host names and addresses and have TCP/IP connectivity between them and the Firewall. The firewall administrator needs to create connections between the Firewall and the trusted Windows NT servers in order to permit traffic to flow between the two.
9. If you will be using network address translation, first contact your ISP to obtain a registered Internet address for use with many-to-one address translation. This address is in addition to the address you requested in step 2 on page 27. Then, go to the *Add NAT Configuration* panel to add the registered Internet address in the *Many-to-One IP Address* field. For more information, see “Chapter 15. Network Address Translation” on page 145.

Following these steps should help you to get a basic firewall configuration up and running. The IBM Firewall provides other functions, such as system logs

to help you ensure the security of your network. See “Chapter 18. Managing Log and Archive Files” on page 189 for more information.

If the Firewall shuts down either normally or abnormally, your configuration data will not be affected because it will be saved to the hard drive and will automatically be reactivated upon rebooting. However, certain firewall log messages will occur indicating that some active connections were interrupted; for example, an active FTP session.

Designating Your Network Interface

This book distinguishes between the secure and nonsecure interfaces, networks, and hosts. Secure interfaces connect the IBM Firewall host to the network of hosts in your internal network, the network that you want to protect. **You must have at least one secure interface and one nonsecure interface for your firewall to work.** Nonsecure interfaces connect the IBM Firewall to one or more outside networks or to the Internet.

All networks attached through a secure interface are considered secure networks. To discriminate between the various subnets attached to the secure interface, use the static filter rules to deny or permit access between several subnets on the same interface based on IP address or an address mask.

To designate secure and nonsecure interfaces, use the System Administration folder on the configuration client navigation tree. All known interfaces (adapters) will be shown and identified as secure or nonsecure.

You must provide a name for each interface before you can perform specific interface filtering.

To identify a network interface as either secure or nonsecure:

1. Select an interface and click **Change**.
2. Repeat as necessary.
3. Click **Close**.

To identify the interface as secure or nonsecure and to provide a meaningful name for that interface, click **Open**. This name will be used by filters for specific interface filtering.

Using the Configuration Client to Define a Security Policy

One of the first things to consider when configuring the IBM Firewall is the general security policy for your installation.

The IBM Firewall provides a dialog box to assist you in setting up your security policy, as shown in Figure 11.

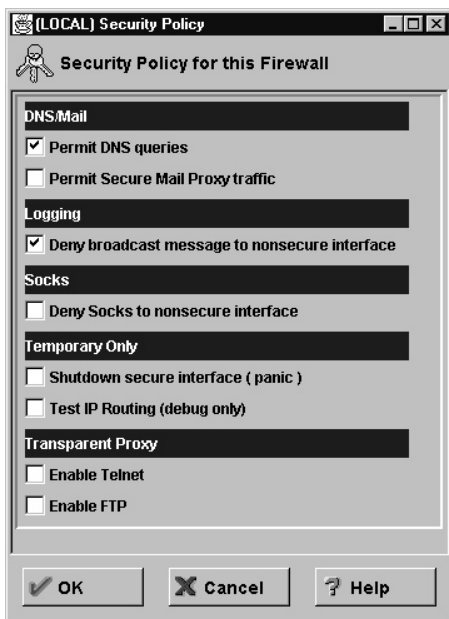


Figure 11. Security Policy

The Security Policy provides a quick and easy way for administrators to set blanket policies for the firewall. Most of the check boxes displayed in the security policy window provide a fast path to selecting certain Predefined Services that will apply to all network traffic received by the Firewall. The exceptions are the Transparent Proxy choices which simply act to enable or disable Transparent Telnet and Transparent FTP.

When you select a security policy, the Firewall builds the filter rules, which you then need to activate. The Firewall enables the services selected and makes them globally available.

Note that any time you select a check box that pertains to a Predefined Service and you click **OK**, you must activate these changes through the Connection Activation window. You do not need to activate the Transparent Proxy selections because these do not pertain to Predefined Services. See "Predefined Services" on page 83 for a list of predefined services.

You are presented with the following list of check boxes from which you can select attributes that reflect the security policy for your site. The attributes selected apply to all addresses on both sides of the IBM Firewall.

- Select **Permit DNS Queries** to allow Domain Name Service resolution requests and replies. Very little communication will take place efficiently if you do not provide DNS resolution. See “Chapter 6. Handling Domain Name Service” on page 35 for details on configuring DNS.
- Select **Secure SMTP Proxy** to allow mail traffic to flow through the Firewall.
- Select **Deny broadcast message to nonsecure interfaces** to prevent broadcast messages from being received at the nonsecure port. If your firewall’s nonsecure interface is connected to the Internet, this service can help reduce the amount of logging on the Firewall.
- Select **Deny Socks to nonsecure interface** to disallow socks traffic to enter the Firewall from the nonsecure network. If you want to allow clients to enter your network from the nonsecure network, you must not turn on this checkbox.
- Select **Shutdown secure interface (panic)** to disallow all traffic to and from the Firewall over the secure interfaces. This is used for emergency purposes only.
- Select **Test IP Routing (debug only)** to allow all traffic to and from Firewall over any interface. Note that if you change the value of this check box, you must save it by clicking **OK** and activate it through the Connection Activation window. **Use of this Service can cause security exposures for your Firewall. Use it with extreme caution.**
- Select **Enable Telnet** to allow Transparent Proxy Telnets.
- Select **Enable FTP** to allow Transparent Proxy FTPs.

Network Objects

Network objects are representations of components that exist in your network such as hosts, networks, routers, virtual private networks, or users. Network objects designate source and destination addresses for services when you create your connections.

Objects can be identified by name, icon representation, type, and description. There are several types of network objects but Host and Firewall are the most common. The default network object shipped with the IBM Firewall is “The World”. This is a global object that encompasses all possible IP addresses. After you have filled in the network configuration worksheets (see “Network Configuration Planning Worksheet” on page 10), you are ready to build objects.

You can create single or group objects. All network objects are defined by an IP address and an address mask (subnet mask) so that it is possible for one object to represent a range of network addresses.

Using the Configuration Client to Define Network Objects

To define a single network object, select **Network Objects** from the configuration client navigation tree. The Network Objects dialog box appears. Double-click **NEW**. The **Add a Network Object** dialog box appears, as shown in Figure 12.

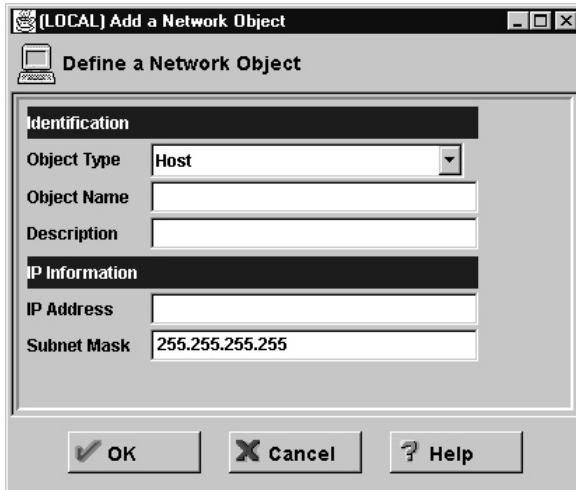


Figure 12. Add a Network Object

1. Enter the object type. Click the **Object Type** arrow to see the object types you can create. For performance reasons, it is better to create network type objects instead of host type objects. The object types you can create are:
 - Host - a particular node on your network with a mask of 255.255.255.255.
 - Network - a collective range of network addresses that is characterized by an address range and a specific subnet mask.
 - Firewall - a single machine with a firewall installed on it with a mask of 255.255.255.255.
 - Router - a host that routes traffic between two or more networks with a mask of 255.255.255.255.
 - Interface - a network adapter on a machine with a mask of 255.255.255.255. It does not have to be an adapter on the Firewall.
2. Fill in the object name. Use a single-byte character set to do this. Double-byte character sets are not supported.
3. Fill in the description. This field is optional but if you do fill it in use a single-byte character set to do this. Double-byte character sets are not supported.

4. Enter a dotted-decimal IP address for this object.
5. Enter a subnet mask that specifies the bits in the address to compare to the address of the IP packet.

Network Object Groups

A group represents a collection of network objects. Groups are used as a convenience in setting up connections and can eliminate repetitive work. An example would be to group some addresses, individually represented by network objects, into a network object group to represent a department. This department can be used as either the source or destination address for a connection.

To define a group of network objects, select Network Objects from the configuration client navigation tree. The **Network Objects** dialog box appears. Double-click **NEW GROUP**. The **Add a Network Object** dialog box appears.

1. Fill in the group name.
2. Fill in a description. This field is optional.
3. Click **Select** to select objects for the group.
4. Click **OK**.

Tip: It is a good idea to encompass contiguous address ranges into a single network object whenever possible. This will improve the performance of the connection rule processing. The following example illustrates this.

```
ACCOUNTING DEPARTMENT
Kevin's machine 191.1.10.1
Susan's machine 191.1.10.3
Helen's machine 191.1.10.5
Peter's machine 191.1.10.7
Bob's machine   191.1.10.9
```

To create a network object for this accounting department, you would enter the IP address information for this group as: 191.1.10.0 with a Subnet Mask of: 255.255.255.0. This network object, accounting department, can be used as either the source or destination for a connection.

Chapter 6. Handling Domain Name Service

This chapter explains how to configure Domain Name Service (DNS) in relation to the IBM Firewall. The goal of DNS is to provide full-domain name service to hosts inside the secure network while providing no information to hosts outside the secure network. This allows users inside the secure network to access all the services the Internet has to offer. However, by refusing to divulge information about the secure network, it makes it more difficult for an intruder to locate a computer to attack.

Three domain name servers are required to accomplish this:

1. One at the IBM Firewall
2. One inside the secure network
3. One outside the secure network

Refer to Figure 13 to see how DNS works with the IBM Firewall.




Figure 13. DNS

The Firewall is configured to act as a gateway between the nameserver(s) for the secure network and those serving the nonsecure network. The official term for the Firewall's role is *caching-only nameserver*, because the Firewall's DNS does not contain any database files itself.

Figure 13 illustrates the Firewall's role. Anytime the Firewall needs to resolve a name for its own use, it asks the secure-side nameservers. Anytime a query is forwarded to the Firewall, it in turn forwards the query to the nonsecure nameservers.

When a client on the secure network asks for secure-side information, it sends its request to the secure-side DNS, who answers. When the same client asks

for nonsecure-side information, it sends the request to the same secure-side DNS. Because the query is for nonsecure information, the secure-side DNS cannot answer, so it forwards the query to the Firewall. In the event that a nonsecure DNS were to forward a request to the Firewall, that request would be forwarded to the nonsecure DNS domain, so again no sensitive information is divulged.

Configuring DNS Using the Configuration Client

Do not use Microsoft DNS Service Manager to configure DNS. Use the IBM Firewall configuration client. To configure DNS, select System Administration from the configuration client navigation tree. Double-click the file folder icon to expand the view. Select **Domain Name Services**. The IBM Firewall displays the current DNS configuration, which you can modify.

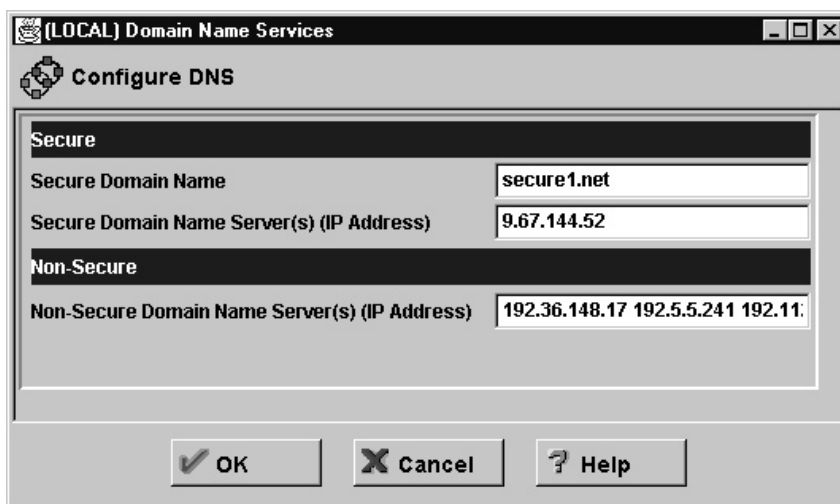


Figure 14. Domain Name Service

Note: When you add DNS, the firewall saves and renames any existing domain-name service configuration files.

1. The **Secure Domain Name** field identifies the domain name which the Firewall will append to any unqualified hostnames.
2. The **Secure Domain Name Server** field refers to the server that resolves names and IP addresses for the hosts protected from the Internet by the IBM Firewall. You can enter dotted-decimal IP addresses, separated by spaces.

3. The **Nonsecure Domain Name Server** field refers to the server(s) provided by your service provider to resolve information about the nonsecure network. You can enter dotted-decimal IP addresses, separated by spaces.

Note: When a nameserver initializes, it sends out a query to obtain the list of root nameservers. Most implementations retain this list in memory. Microsoft's implementation however, writes this list back to the configuration file. This does not modify the behavior of the nameserver, but will change the values displayed in the **Nonsecure Name Server** field. This is not cause for concern.

If the IBM Firewall is not directly attached to the Internet, add a forwarders statement to the `/etc/fwnamed.boot` file. The forwarders statement should point to a DNS on the Internet.

Configuring the Secure Name Server

For examples of how to configure the secure name server, you can refer to Chapter 4 of *DNS and Bind*. See "Bibliography" on page 199. Or, you can follow the examples given below.

In the example below, the Secure Domain Name is `sec-a.com`.

```
Secure-Network ----- Firewall ----- Nonsecure Network
(sec-a.com)                (border5)                (IBM network)
```

There are three secure mail servers behind the Firewall. You have to set up the MX records for these servers. If you do not have mail servers behind the Firewall, then you do not have to create the MX records for the mail domains.

```
50.100.143.67 limited.sec-a.com (with mail domain sec4.com)
50.100.143.68 gap.sec-a.com     (with mail domain sec6.com and it's the
Secure DNS)
50.100.143.69 express.sec-a.com (with mail domain sec5.com)
```

This is the Firewall secure interface.

```
50.100.143.65 border5.sec-a.com
```

Setting Up DNS Data

The file mapping host names to addresses is called `named.data`. In *DNS and Bind*, it is called `db.DOMAIN`. The file mapping addresses to host names is called `named.rev`. In *DNS and Bind*, it is called `db.ADDR`. There are several other data files: `named.cache` and `named.127.0.0`. In *DNS and Bind*, they are called `db.cache` and `db.127.0.0`.

If you have mail servers behind the Firewall as in this example, you have to add the MX records for each mail domain into the `named.data` or create a separate database file for each mail domain. In this case, the files are called `named.sec4.com`, `named.sec5.com`, and `named.sec6.com`.

To tie all the database files together, a name server needs a startup file. For BIND version 4, this file is usually `/etc/named.boot`. For BIND version 8, this file is usually `/etc/named.conf`.

Here are the completed Data Files:

named.boot

```
primary      sec-a.com /etc/named.data
primary      sec4.com  /etc/named.sec4.com
primary      sec5.com  /etc/named.sec5.com
primary      sec6.com  /etc/named.sec6.com
primary      in-addr.arpa /etc/named.rev
cache        .         /etc/named.ca

forwarders   50.100.143.65
```

named.data

```
sec-a.com.   IN SOA gap.sec-a.com. root.gap.sec-a.com. (
                1000 3600 300 3600000 86400)

sec-a.com.   IN NS      gap.sec-a.com.

localhost.sec-a.com. IN A 127.0.0.1
loopback.sec-a.com. IN A 127.0.0.1
gap.sec-a.com.   IN A 50.100.143.68
limited.sec-a.com. IN A 50.100.143.67
express.sec-a.com. IN A 50.100.143.69
border5.sec-a.com. IN A 50.100.143.65
```

named.rev

```
@   IN SOA gap.sec-a.com. root.gap.sec-a.com. (
                1000 3600 300 3600000 86400)

143.100.50   IN NS      gap.sec-a.com.

1.0.0.127 IN PTR localhost.sec-a.com.
1.0.0.127 IN PTR loopback.sec-a.com.
68.143.100.50 IN PTR gap.sec-a.com.
67.143.100.50 IN PTR limited.sec-a.com.
69.143.100.50 IN PTR express.sec-a.com.
65.143.100.50 IN PTR border5.sec-a.com.
```

named.sec4

```
sec4.com. IN SOA gap.sec-a.com. root.gap.sec-a.com. (
    1000 3600 300 3600000 86400)
```

```
sec4.com. IN MX 10 limited.sec-a.com.
```

named.sec5

```
sec5.com. IN SOA gap.sec-a.com. root.gap.sec-a.com. (
    1000 3600 300 3600000 86400)
```

```
sec5.com. IN MX 10 express.sec-a.com.
```

named.sec6

```
sec6.com. IN SOA gap.sec-a.com. root.gap.sec-a.com. (
    1000 3600 300 3600000 86400)
```

```
sec6.com. IN MX 10 gap.sec-a.com.
```

You can put the MX record `sec4.com. IN MX 10 limited.sec-a.com.` into the `named.data` file. But if you do not have a mail server behind the Firewall, then you do not have to worry about the MX record data files.

The Root Cache Data

For the `named.ca` file, you can use anonymous ftp to `ftp.rs.internic.net` to retrieve the file `named.root` from the domain subdirectory. Rename the `named.root` file to `named.ca`.

Point the Unresolved Queries to the Firewall

The secure name server must be configured to forward unresolved queries to the Firewall. If you have a standard BIND implementation, add a `forwarders` statement and a `cache` statement to the boot file on your secure name server as shown in the `named.boot` example.

```
forwarders    50.100.143.65
cache         .      named.ca
```

Create the cache file, `named.ca`, to point to the Firewall:

```
. 99999999 IN NS border5.sec-a.com
border5.sec-a.com 99999999 IN A 50.100.143.65
```

where `sec-a.com` is the domain name used from the secure side and `50.100.143.65` is the Firewall's IP address.

In addition, you might want to add your firewall's host name to the DNS databases. This way your users can access the Firewall's Socks server, HTTP proxy, Telnet proxy, and FTP proxy using the Firewall's hostname instead of

its IP address. This requires two additional steps as shown in the example above. (Please refer Chapter 4 of DNS and BIND.)

First add an A record to the domain database file named `.data`:

```
border5.sec-a.com. IN A 50.100.143.65
```

Then add a PTR record to the reverse-lookup file named `.rev`:

```
65.143.100.50      IN PTR border5.sec-a.com.
```

If you do not use DNS for your secure network, your firewall must still be able to resolve its own information. Configure the firewall as described for the normal case, but list the firewall's secure interface in the Secure Domain Name Server field. Then add the following line to `/etc/fwnamed.rev`:

```
primary 143.100.50.in-addr.arpa /etc/fwnamed.boot
```

Then create `fwnamed.rev` to resemble the following:

```
143.100.50.in-addr.arpa      IN SOA border5.sec-a.com. root.border5.sec-a.com.
(
    9      ; Serial
    86400  ; Refresh after 1 day
    300    ; Retry after 5 minutes
    654000 ; Expire after 1 week
    3600)  ; Minimum TTL of 1 day
143.100.50.in-addr.arpa.    IN NS   border5.sec-a.com
65.143.100.50.in-addr.arpa  IN PTR  border5.sec-a.com
```

Configuring the Secure Clients

Clients on the secure network must be configured to send their queries to the secure nameserver, not to the Firewall. This is important because it ensures that no secure-side information is stored in the Firewall's in-memory cache. Also, it saves workload on the Firewall because the Firewall will not get involved unless a query involves forwarding a query from the secure side to the nonsecure side.

If you do not use DNS for your secure network, your clients will have to point to the Firewall as their nameserver.

Publishing Services to the Public

Many organizations want to publish particular services to the Internet public. Often, these services include e-mail and Web servers, although any type of TCP/IP server could be used. In order to make such services available, you must not only place the server on the network where it can be reached, but you must also list that server with the public DNS, so that users can obtain the right information.

There are two ways to accomplish this. Either your service provider will list your servers as a part of their domain (and hence on their nameservers), or you must provide your own nameserver and register it with the Internet. It is by far easier for your Internet Service Provider (ISP) to provide this service for you. If you can choose this option, you need to provide them with the hostnames and IP addresses you wish to have listed. For example, if you operate your public Web server as *www.public.com* at IP address is *50.100.150.200*, you need to ask your ISP to list *www.public.com* at *50.100.150.200*.

In addition, if you wish to receive e-mail, you should ask your ISP to list your firewall as the *mail exchanger* for your public e-mail domain. The ISP needs to know the hostname (*gateway.public.com*), its IP address (*50.100.150.201*), and the domain name by which you want to receive mail (*public.com*).

If your ISP is not willing to provide these services for you, then you will have to do it yourself. Here again, you have two additional choices. You can place a DNS server in your DMZ or you can use your firewall as that nameserver. Using the firewall does not open additional security risks because the database files you will put there do not contain any information about your secure network. The only information that will be stored will pertain to the public services you choose to offer.

The details involved in setting up a DNS server are contained in Chapter 4 of *DNS and BIND*, which is listed in the “Bibliography” on page 199. That chapter is highly-recommended reading, as are the preceding chapters, if necessary. Setting up a DNS server is not a trivial task and is often best left to experts. If you have such an expert available, seriously consider taking advantage of that expertise.

See “Sample Configurations” on page 42 for more information.

Installing Microsoft's DNS Server

To install Microsoft's DNS Server, go to the control panel, click **Network**, click **Services tab**, click **Add**, and select **Microsoft DNS Server**. You will need the installation CD-ROM.

Sample Configurations

This section illustrates some sample configurations in which a firewall might be deployed. Most of these examples focus on the configuration necessary for DNS operation. It is unlikely that one of these examples illustrates your network, so take care to understand each example and to apply the appropriate concepts to your particular installation.

Example 1: DNS Server in a DMZ on the Nonsecure Interface

The first example illustrates the files needed to operate the nameserver in a DMZ which is located inside the nonsecure network, as shown in Figure 15.

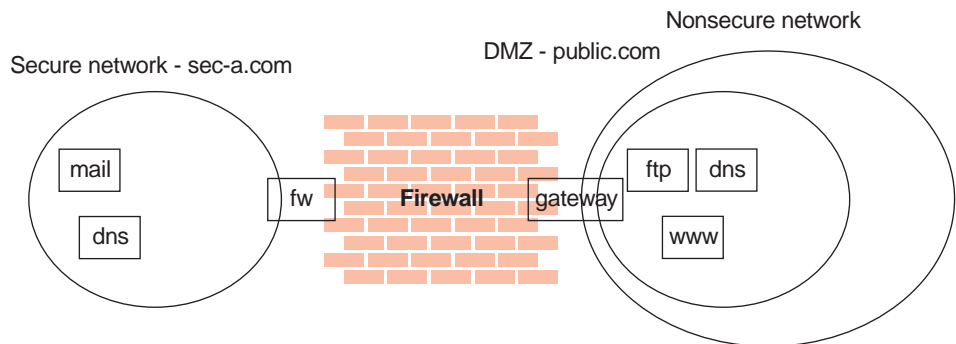


Figure 15. Nameserver in DMZ Inside Nonsecure Network

This figure illustrates a private network, *sec-a.com*, behind an IBM Firewall whose secure interface is named *border5.sec-a.com* and whose nonsecure interface is named *gateway.public.com*. The company's DMZ is attached to the nonsecure interface and contains a nameserver *dns.public.com*, an FTP server *ftp.public.com*, and a Web server *www.public.com*. The files on *dns.public.com* to implement this scenario are as follows:

db.public

```
public.com.  IN SOA dns.public.com. admin.public.com. (
                1                ; serial number
                10800             ; refresh after 3 hours
                3600              ; retry after 1 hour
```

```

        604800      ; expire after 1 week
        86400 )    ; minimum TTL 1 day
;
; Nameservers
;
public.com      IN NS  dns.public.com.
;
; Hosts in the DMZ
;
dns.public.com.  IN A  50.100.150.202
gateway.public.com. IN A  50.100.150.201
www.public.com.  IN A  50.100.150.200
ftp.public.com.  IN A  50.100.150.203
;
; Mail-related entries
;
public.com.      IN MX  0  gateway.public.com.
public.com.      IN CNAME gateway.public.com.

```

db.50.100.150

```

150.100.50.in-addr.arpa.  IN SOA dns.public.com. admin.public.com. (
        1          ; serial number
        10800     ; refresh after 3 hours
        3600      ; retry after 1 week
        604800    ; expire after 1 week
        86400 )   ; minimum TTL 1 day
202.150.100.50.in-addr.arpa.  IN NS dns.public.com.
203.150.100.50.in-addr.arpa.  IN PTR ftp.public.com.
202.150.100.50.in-addr.arpa.  IN PTR dns.public.com.
201.150.100.50.in-addr.arpa.  IN PTR gateway.public.com.
200.150.100.50.in-addr.arpa.  IN PTR www.public.com.

```

db.127.0.0

```

0.0.127.in-addr.arpa.  IN SOA dns.public.com. admin.public.com. (
        1          ; serial number
        10800     ; refresh after 3 hours
        3600      ; retry after 1 week
        604800    ; expire after 1 week
        86400 )   ; minimum TTL 1 day
0.0.127.in-addr.arpa.  IN NS  dns.public.com.
1.0.0.127.in-addr.arpa. IN PTR localhost.

```

db.cache

The best choice for this file is to FTP the current root nameserver list from *ftp://ftp.rs.internic.net/domain/named.root*.

boot

```

primary public.com          db.public
primary 150.100.50.in-addr.arpa db.50.100.150
primary 0.0.127.in-addr.arpa db.127.0.0
cache .                     db.cache

```

To set the traffic filter to allow the appropriate DNS traffic, enable *Permit DNS Queries* on the **Security Policy** panel.

Example 2: DNS in a DMZ on a Dedicated Interface

In the second example, the DNS for the DMZ is still on a dedicated nameserver, but this time the DMZ is attached to a distinct interface instead of the same interface as the nonsecure network.

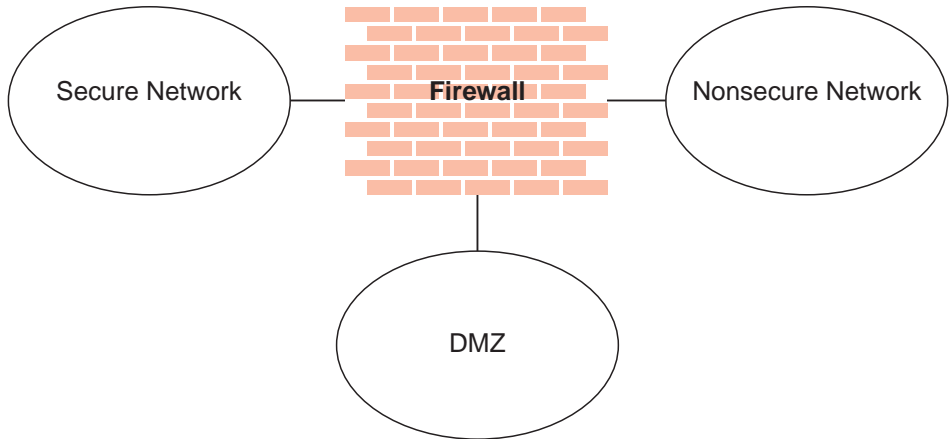


Figure 16. DNS in a DMZ on a Dedicated Interface

The DNS data files on *dns.public.com* are the same as in the preceding example. In order to make that nameserver accessible to the public network, though, it is necessary to either open the traffic filter or to perform a zone transfer to copy the data files to the Firewall.

To open the traffic filter, copy the three rule templates entitled *DNS Server queries*, *DNS Replies*, and *DNS Client queries*. Change the routing setting on each rule from *local* to *routed*. Then include the three new rule templates in a service and set the flow indicators as follows:

- DNS Client queries: --->
- DNS Replies: <---
- DNS Server queries: --->
- DNS Server queries: <---

Include this service in a connection which uses *The World* as the source object and *dns.public.com* as the destination object.

To perform a zone transfer, you need to both set the traffic filter and instruct the nameservers to copy the appropriate files. To set the traffic filter:

1. On the **Security Policy** panel, enable *Permit DNS Queries*.
2. Add a connection from *dns.public.com* (source object) to the Firewall's DMZ interface (destination object), which includes the service entitled *DNS Transfers*.

To activate the zone transfer, add the following lines to the Firewall's *boot* file in `c:\winnt\system32\dns`:

```
secondary public.com      50.100.150.202 db.public
secondary 150.100.50.in-addr.arpa 50.100.150.202 db.50.100.150
```

Then go to the Service Control Manager and stop and restart the DNS Server service.

Example 3: Using the Firewall as the Secure Nameserver

To use the Firewall as your secure name server, place the database files which would normally reside on the secure server, on the Firewall. Then your clients can point to the Firewall as their DNS server. The risks associated with this approach are that the DNS server cannot tell a request from the secure side from a request from the nonsecure side. Accordingly, it will provide this secure-side information to any client who asks; you no longer can hide your secure DNS information.

If you want your public side information on the other machine outside your firewall (for example, an ISP or DMZ), so that your firewall just holds the secure side information, see "Publishing Services to the Public" on page 41.

To implement this approach, follow these steps:

1. Place the database files: `named.data`, `named.rev`, `named.sec4.com`, `named.sec5.com`, and `named.sec6.com`, which would normally reside on the secure server, onto the Firewall. And change the SOA in the database files to point to the Firewall `border5.sec-a.com` as shown in the example in "Configuring the Secure Name Server" on page 37.

named.data

```
sec-a.com.      IN SOA border5.sec-a.com. root.border5.sec-a.com. (
                    1000 3600 300 3600000 86400)
```

```
sec-a.com.      IN NS      border5.sec-a.com.
```

```
localhost.sec-a.com.      IN A 127.0.0.1
```

```
loopback.sec-a.com.      IN A 127.0.0.1
```

```
gap.sec-a.com.      IN A 50.100.143.68
```

```
limited.sec-a.com.      IN A 50.100.143.67
```

```
express.sec-a.com.      IN A 50.100.143.69
```

```
border5.sec-a.com.      IN A 50.100.143.65
```

named.rev

```
@ IN SOA border5.sec-a.com. root.border5.sec-a.com. (
1000 3600 300 3600000 86400)
```

```
143.67.9 IN NS border5.sec-a.com.
```

```
1.0.0.127 IN PTR localhost.sec-a.com.
```

```
1.0.0.127 IN PTR loopback.sec-a.com.
```

```
68.143.100.50 IN PTR gap.sec-a.com.
```

```
67.143.100.50 IN PTR limited.sec-a.com.
```

```
69.143.100.50 IN PTR express.sec-a.com.
```

```
65.143.100.50 IN PTR border5.sec-a.com.
```

named.sec4.com, named.sec5.com and named.sec6.com are the MX records for the mail servers behind the firewall. You can put the line for the MX record of each file into the named.data, so you do not have to create a separate file just for the MX record. For example:

```
sec4.com. IN MX 10 limited.sec-a.com
```

2. Edit the /etc/fwnamed.rev file or \WINNT\system32\DNS\fwnamed.rev file. For example, the firewall's secure interface is named border5.sec-a.com and its IP address is 50.100.143.65, and there are three secure clients behind the Firewall (50.100.143.67, 50.100.143.68, 50.100.143.69), which copy the information from /etc/named.rev.

fwnamed.rev

```
143.100.50.in-addr.arpa. IN SOA border5.sec-a.com root.border5.sec-a.com (
1 10800 3600 604800 86400 )
```

```
143.100.50.in-addr.arpa. IN NS border5.sec-a.com.
```

```
68.143.100.50.in-addr.arpa. IN PTR gap.sec-a.com.
```

```
67.143.100.50.in-addr.arpa. IN PTR limited.sec-a.com.
```

```
69.143.100.50.in-addr.arpa. IN PTR express.sec-a.com.
```

```
65.143.100.50.in-addr.arpa. IN PTR border5.sec-a.com.
```

3. Add the following lines to the Firewall's /etc/fwnamed.boot file or \WINNT\system32\DNS\boot file:

fwnamed.boot

```
; Created by IBM Firewall 1999182133
```

```
cache . /etc/fwnamed.ca
```

```
primary 0.0.127.in-addr.arpa /etc/fwnamed.loc
```

```
primary 143.100.50.in-addr.arpa /etc/fwnamed.rev
```

```
primary sec-a.com /etc/named.data
```

```
primary sec4.com /etc/named.sec4.com
```

```
primary sec5.com /etc/named.sec5.com
```

```
primary sec6.com /etc/named.sec6.com
```

```
forwarders 50.100.144.84
```

If the Firewall is not on the Internet, add the statement for the forwarders in /etc/fwnamed.boot, Which 50.100.144.84 is the DNS on the Internet.

4. For the Secure Domain Name field on the configuration client, list the Firewall's secure interface, which is the domain name you will be using on your secure network.

For the Nonsecure Nameserver, list the nameserver provided by your ISP, as usual.

5. Testing the DNS setup with nslookup:

```
% nslookup 50.100.200.3
% nslookup sendmail.dmz.dom
% nslookup
>set type=mx
>us.ibm.com
```

The result should return the machine's hostname and IP address.

6. Testing the DNS setup with connecting to the proxy servers on the Firewall.

You should be able to connect to the proxy servers (telnet, ftp, httpd) on the Firewall by using the Firewall's hostname or IP address.

7. Testing the DNS setup with mail.

This means that any secure-side client that wants to connect to the configuration server or any of the proxy servers on the Firewall, can refer to the Firewall by IP address or hostname

Chapter 7. Secure Mail Proxy

The IBM Firewall Secure Mail Proxy provides a gateway for SMTP traffic. It relays messages from the secure mailserver to the nonsecure side, hiding sensitive domain names as it goes. It relays messages from the nonsecure side into the secure mail domain and insulates the secure network from attacks.

The Secure Mail Proxy relays messages in real time from the sender to the receiver. This is to avoid the risks and complexity involved with maintaining a message queue on the Firewall. This necessitates certain configuration requirements upon the adjacent mail domains. Read this chapter thoroughly as you design your implementation.

How the Secure Mail Proxy Works

The Secure Mail Proxy acts as a real-time gateway between two or more e-mail domains. In contrast with a traditional SMTP relay, messages are not stored on the Firewall before being forwarded to the destinations. The SMTP conversation is interpreted as it happens, and the Secure Mail Proxy conversation is proxied on to each of the necessary destination servers, command by command.

When an SMTP server opens an SMTP conversation with the Firewall's Secure Mail Proxy, the SMTP conversation takes place between the proxy and the sending server until the sending server sends the list of recipients. Then, as each recipient is sent, the Secure Mail Proxy opens a new SMTP conversation with each of the necessary recipient servers. Then, as the body of the message is sent, it will be fanned out (as it comes in) to each of the recipient servers.

The benefits are:

- Because messages are not stored on the Firewall before transmittal, the messages need not be stored for long periods of time.
- The complexity of a queue-management component is unnecessary.
- Misleading positive results are not generated because of synchronous transmission. If an error is encountered with one or more recipient, the sending Secure SMTP server can be notified of that error immediately. See "Error Handling" on page 50.

Error Handling

The SMTP world is an imperfect one. Errors occasionally occur. The proxy works in conjunction with a traditional store and forward mail server to recover from rare timing conditions and when it is unable to establish connections to a mail server.

The Firewall's Secure Mail Proxy tries to be invisible to errors, so that a sending SMTP server is free to respond to any error it encounters according to the SMTP server's own implementation, with no interference from the Firewall's proxy. This approach is unsuitable for certain types of errors that might occur. In particular, errors which take place during transmission of the body of the message are difficult to handle, because such an error will usually affect only one recipient server, while the remaining recipient servers are unaffected.

To accommodate such a case, the Secure Mail Proxy will keep a copy of each message as it is being sent. In the case of an error, this copy will be sent after the successful transmission to all recipient servers that did not encounter an error to a configured overflow server. This overflow server must be a conventional store-and-forward type server, and must be configured to relay messages which might be addressed to either secure-side or nonsecure-side destinations. It will be this server which will implement the policy regarding retry attempts and delivery-failure notification.

Although the proxy does keep a copy of each note, it is important to note the distinction between the proxy's behavior and that of a conventional store-and-forward server. After the normal SMTP conversation, if no errors are encountered, the Firewall's copy of the message is immediately discarded. In the case of a store-and-forward server, the message must persist on the Firewall for the entire duration of:

- The SMTP conversation in which the message is received from the sender
- The latency time involved with the queueing system on the server
- The one or more SMTP conversations required to transmit the message to each receiving SMTP server, including potential additional queue latency while the list of recipient servers is processed.

In the case of an error, the store-and-forward server must store the message long enough to implement its retry policy, which could require hours or days. In the case of an error on the Firewall, the message is relayed immediately to the overflow server, and is therefore no longer consuming Firewall resources.

Note: The proxy keeps its copy of each message in the Firewall's tmp subdirectory. There is only one copy of each message kept, regardless of

the number of recipients on each message. These files are discarded immediately either after successful transmission or after being relayed to the overflow server.

Fan-Out Limit

As described above, when the Secure Mail Proxy receives a message destined to more than one destination domain, it will *fan-out* the message and deliver it simultaneously to each destination domain's mail server. The proxy will only open a certain number of these outbound connections; destinations exceeding that limit will be forwarded to the overflow server. Fan-out is limited due to the proxy's packet-by-packet transmission model. Each time a packet is received, it must be sent to the entire set of receiving servers. If the set of receiving servers is allowed to grow too large (or if that set contains exceptionally slow servers), the sender might time out, because the proxy cannot process the transmission quickly enough.

This fan-out limit **does not** affect the following types of messages:

- Messages sent to many users on a single domain (or a few domains)
- Messages sent to listservers (such as Majordomo)
- Messages sent into the secure network (usually there are fewer secure-side domains than the fan-out limit)

The fan-out limit **does** affect the following types of messages:

- Messages sent from a secure-side listserver to a wide assortment of nonsecure-side participants
- Messages sent to a large number of nonsecure domains

Overflow Server

The Securemail SMTP proxy does not store and forward messages in order to provide higher levels of security and firewall reliability. In most cases it is desirable to have the store and forward activity occur on the native mail system and not on an intermediate server. However, there are three conditions where it may be desirable to create a connection to an overflow server to handle processing that the proxy is not able to handle.

1. When the proxy receives a message it attempts to create connections to the outbound hosts when RCPT commands are received. As each new recipient (RCPT) command is received it either opens a new connection or reuses an existing connection. The opening of outbound connections is limited by a configurable MAX_FAN_OUT, which is limited to 32 to keep a single message from dominating proxy and machine resources.

When this limit is reached then if an overflow server is configured the recipients that do not fit into existing connections will be sent to the overflow server. Most messaging systems will not require this usage since

a RCPT command issued to a recipient, which would have gone to a nonconfigured overflow server will have a temporary error returned and the message should be retired for those recipients later.

There are, however some mail systems where the retry logic may not work transparently and it is desirable to off-load the processing to an overflow server. An alternative is to configure the mail system to limit the number of recipients per message to avoid attempting to go to the overflow server.

2. The overflow server can serve as a store and forward server for queuing messages when the destination server is not available. With most messaging systems the temporary error returned as a response to the RCPT command when a host is known to DNS but is not available for a connection, causes the message to be requeued for later retransmission. However, for some messaging systems the requeue is visible to the users as a protocol error and it is desirable to have the messages moved off the messaging system on top of another server regardless of the success of the final delivery. Also, while most messaging systems have reasonable queuing mechanisms there may be cases where it is preferable to have retransmission queuing occur at another server other than the messaging server; the overflow server can serve this purpose.
3. If a failure occurs during the transmission of the DATA segment of a message after the RCPT commands have flowed and responded to, then the overflow server can serve the function of receiving the recipients for retransmission. If no overflow server is installed, then the message will be resent to all of the recipients, which is normally not an issue with messaging systems because they usually contain code to detect and discard duplicate messages. However, with some messaging systems it may be preferable to use an overflow server for those cases in particular where large messages are routinely sent over unreliable lines.

Why You Might Want to Install the Overflow Server

You might want to install an overflow server for the the following reasons:

- Your messaging system has no ability to limit the number of recipients per message being sent outbound and does not tolerate temporary errors acceptably.
- Your messaging system does not handle large queues or retransmission of messages when they cannot be immediately delivered gracefully and you have another store and forward messaging component available.
- Transmission of large messages over unreliable lines is causing duplicate messages frequently enough to justify having a fallback server.

Generally, if your messaging system can tolerate and properly requeue messages, which are responded to with temporary error conditions, then the installation of an overflow server is not required.

Advantages to Not Installing the Overflow Server

If you can avoid installing an overflow server, then you have the following advantages:

- Messages are only queued in the messaging system thus making message disposition and tracking questions less complicated.
- Maintenance of queued messages outside of the messaging system is eliminated.

Basic Configurations for the Overflow Server

There are three basic configurations for an overflow server:

1. Use of the primary messaging system as the overflow sever
2. Use of a separate messaging component like sendmail on the firewall
3. Use of a separate messaging component on another server

Use of Primary Messaging System as the Overflow Server: The communication between components occurs on the standard port 25. Note that the primary messaging system must have the ability to delay the retransmission when messages are sent over the overflow path to stop mail loops.

Use of a Separate Messaging Component like Sendmail on the Firewall: The store and forward messaging component must listen on a non-standard port so as not to interfere with the proxy listening on port 25. Both the proxy and the store and forward component are able to send outbound, but only the secure mail proxy is able to receive. Filters should be set up to exclude traffic to the overflow port from the outside to avoid any security related issues with the store and forward component if, for example, sendmail is used.

This configuration has the advantage of not requiring another server, but has a potential for causing instability with the firewall if the store and forward queues are not isolated from causing disruption if the queuing space becomes full.

Use of a Separate Messaging Component on Another Server: The store and forward either has the ability to send outbound directly or it can tunnel the traffic back through the proxy (assuming that the store and forward component will not get into loops sending messages by being able to delay retransmission of messages where a host is temporarily not available).

This configuration requires another server but has the advantage of not storing any content on the firewall itself and minimizing disruption of the firewall's operations.

Configuring the Secure Mail Proxy Using the Configuration Client

Before you configure the Secure Mail Proxy, read the following configuration information which will help you when designing your installation.

- Because the proxy is rewriting private domain names into public domain names, the set of private domains and the sets of users on those private domain names must be arranged to accommodate the ambiguity which can result from the domain-name rewrite operation. For example, if there is a user "joe@hq.private.com" and another user "joe@sales.private.com", and if the proxy is configured to rewrite "private.com" into "public.com", then both of these users will appear to the outside world as "joe@public.com".

To solve this ambiguity, you would configure the proxy to rewrite "hq.private.com" into "hq.public.com", and "sales.private.com" into "sales.public.com", rather than rewriting "private.com" into "public.com".

Substitution is performed on all mail sections except the subject and body sections of the mail.

- If your network contains multiple Firewalls or other outbound mail gateways, particularly if these gateways (and the private networks) contain expensive network media (such as private transatlantic trunk lines), it is easy to introduce a situation in which replies get routed by the Internet to the wrong gateway, causing unnecessary cost on this expensive media. For example, a company with a North American Firewall and a European Firewall might choose to substitute "private.com" to "public.com" across all of their Firewalls. In such a case, a European customer of this company might send a note to "sales@public.com", which could (based upon MX records) be sent to the North American Firewall instead of the European Firewall causing the company to unnecessarily pay for the message to be routed over the transatlantic T1 back to Europe. To avoid this situation, consider using different public domain names for different proxies, if an expensive medium separates them (in this example, perhaps "na.public.com" and "ec.public.com").

To configure the Secure Mail Proxy, select **System Administration** from the configuration client navigation tree. Double-click the file folder icon to expand the view. Select **Secure Mail Proxy**. Double-click the file folder icon to expand the view. There are various options you can work with.

List or Add a Mail Domain Entry

To list or add a mail domain, select **Managed Mail Domains**. The IBM Firewall displays the list of both secure and nonsecure managed mail domains.

To add a secure or nonsecure mail domain, select **NEW** and click **Open**. The **Domain Name** dialog box appears.

1. Fill in your domain name.
The **Domain Name** field contains the name by which the mail domain being described is known to users on the secure side of the firewall. The value for this field must be unique among all entries.
Any subdomains associated with a particular domain name will automatically be included.
2. Check *Secure domain* if you want this domain name to be secure.
3. The **Mail Exchanger** field contains the host name of the mail server to which this entry applies. This server must be on one of the secure networks. You can have up to five mail exchangers in your list. The order of the mail exchangers is important. Use the Open, Add, Edit, Delete, Move Up and Move Down buttons to add to or adjust your list of mail exchangers.
4. If you want your mail exchangers to be load balanced, check the load balancing box.
You can have up to five mail exchangers. If you do not select load balancing, the first mail exchanger will be contacted. If it fails, the second mail exchanger will be contacted and so forth. If you select load balancing, the mail exchangers are randomly selected.

Change a Mail Domain Entry

To change a mail configuration entry, select an entry in the list and click **Open**. The **Managed Domain Details** dialog box appears.

Delete a Mail Domain Entry

To delete a mail configuration entry, select an entry in the list and click **Delete**. You will get a delete warning. Click **OK** to delete.

Excluded Mail Domains

To create a list of excluded domains, select **Secure Mail Proxy**.

1. Double-click the file folder icon to expand the view and select **Excluded Mail Domains**. A list of domains that are excluded from being able to send messages through the Firewall is displayed.
2. To add to this list, select **NEW** and click **Open**. The **Add** dialog box appears.
3. Fill in the domain name you want to exclude.
Any subdomains associated with a particular domain will automatically be excluded.
4. Click **OK**.

Domain Name Hiding

You can hide your secure domains behind public domain names for additional security. The **Public Mail Domain Name** field contains the name by which the mail domain being described is known to users on the nonsecure side of the firewall. This name will be substituted in place of the secure domain name, in order to hide the topography of the secure network.

To create a hidden domain pair, select **Secure Mail Proxy**. Hidden domain name pairs are secure because they are on the secure side of the Firewall.

1. Double-click the file folder icon to expand the view and select **Domain Name Hiding**.
2. Select **NEW** and click **Open**. The **Add Hidden Mail Domain** dialog box appears.
3. Enter your secure mail domain name. The secure mail domain name specifies the name of your network, the domain protected by the Firewall. Note that the fully qualified domain name is required. This domain name must also exist as a secure domain in the Managed Mail Domain.
4. Enter the public mail domain name. The public mail domain specifies the name the nonsecure network knows your secure network by. For example, if your network is **secnet.idaho.edu** your public domain name is **idaho.edu**. The firewall name server uses this domain and its name servers when resolving outside host names. Note that the fully qualified domain name is required. This domain name must also exist as a secure domain in the Managed Mail Domain List.
5. Click **OK**.

Proxy Characteristics

You can configure mail proxy characteristics, like anti-SPAM and anti-forgery security enhancements for messaging.

To configure mail proxy characteristics, select **Secure Mail Proxy**. Double-click the file folder icon to expand the view. Select **Proxy Characteristics**. The **Proxy Characteristics** notebook dialog appears, as shown in Figure 17 on page 57.

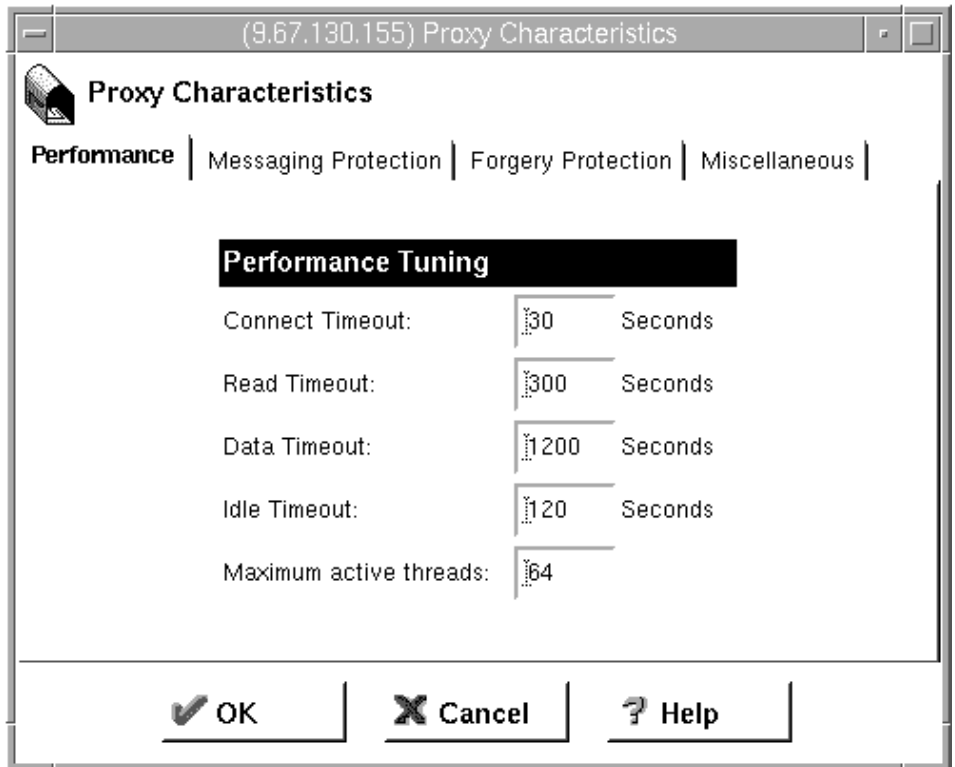


Figure 17. Mail Notebook Control

1. Select **Performance** if you want to do some performance tuning for your mail proxy. You can adjust the following parameters:

Connect Timeout

How many seconds to wait while trying to establish a connection to another system.

Read Timeout

How many seconds to wait while trying to read data from another system.

Data Timeout

How many seconds to wait while trying to read data from the transfer of the content of a message.

Idle Timeout

How many seconds to wait before timing out and terminating an idle session.

Maximum Active Threads

Enter a value for the number of mail clients you would like to have simultaneously without waiting.

2. Select **Messaging Protection** if you want to enhance the security of your mail. The anti-SPAM security enhancement is included here with the following configuration options:
 - On - Use the algorithm described to verify whether to accept the message
 - Off - Do not use the algorithm described
 - Warn - Use the algorithm described but only log a warning message instead of rejecting the message, which is useful for determining if an option should be enabled

You can set the following parameters:

Block invalid originators

Select Yes, No, or Warn. If you select Yes, an attempt is made to resolve the domain name of the originator using DNS to see if the message is replyable. If the message is not replyable, then the message may be rejected. Most SPAM is sent with non-replyable addresses.

Maximum recipients

Enter a value for the maximum number of recipients for a single mail item.

Maximum size

Enter a value for the maximum size of the mail in bytes.

Block relay mail

Choose Yes or No. In most cases no relay of messages should be occurring through a firewall unless either the originator or a recipient is within the domains managed by the organization hosting the firewall. Otherwise, the messaging transfer is not to the benefit of the owner of the firewall. For most deployments, the firewall will have this setting turned on. However, there are some deployment scenarios where the firewall between the DMZ and the internal network would want to relay all messages. See "Sample Configurations" on page 42 for more information on DMZs. You configure the domain mapping feature to determine whether or not a domain is local.

Maximum fan-out

The maximum number of outbound sessions per each inbound connection. The default is 3.

3. Select **Forgery Protection** if you want to enhance the security of your mail. The anti-forgery security enhancement is included here with the following configuration options:
 - On - Use the algorithm described to verify whether to accept the message
 - Off - Do not use the algorithm described
 - Warn - Use the algorithm described but only log a warning message instead of rejecting the message, which is useful for determining if an option should be enabled

You can set the following parameters:

Block HELO Spoofing

Select Yes, No, or Warn. If you select Yes, the secure mail proxy will reject connections where the hostname of the sending system, as determined by reverse DNS, does not match the hostname provided by the sender in the HELO command. The default is No.

Block originator forging

Select Yes, No, or Warn. Selecting Yes enables checking for the believability of the identity of the originator based upon whether the IP address of the sending message transfer agent is in the same classification as the domain name of the originator. For instance, if the IP address is local but the domain name is not local, then impersonation could occur and likewise if the domain name is local but the IP address is not local. This simple algorithm can block a high percentage of the casual prank messages such as an internal employee trying to send a message as if it came from someone famous outside of the organization, or someone outside of the organization trying to send a message from the president of the organization. However, in the case of list servers, this feature would need to be disabled.

Block host without DNS entry

Choose Yes, No, or Warn.

4. Select **Miscellaneous** if you want to list overflow mail servers or enable or disable domain name hiding.

Configuring the Overflow Server

The overflow server should be placed on the secure side of the network.

To set up the overflow server, click **Miscellaneous** on the **Proxy Characteristics** panel. You are prompted for the hostname (or address) and port number for the overflow server. You can use "localhost" or "127.0.0.1" if the overflow server resides on the Firewall machine. This IP address is not filtered by the Firewall's traffic control.

Click **OK** to save your changes, or **Cancel** if you change your mind.

Configuring the Secure Servers

You must configure your secure mail servers to list the Firewall as their gateway for unknown domains. This causes mail intended for the nonsecure network to be forwarded to the Firewall. Also, each server must be configured to accept messages addressed to their public domain name in addition to their private domain name. When the Firewall forwards a note from the nonsecure network, all recipients will be listed with their public-side domain names.

If you have more than a single distinct mail domain inside your secure network, you must also configure each server to forward mail intended for another secure-side domain directly to that server, not through the Firewall. This relieves the Firewall of unnecessary workload and allows the Firewall's real-time delivery mechanism to function properly.

Configuring the Public Domain

The only configuration necessary in the nonsecure network is to list your Firewall as the mail exchanger for your network. Ask your service provider to add the necessary information to their DNS servers. See "Chapter 6. Handling Domain Name Service" on page 35 for additional specifics regarding the mechanics involved.

The objective is to list your Firewall as the *mail exchanger* for each public domain name for which you want to accept mail. For example, if you use the domain name *private.com* inside your secure network and *public.com* outside your secure network, you might name your firewall *gateway.public.com*. In such a case, you would ask your provider to list the Firewall's hostname and IP address as a host (which will usually be listed with "A" records and "PTR" records). Then, because you want to accept mail addressed to *user@public.com*, you would ask your provider to add an MX record for the domain *public.com* which lists *gateway.public.com* as the mail exchanger for that domain. If you also want to receive mail addressed to *user@somethingelse.com*, you can list an additional MX record which also points to the Firewall.

Using an SMTP Server Instead of the Secure Mail Proxy

The following sections tell you how to disable the Secure Mail Proxy and how to configure an SMTP server.

Disabling the Secure Mail Proxy

To disable the Secure Mail Proxy in order to avoid conflicts with another SMTP server product, disable the Secure Mail Proxy service from the **Service Control Manager**. From the Windows **Start** menu, select **Settings, Control Panel, Services**. Scroll to select *IBM Firewall SMTP Proxy*. Click **Startup**. In the **Startup Type** field, select **Disabled**. Click **OK**.

Configuring an SMTP Server

You need to consider several aspects when installing a full SMTP server in place of the Secure Mail Proxy. This section describes the security features of the Secure Mail Proxy, in an attempt to allow you to configure your SMTP server to perform similar functions. Certain SMTP server products might be unable to perform some of these tasks, so study the choices available and your needs carefully before purchasing a product.

There are certain attacks which attempt to overflow or otherwise corrupt the mail queue. Although no full-blown server can operate without a mail queue, the risks associated with the mail queue are reduced if you can dedicate a disk volume exclusively to that task. This minimizes the chances that an overflowed queue would impact other operations of your firewall.

It is also important that your mail server hide information about the secure network. The Secure Mail Proxy substitutes all private-side hostnames to the public domain name. This removes information that could be used to map your network.

Chapter 8. Controlling Traffic through the Firewall

This chapter tells you how to use the configuration client to control network traffic through the Firewall. Using expert filters, the firewall filters packets at the session level based upon multiple criteria such as time of day, IP address, and subnet. The filter acts between the secure and nonsecure network interfaces. Filters do not impact the firewall routing tables.

By default the Firewall does not allow any traffic to flow between the secure and nonsecure network. You must create connections to allow specific types of traffic to flow between the secure and nonsecure networks.

Versatile Filters Configuration

You can include or exclude particular connections, which will take affect when the filter rules are regenerated.

Using the Configuration Client to Build Connections

You can rearrange the dynamic and static filter rules depending upon how much granular control you would like to have. Static filter rules provide you with more granular control.

You use the components of the configuration client illustrated in Figure 18 on page 64 to create network objects, rule templates, services, and connections.

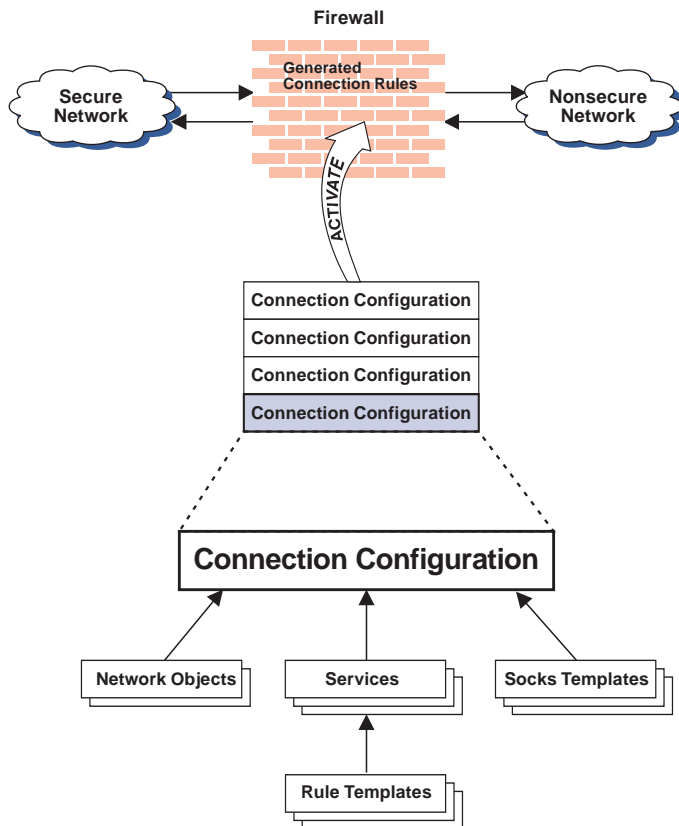


Figure 18. Building Connections

Connections

Associate network objects with services or socks templates to define the types of communications allowed between endpoints. Each connection defines a specific type of IP traffic to be allowed or denied between a source and destination network object.

Services

Are built of one or more rule templates. Defines the type of IP traffic that is permitted or denied between a source and destination object. For example, you could construct a service to permit Telnet or deny Ping. (One of the FTP services is comprised of eight rule templates). The IBM Firewall comes with a set of default services. You cannot delete these preloaded default services but you can modify certain fields. However, if these predefined services do not meet your needs you can add to services by using the rule templates to create new rules. See “Defining Services” on page 86 for more information.

Rule Templates

Provide instructions to the Firewall to permit or deny IP packets based upon their various attributes.

Socks Templates

Provide instructions to the firewall socks daemon to permit or deny IP packets based upon their various attributes.

Network Objects

Represent the various network components, like hosts, users, and subnets, that interact with the Firewall. They are defined by an IP address and an address mask, so it is possible for one object to represent a whole range of network addresses. Network objects can be grouped.

Network Object Groups

Represent one or more network objects. They are used as a convenience in setting up connections and can eliminate repetitive work. An example would be to group several addresses together into a network object group to represent a department. This network object group can then be used as either the source or destination for a connection.

Building Connections Using Predefined Services

In order to permit or deny specific types of communications between two named network objects or network object groups that serve as endpoints, you need to build a connection.

After you have defined your network objects, you create connections. Select one network object or group to be the source and another network object or group to be the destination for the traffic flow through the Firewall.

To build a connection, select Traffic Control from the configuration client navigation tree and double-click the file folder icon to expand the view. Select **Connection Setup**. The **Connections List** dialog box appears. Select **NEW** and click **Open**. The **Add a Connection** dialog box appears, as shown in Figure 19 on page 66.

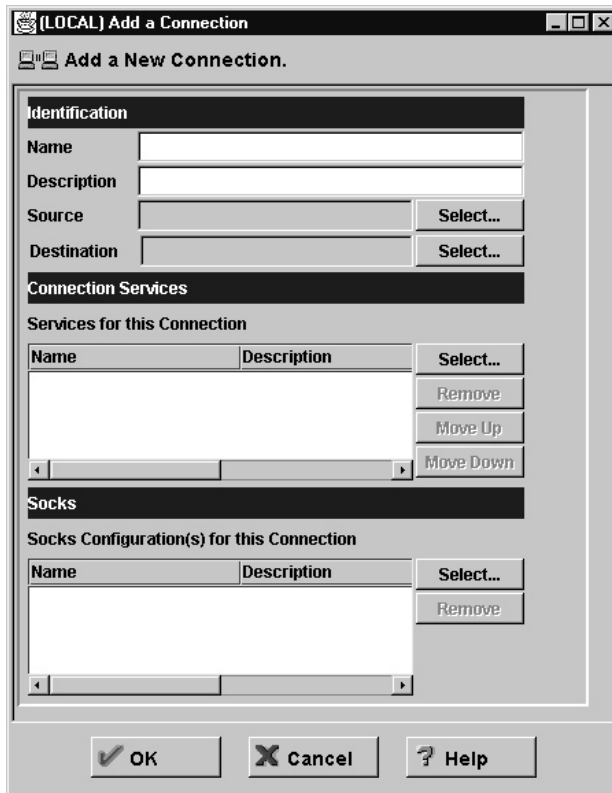


Figure 19. Add a Connection

1. Fill in a name for the connection.
2. Fill in a description of the connection.
3. For the source field, click **Select** and choose a network object from the **Network Object** dialog list.
4. For the destination field, click **Select** and choose a network object from the **Network Object** dialog list.
5. To choose the services for this connection, click **Select** and choose the type of traffic you wish to control between the endpoints.
6. Choose one or more services from the list to add the service to the Connection.
7. You can reorder the list by selecting a service and clicking **Move Up** or **Move Down**. See "Ordering Connections" on page 67.
8. You can remove a service by selecting it and clicking **Remove**.
9. If you use **Socks Configuration for this Connection**, follow steps 5–7 to make socks connections.
10. After you have everything defined, click **OK**.

11. When you save a new connection it remains excluded until you include it. To include one or more connections, go to the **Connections List** dialog box. Select one or more connections and click **Include** to initiate an explosion. This will make the connections active. This will also cause any related Socks entries to be added to the Socks configuration file. If you want to activate all connections at once, you can use the **Connection Control** dialog box. See “Connection Control”. You can not modify active connections.
12. To exclude one or more connections, go to the **Connections List** dialog box. Select one or more connections and click **Exclude**. This will deactivate the connections. This will also cause any related Socks entries to be deleted from the the Socks configuration file. If you want to deactivate all connections at once, you can use the **Connection Control** dialog box. See “Connection Control”.

Ordering Connections

Most IBM Firewall users have less than 1000 rules. The more rules you have, the greater the impact there will be on performance.

When a packet is received at a network interface, whether going into or out of the firewall host, rules are applied starting at the top of the generated connection rules. When the information from the packet exactly matches the information in a rule, the action (permit or deny) is taken. If the entire configuration is searched without a match, the request is denied.

Place more specific connections closer to the top and less specific connections closer to the bottom. For example, you might have a Department ABC, with an address of 1.1.10.X and a machine that is used as a server inside of Department ABC, with an address of 1.1.10.7. If you want to exclude machine 1.1.10.7 because it is a server that should not be used for telnet traffic, you must place the connection Deny telnet for Dept ABC server before the Permit telnet for Dept ABC connections. If you reverse the order of the connections, the deny connection will never be encountered.

Connection Control

Note: Before you activate connections, make sure your secure interface is defined.

Select **Connection Control** from the configuration client navigation tree to do any of the following:

Regenerate and Activate Connection Rules

The Firewall builds the static filter rules from the active ion connections rules.

Deactivate Connection Rules

The Firewall is now protected by the default rules.

List Current Connection Rules

You see the most recently generated connection rules set. If you previously deactivate rules, they are not being used.

Validate Rule Generation

The rules you have created are either valid or invalid.

Enable Logging

The Firewall logs selected traffic to the firewall log facility.

Disable Logging

Stops the Firewall logging.

The **Connection Control** dialog box appears, as shown in Figure 20.

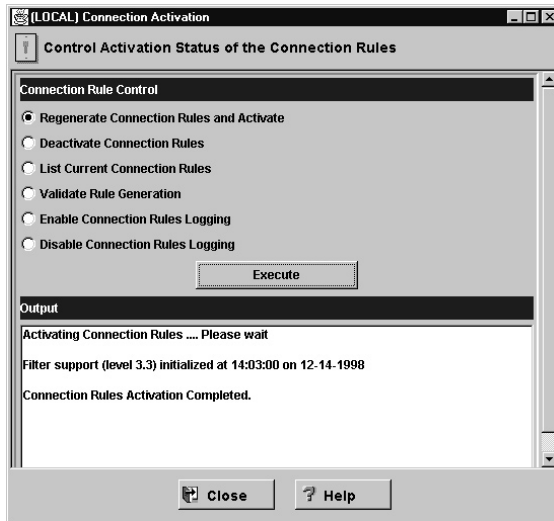


Figure 20. Connection Control

After you make a selection, click **Execute**.

Logging

Dynamic filter rules are automatically logged if you turn on firewall logging.

The following is a sample of the logging output when you regenerate and activate connection rules.

```
Feb 03 13:46:53 1998 mr16n18: ICA9037i: Firewall interfaces being updated
automatically on Tue Feb 3 13:46:53 1998.
```

```
Feb 03 13:46:55 1998 mr16n18: ICA1032i: Filter rules updated at
13:46:55 on Feb-03-1998
```

```
Feb 03 13:46:55 1998 mr16n18: ICA1037i: #:1 permit 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0
udp eq 53 eq 53 both local both l=n f=y t=0 e=none a=none
Feb 03 13:46:55 1998 mr16n18: ICA1037i: #:2 permit 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0
udp gt 1023 eq 53 both local both l=n f=y t=0 e=none a=none
Feb 03 13:46:55 1998 mr16n18: ICA1037i: #:3 permit 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0
udp eq 53 gt 1023 both local both l=n f=y t=0 e=none a=none
Feb 03 13:46:55 1998 mr16n18: ICA1037i: #:4 permit 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0
tcp gt 1023 eq 25 both local both l=y f=y t=0 e=none a=none
Feb 03 13:46:55 1998 mr16n18: ICA1037i: #:5 permit 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0
tcp/ack eq 25 gt 1023 both local both l=y f=y t=0 e=none a=none
Feb 03 13:46:55 1998 mr16n18: ICA1037i: #:6 permit 9.67.144.49 255.255.255.240
9.67.130.154 255.255.248.0 all any 0 any 0 both both both l=y f=y t=0 e=none a=none
Feb 03 13:46:55 1998 mr16n18: ICA1037i: #:7 permit 9.67.130.154 255.255.248.
0 9.67.144.49 255.255.255.240 all any 0 any 0 both both both l=y f=y t=0 e=none a=none
Feb 03 13:46:55 1998 mr16n18: ICA1037i: #:8 permit 9.67.144.49 255.255.255.240
9.67.131.250 255.255.255.255 tcp gt 1023 eq 21 secure local inbound
l=n f=y t=0 e=none a=none
Feb 03 13:46:55 1998 mr16n18: ICA1037i: #:9 permit 9.67.131.250 255.255.255.255
9.67.144.49 255.255.255.240 tcp/ack eq 21 gt 1023 secure local outbound
l=n f=y t=0 e=none a=none
Feb 03 13:46:55 1998 mr16n18: ICA1037i: #:10 permit 9.67.131.250 255.255.255.
255 9.67.144.49 255.255.255.240 tcp eq 20 gt 1023 secure local outbound
l=n f=y t=0 e=none a=none
Feb 03 13:46:55 1998 mr16n18: ICA1037i: #:11 permit 9.67.144.49 255.255.255.
240 9.67.131.250 255.255.255.255 tcp/ack gt 1023 eq 20 secure local inbound
l=n f=y t=0 e=none a=none
Feb 03 13:46:55 1998 mr16n18: ICA1037i: #:12 permit 9.67.144.49 255.255.255.
240 9.67.131.250 255.255.255.255 tcp gt 1023 gt 1023 secure local inbound
l=n f=y t=0 e=none a=none
.
.
.
Feb 03 13:46:58 1998 mr16n18: ICA1037i: #:100 deny 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0
all any 0 any 0 both both both l=y f=y t=0 e=none a=none
```

Determining the Rule States

The IBM Firewall rules can be in one of these states:

1. The configuration is not active.

You have not yet used the configuration client to activate the configuration or you have deactivated the configuration. This is the state of the configuration when you first install the IBM Firewall and boot your

system or deactivate filter rules. Default filters are in place to protect your network from intrusion when you first install the Firewall.

Firewall Access:

- The default filter configuration permits all local inbound traffic and permits all outbound traffic.
2. The configuration is active but has errors.

You have activated the configuration. Either there are errors (nonvalid rules) in the configuration or nothing has been configured. Errors and warnings are displayed in the Activation output window.

Firewall Access:

- Permit all local inbound traffic.
 - Permit all outbound traffic.
3. The configuration is active and valid. Note that there may have been some warnings, most notably, duplicate filter rules.

You have activated the configuration that you defined using the traffic control section of the configuration client.

Note: The configuration file can be valid and still contain no rules. In this case, an implied “deny all access” rule is in effect.

Firewall Access:

- Access determined by the configuration file.
Each packet that is received by, or is about to be sent by, any network interface is examined and its contents compared against each rule in the generated connection rules. When a match is found, the action (permit or deny access) on that rule is carried out.
- If no rules match the packet, there is an implied “deny all” rule that denies access.

Chapter 9. Examples of Services

This chapter describes how to configure the Firewall to perform certain common tasks. The tasks listed are examples only, but after understanding these, you should be able to configure your firewall to use any service that has been provided.

Planning Considerations

The Firewall's traffic control is organized in terms of connections that define the types of communication allowed or prohibited between pairs of endpoints. Therefore, it is critical to plan your connections in terms of these endpoints.

As described in "Chapter 8. Controlling Traffic through the Firewall" on page 63, endpoints are represented to the Firewall by network objects. If you have not already done so, you should complete the network planning worksheet in "Chapter 2. Planning" on page 9 and create the network objects necessary to represent your network.

The examples in this chapter use the following network objects:

Secure Interface

The secure interface of the Firewall.

Nonsecure Interface

The nonsecure interface of the Firewall.

Secure Network

The range of addresses that are accessible through the Firewall's secure interface. This could be a network object group that could contain several distinct domains, each of which is represented by its own network object.

The World

The nonsecure network.

Each desired type of communication must be viewed in terms of the endpoint-to-endpoint communication involved. In this stage, consider whether your firewall will be providing these communications by proxy or whether the Firewall will route these communications.

If the firewall acts as a proxy, then the firewall will perform the necessary work on behalf of the secure user and the nonsecure host(s) will never know

that the secure host exists. If the firewall is to route the traffic, then the secure host and the nonsecure host will speak directly to each other.

If you will use the Firewall as a proxy, then the endpoints of your communication will include the firewall, as shown in Figure 21.




Figure 21. Telnet Proxy and Direct Telnet Connection

Example of Telnet Proxy

This first example is of a straightforward outbound telnet proxy connection. In this example, users on the secure network will be allowed to use the firewall's Telnet Proxy to access telnet services on the hosts in the nonsecure network.

As described in Figure 21, two connections are taking place:

1. The client inside the secure network is connected to the firewall's Telnet Proxy.
2. The firewall's Telnet Proxy is, on behalf of the secure user, connected to the host in the nonsecure network.

To configure the Firewall's traffic control for this communication, we need to set up two connections:

Table 1. Telnet Proxy

Source Object	Destination Object	Services Required
Secure Network	Secure Interface	Telnet Proxy out 1/2

Table 1. Telnet Proxy (continued)

Source Object	Destination Object	Services Required
NonSecure Interface	The World	Telnet Proxy out 2/2

Example of Filtered Telnet

Contrast the above example with a simple filtered telnet connection. In this case, the client on the secure side will connect directly with the host on the nonsecure side.

Table 2. Filtered Telnet

Source Object	Destination Object	Services Required
Secure Network	The World	Telnet direct out

As noted before, this configuration will expose the addresses of your secure clients as they connect to nonsecure hosts.

Example of Proxy HTTP

Most installations will want to allow at least some of their secure clients to surf the Web. The IBM Firewall provides a predefined HTTP outbound direct service to allow routed HTTP, which functions exactly like the filtered Telnet example. In addition, the Firewall provides an HTTP proxy.

The HTTP protocol differs from Telnet in that it may encapsulate other protocols. Even for simple surfing, most users will require not only HTTP but also FTP services. To provide the full range of HTTP function, Gopher and WAIS should also be permitted, although these are used much less frequently.

Note, though, that when these additional protocols are used, they are wrapped in HTTP between the client and the proxy. Therefore the communication would be similar to the diagram in Figure 22 on page 74.

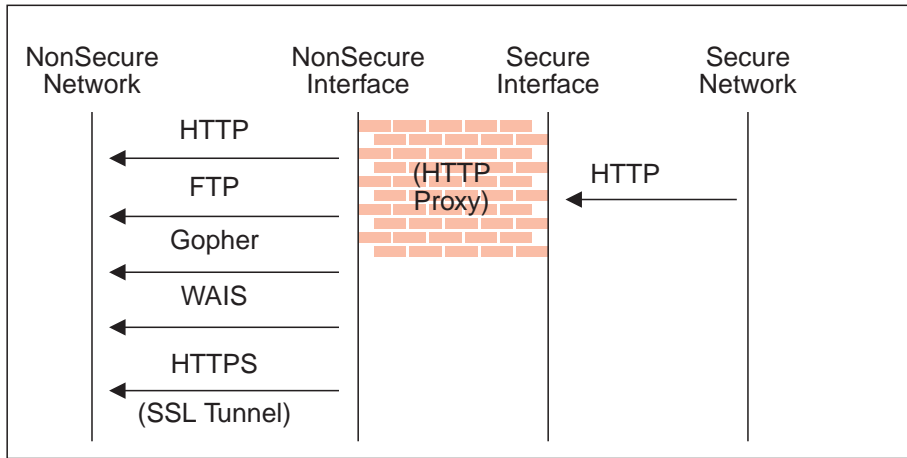


Figure 22. Proxy HTTP

Because we have two pairs of endpoints involved, we must code two connections.

Table 3. Proxy HTTP

Source Object	Destination Object	Services Required
Secure Network	Secure Interface	HTTP proxy outbound 1/2
NonSecure Interface	The World	Select from... <ul style="list-style-type: none"> • HTTP proxy out 2/2 • FTP proxy out 2/2 • Gopher proxy out 2/2 • WAIS proxy out 2/2 • HTTPS proxy out 2/2

For more information on HTTP Proxy, see “Chapter 12. Configuring Proxy Servers” on page 99.

Example of Socks

Socks presents a similar challenge to that of the HTTP proxy in that the socks daemon handles many different protocols and encapsulates them into a single data stream between the Firewall and the client. Socks is more flexible than the HTTP proxy because it can accommodate any TCP- or UDP-oriented protocol and because the Firewall can be configured independently of the filters to further control communications.

Because of this added flexibility, configuring socks requires a third connection in addition to those we demonstrated with the HTTP proxy. The two basic connections will allow the packets to flow to and from the Firewall; the third connection is required to tell the socks daemon to proxy the requests once it receives the packets.

Table 4. Socks

Source Object	Destination Object	Services Required
Secure Network	Secure Interface	Socks 1/2
NonSecure Interface	The World	Select from... <ul style="list-style-type: none"> • HTTP proxy out 2/2 • FTP proxy out 2/2 • Telnet proxy out 2/2 (Any second-half proxy service for which you wish to provide support)
Secure Network	The World	In the Socks Configuration window, select from... <ul style="list-style-type: none"> • permit socksified HTTP • permit socksified FTP • permit socksified Telnet

Of course, the clients inside your secure network must be socksified and must be configured to use your firewall as their socks server.

For more information on Socks, see “Chapter 11. Configuring the Socks Server” on page 89.

Example of Virtual Private Networks Using Static Filter Rules

To establish a Virtual Private Network requires an intricate configuration. In this case, the packets being sent between the client and the host are encapsulated for their journey between the two firewalls. For this reason, each packet passes through the filter mechanism twice: once in its encapsulated form and once in the clear. On each iteration, the packet looks completely different, and therefore requires a different connection to permit its passage.

The client, in the secure network, will be sending packets addressed to the remote host. These packets will be permitted by the Service *VPN Traffic 1/2*. Next, the same packet, still addressed to the remote host from the client in the secure network, will be encrypted and/or encapsulated based upon the tunnel being used by the service *VPN Traffic 2/2*. (It is recommended to copy this

service once for each tunnel being used. Each copy would reference a single tunnel ID, and any connections to that VPN would include the appropriate copy of this service). Once the packet has been encrypted and/or encapsulated, the Firewall sends the encapsulated packet to the remote firewall directly, where the packet will be decapsulated and/or decrypted and sent to its destination.

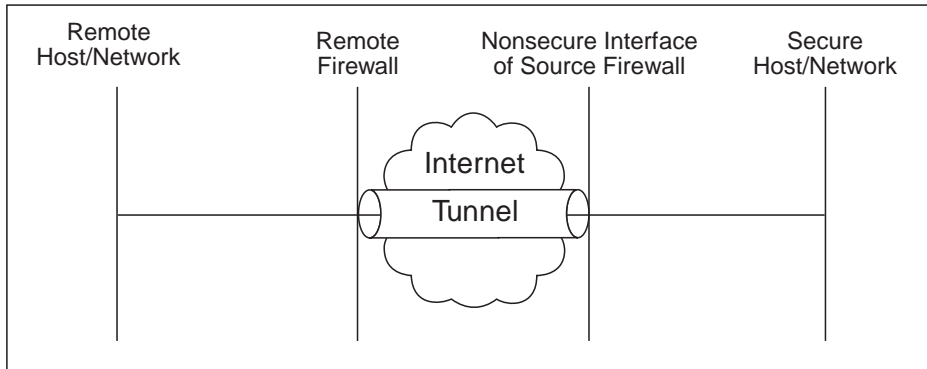


Figure 23. Virtual Private Networks

Such a configuration requires the following connections:

Table 5. Virtual Private Networks

Source Object	Destination Object	Services Required
Secure Host/Network	Remote Host/Network	<ul style="list-style-type: none"> • VPN traffic 1/2 • VPN traffic 2/2
NonSecure Interface of the Source Firewall	Remote Firewall	VPN encapsulation

For more information on VPNs, see “Manual Tunnels” on page 127.

Chapter 10. Customizing Traffic Control

This chapter helps you to define filter rules and services. Services are a collection of rules or a set of instructions to permit or deny a particular type of traffic through the Firewall, for example, a telnet session. You can add to services by using the rule templates to create new rules. You can also delete services. Socks services apply to socksified connections.

The IBM Firewall comes preloaded with a default set of services. You can tailor any predefined services to your particular needs or create new services.

Using the Configuration Client to Create Rule Templates

Use this procedure to add a new rule to the list of available rule templates.

1. From the configuration client navigation tree, select Traffic Control and double-click the file folder icon. Select **Connection Templates** and then select **Rules**.
2. On the **Rules List** dialog box, double-click **NEW**.

The IBM Firewall displays an **Add IP Rule** dialog box, as shown in Figure 24 on page 78 so that you can define a rule.

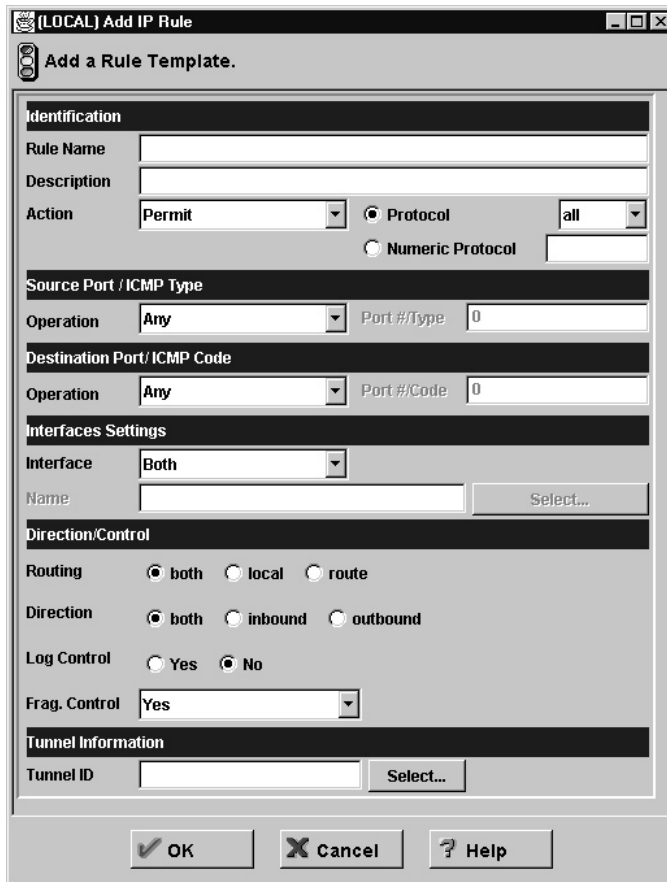


Figure 24. Add IP Rule

3. Enter the Rule Name.
4. Enter the Rule Description. This field is optional.
5. Click the action arrow and choose to either permit or deny access to the Firewall.
6. Click the protocol arrow and select from the following list:
 - all** Any protocol will match this rule.
 - tcp** The packet protocol must be Transmission Control Protocol (TCP) to match this rule.
 - tcp/ack** The packet protocol must be TCP with acknowledgement to match this rule.

- udp** The packet protocol must be User Datagram Protocol (UDP) to match this rule.
- icmp** The packet protocol must be Internet Control Message Protocol (ICMP) to match this rule.
- ospf** The packet protocol must be Open Shortest Path First protocol (OSPF) to match this rule. When **ospf** is specified as the protocol, the source port operation and source port value is used for the **ospf** record type value. Filtering can also be performed on the **ospf** type. A type value of **any** can be specified and the destination port fields must be specified as **any 0**. Anything else is ignored.
- ipip** The packet protocol must be IP-in-IP protocol (IPIP) to match this rule. When IPIP is specified, the port fields must be specified as **any 0**.
- esp** The packet protocol must be encapsulating security protocol used by the virtual private network for sending encrypted or authenticated IP packets, to match this rule.
- ah** Authentication header protocol is the packet protocol used by the virtual private network for sending IP packets which have an associated authentication token.

7. The numeric protocol allows you to specify a protocol by using its decimal value (according to RFC-1700). Valid values are in the range of 1 to 252. Note that port fields for this rule must be specified as 0 (signifying any port) when using this option. See RFC-1700 for a list of all protocols. Or, you can access the Internet Assigned Numbers Authority (IANA) directly with a browser.
8. The operation and port number operands are used together. The source and logical operations state a relationship between the port number (destination or origin) for the packet and the source port# and destination port# operands. For example, if the packet destination port is port 20, and the destination operation and destination port# are "ge 15", the packet matches (20 is greater than or equal to 15).

If you use a source or destination operation of **any**, the filter does not look at the port number; any port will match. The port number cannot be changed in this case.

For the ICMP protocol, rather than specifying a source port, specify an ICMP type and in place of a destination port, specify an ICMP code. The logical operator specified is applied to the type or code and, as for ports, an operator of **any**, means that any type or code value will match the rule. The port number cannot be changed in this case.

The values for operation are:

- Any

- Equal to
- Not equal to
- Less than
- Greater than
- Less than or equal to
- Greater than or equal to

Here are some of the more important ports to protect. The values for port numbers must be in the range 1 through 65535:

Port	Use
20	FTP data
21	FTP control
23	Telnet
25	Mail
53	Domain Name Server
70	Gopher
80	HTTP
111	RPC
161	SNMP
1080	socks

Here are some of the ICMP types and codes:

Type	Code and Description
0	0 - Ping reply
8	0 - Ping request
3	1 - Host unreachable
3	3 - Port unreachable
3	4 - Fragmentation needed but not allowed
5	1 - Redirect for host

9. Click the **Interface** arrow to select the type of interface (adapter).

both For packets coming or going on either the secure or the nonsecure interface

secure For packets coming or going on the secure interface

nonsecure
For packets coming or going on the nonsecure interface

specific

Use with the interface name field when selecting an interface, if you have assigned a name to the interface.

10. If you choose **specific**, for the interface type, the name of the specific interface will appear in the Name field.

11. Click the desired routing:

both Applies to all traffic.

local Implies that the packet is local to the firewall host. This means that:

- Incoming local packets are packets that are received by the interface and are destined for this firewall host; they will not be routed to another host. Their destination is local.
- Outgoing packets are transmitted from the interface, but originate on the firewall host. Their origin is local.

route Implies that the packet is routed by the firewall host. This means that:

- Incoming local packets are packets that are received by the interface and are destined for some other host; they will not remain on the Firewall. Their destination is remote.
- Outgoing packets are transmitted from the interface, and originated on some other host. Their origin is remote.

12. Click the desired direction:

both For packets going out from or into the selected interface

inbound

For packets coming into the selected interface from the network

outbound

For packets going out from the selected interface to the network

13. If you choose **Yes** for the Log Control field, every packet that matches that rule is recorded in the firewall log with priority level Error. If this parameter is not specified, the default is no.

14. Click the **Fragment Control** arrow to choose the desired fragment control. For IP packet information to match a rule fragmentation control specification, the control is interpreted as follows:

Yes The rule will match fragment headers, fragments and non-fragments. For fragments, the port information will be ignored and assumed to match.

Only Only fragments and fragment headers can match. For fragment headers, port information must match. For fragments, port information will be ignored.

No Only non-fragments can match. Fragment headers and fragments are excluded by this parameter.

Headers

Only non-fragments and fragment headers can match. Fragments are excluded by this parameter.

If this parameter is not specified, the default for both "permit" rules and "deny" rules is Yes.

Note: Regardless of the setting of this control, IP fragments with an offset of one (1) are discarded. This action eliminates a known attack of using packet fragments to overlay TCP header flags.

15. If this rule will be used with a tunnel, you can choose a tunnel ID. Click **Select** and choose a tunnel ID from the Select a Tunnel screen. Click **Apply**.

For a packet header to match a defined IP rule, the packet information must match all the parameters specified in the coded rule. For packet fragments, all parameters except port information is used to determine a match.

If the fragments were not permitted by an earlier rule, which had Yes or Only coded, the packet fragments will be denied by the final rule that is always appended to the bottom of the rule file.

Change IP Rule Configuration Entry

To modify an IP rule that you have created:

1. Double-click on an existing rule in the **Rules List**. The **Modify IP Rule Configuration** dialog box appears.
2. Modify the appropriate fields as described in "Chapter 10. Customizing Traffic Control" on page 77 and click **OK** to apply the changes.

Delete Rule Configuration Entry

To delete a rule, select a rule from the **Rules List** and click **Delete**.

Predefined Services

The IBM Firewall comes preloaded with a default set of services. Services are a collection of rules or a set of instructions to permit or deny a particular type of traffic through the Firewall, for example, a telnet session. You can add to services by using the rule templates to create new rules. The default services number 1 through 500.

The preloaded default services are:

All non-secure

Deny all traffic across nonsecure interface

All permit

Permit all traffic(for debugging purposes only)

All permit, in one direction

Permit all traffic(for debugging purposes only)

All secure

Deny all traffic across secure interface (in case of security violation)

All shutdown

Deny all packets (shutdown or debug)

Anti Spoofing

Deny inbound nonsecure packets with secure source address

Broadcasts

Deny broadcast messages to nonsecure interface

Config Client non-secure

Permit use of the configuration client from nonsecure network

Config Client secure

Permit use of the configuration client from secure network

CU-SeeMe

CU-SeeMe Video on default ports 7649 and 7648

DNS queries

(SECURITY POLICY) Permit DNS queries

DNS transfers

Permit DNS zone transfers (for secondary name servers)

Domain Controller Authentication

Allow use of Domain Controller for user authentication

FTP proxy in 1/2

Permit FTP inbound from nonsecure network to Firewall

- FTP proxy in 2/2**
 - Permit FTP inbound from Firewall to secure network
- FTP proxy out 1/2**
 - Permit FTP outbound from secure network to Firewall
- FTP proxy out 2/2**
 - Permit FTP outbound from Firewall to nonsecure network
- Gopher proxy in 2/2**
 - Permit gopher from Firewall to secure network
- Gopher proxy out 2/2**
 - Permit gopher from Firewall to nonsecure network
- HTTP deny non-secure**
 - Deny HTTP to nonsecure interfaces
- HTTP direct out**
 - Permit HTTP from secure network directly to nonsecure network
- HTTP proxy in 2/2**
 - Permit HTTP from Firewall to secure network
- HTTP proxy out 1/2**
 - Permit HTTP (port 8080) from secure network to the Firewall
- HTTP proxy out 2/2**
 - Permit HTTP from Firewall to nonsecure network
- HTTPS direct out**
 - Permit HTTPS (SSL) from secure network to nonsecure network
- HTTPS Proxy in 2/2**
 - Permit HTTPS (SSL tunnel) from Firewall to secure network
- HTTPS proxy out 2/2**
 - Permit HTTPS (SSL tunnel) from Firewall to nonsecure network
- Mail** (SECURITY POLICY) Permit Mail traffic through the Firewall
- NetBT Name Services broadcasts**
 - Allow NetBIOS over TCP/IP Name Services broadcasts
- Ping** Permit Ping outbound secure network to anywhere
- SDI authentication**
 - Permit connection to SecurID ACE server in the secure network
- Socks 1/2**
 - Permit use of Socks from secure network to the Firewall
- Socks deny non-secure**
 - Deny Socks from nonsecure adapters

Socks in 1/2

Permit use of Socks from nonsecure network to the firewall

Socks Monitor

Allow connection from Socks Monitor Clients

Socks Server Chaining

Allow Firewall Socks Server to chain to remote Socks server

Telnet direct out

Permit Telnet outbound from secure network to nonsecure network

Telnet proxy in 1/2

Permit Telnet inbound from nonsecure network to the Firewall

Telnet proxy in 2/2

Permit Telnet in from the Firewall to the secure network

Telnet proxy out 1/2

Permit Telnet out from secure network to Firewall

Telnet proxy out 2/2

Permit Telnet out from Firewall to nonsecure network

VDOLIVE Direct In

Permit nonsecure client to secure server

Note, users must configure individual player properties to use only UDP port 7001.

VDOLIVE Direct Out

Permit secure client to nonsecure server

VPN encapsulation

Permit encrypted data between firewalls

VPN traffic 1/2

Permit routed traffic on secure interface (non-encrypted)

VPN traffic 2/2

Permit routed traffic on a non-secure interface (encrypted)

WAIS proxy in 2/2

Permit WAIS (z39.50) from the Firewall to the secure network

WAIS proxy out 2/2

Permit WAIS (z39.50) from the Firewall to the nonsecure network

Defining Services

After you have defined a rule(s), you need to add the rule(s) to a service. Select Traffic Control from the configuration client navigation tree and double-click on Connection Templates, then select Services. The Services List dialog box appears. Double click NEW to get the Add Service dialog box, as shown in Figure 25.

The screenshot shows the "(LOCAL) Add Service" dialog box. It contains the following sections and controls:

- Identification:** Text boxes for "Service Name" and "Description".
- Service Composition:** A table for "Rule Objects" with columns "Flow", "Name", and "Description". To the right are buttons: "Select...", "Remove", "Move Up", "Move Down", and "Flow".
- Service Override Values:** "Override Log Control" and "Override Frag. Control" are dropdown menus set to "no override". "Override Tunnel ID" is a text box with a "Select..." button.
- Time Controls:** "Control By Time of Day" has "Begin" and "End" text boxes. "Control By Days" has a "Week Days" dropdown and "Begin" and "End" dropdown menus set to "Sun". "Time Control Action" has two radio buttons: "Activate Service During Specified Times" (selected) and "Deactivate Service During Specified Times".
- Buttons:** "OK", "Cancel", and "Help" at the bottom.

Figure 25. Add a Service

Using the Configuration Client to Create Services

1. Enter the service name.
2. Enter a description.

3. The **Override Log Control** field provides a means of overriding the log control setting in the rule templates that have been selected for this service. For example, if you include a set of rule templates that ordinarily have log control set to no, you can override this setting to be yes for the purposes of this service. The override setting will act on all of the rules in this service. In the **Override Log Control** field, enter one of the following choices:
 - no override - override is turned off, the settings in the rules themselves still apply
 - yes - write a log record when any rule in this service is matched
 - no - do not write a log record when any rule in this service is matched

When a log record is written for a filter rule, the values shown in the log record are the actual values from the IP packet. Logging matched filter rules can provide valuable information about the content of IP packets seen by the Firewall, for example, actual protocol and port numbers.

4. The **Override Frag. Control** field provides a means of overriding the Fragmentation Control setting in the rule templates that have been selected for this service. For example, if you include a set of rule templates that ordinarily have Frag, Control set to no, you can override this setting to be yes for the purposes of this service. The override setting will act on all of the rules in this service. In the **Override Frag. Control** field, enter one of the following:
 - no override - override is turned off, the settings in the rules themselves still apply
 - yes - match any IP packet, for example, non-fragments, fragment headers and fragments without headers
 - no - match only non-fragment packets, do not match the fragment headers or fragments without headers
 - only - match only fragment headers and fragments without a header, do not match non-fragments
 - headers - match only non-fragments and fragment headers, do not match fragments without headers
5. The **Override Tunnel ID** field provides a means of overriding the tunnel ID setting in the rule templates that have been selected for this service. For example, if you include a set of rule templates that ordinarily have no Tunnel Setting, you can override this setting to include a Tunnel ID for all of the rules in this service. If you leave the field blank, override is turned off. The settings in the rules themselves still apply. In the **Tunnel ID** field, select a tunnel by clicking Select.
6. The time controls allow you to associate a time range with each service. Therefore, this service will only be valid in a specified time period. If there is no time specification for a service, that service is valid all the time.

Control by Time of Day

Select if you want this service to be activated or deactivated according to begin and end times during the day. Use a 24-hour format. If this field is not enabled, the Time of Day fields will be in effect 24-hours a day.

Control by Days

Select if you want this service to be activated or deactivated according to a schedule based upon either days of the week or calendar dates. Note that whether a service is activated or deactivated depends on the value of the Time Control Action field.

Time Control Action

Choose **Activate Service During Specified Times** if you want this service to be activated during the specified times. This service will be deactivated during the times outside of those specified.

Choose **Deactivate Service During Specified Times** if you want this service to be deactivated during the specified times. This service will be activated during the times outside of those specified.

7. Click **Select** to choose the rules that comprise this service.
8. Use the Flow toggle to determine how the Source and Destination values of the Connection should be assigned to the filters as they get written to the Rule Base file.
 - > Left to Right indicates that the Source and Destination of the Connection gets written directly to the rule as it is written to the Rule Base File.
 - <--- Right to Left indicates that the Source and Destination of the Connection gets reversed when it is written to the Rule Base File.
9. When a packet is received, the IBM Firewall compares the information in the packet to the rules in the rules configuration file starting at the top of the file. It stops comparing when the first match is found and performs the action contained in the rule.

Once you have added a series of rules to the service, you can change their order. Select a rule from the **Service Objects** list and click the **Move Up** or **Move Down** buttons to reposition the rule. Or you can remove a rule by clicking **Remove**. The configuration client displays a refreshed list of rules. Click **OK** to save your changes.

Chapter 11. Configuring the Socks Server

Socks is an Internet standard for circuit-level gateways. You use the Socks server for address translation if your application uses TCP, such as Web browsers, FTP, or Telnet applications. Socks can help you access the Internet while hiding your internal IP addresses.

For outbound requests, from a secure client to a nonsecure server, the Socks server has the same objectives as a proxy server, that is to break the session at the Firewall and provide a secure door where users can be allowed to access the external, nonsecure network while protecting the addressing and structure of the internal network. The Socks server has the advantage of simplicity for the user, with little extra administrative work.

The Socks server can intercept all outbound TCP requests that would cross between your network and the Internet. The Socks server provides a remote application program interface so that the functions executed by client programs in secure domains are piped through secure servers at the firewall workstations, hiding the client's IP address. Access is controlled by filters that are associated with Socks rules.

The Socks server is similar to the proxy server. But while the proxy server actually performs the TCP/IP function at the Firewall, the Socks server just identifies the user and redirects the function through the Firewall. The actual TCP/IP function is performed at the client workstation, not at the firewall. This saves processing in the Firewall. The users in the secure network can use the many TCP/IP products that support the Socks standard. Figure 26 illustrates the Socks server relaying an HTTP request from a client in the secure network to the Internet, a nonsecure network.

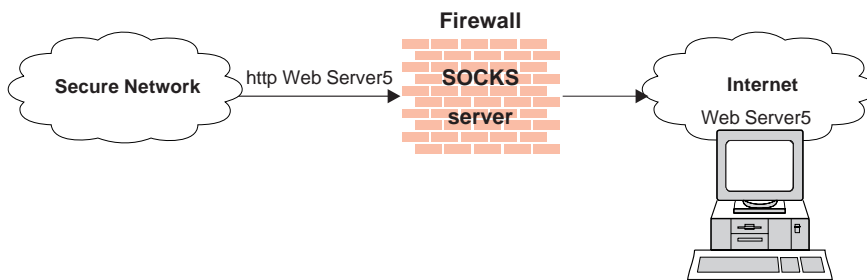


Figure 26. The Socks Server

The Socks server effectively hides the client's internal IP addresses from the outside world.

Authentication

The IBM Firewall provides the Socks Protocol Version 5, which enables clients inside the secure network to go through an authentication stage before communicating with applications in the nonsecure network. It also provides for authenticated generic proxy and the proxy of some streaming audio and video protocols.

A Socks Protocol Version 5 client must be authenticated using one of the server's supported authentication methods.

The IBM Firewall maintains an authentication database, that contains information about the authentication method assigned for every user wanting to use the Socks server.

The IBM Firewall supports the following authentication methods:

- User/Password
- Challenge Response Authentication Method (CRAM)

It can also support *no authentication* as a method, in which case a client can pass through without being authenticated.

When a Socks client initiates a request for service, the Socks server searches the database to determine the appropriate authentication method for that client. See "Authentication Methods Supported" on page 91 for further information.

The Socks daemon runs as a Windows NT Service, automatically starting when the system is started. In addition, a Watch Agent is provided to allow monitoring of the server. You can start the Watch Agent manually.

Three Authentication Profiles

The IBM Firewall provides a smooth migration path in the form of three authentication profiles so that customers can continue to use installed Socks Protocol Version 4 clients as they introduce Socks Protocol Version 5 clients.

1. The most permissive profile does not enable outbound authentication and permits any user, whether using a Socks Protocol Version 4 or Socks Protocol Version 5 client to connect. In this scenario inbound connections are denied.
2. The migration profile allows Socks Protocol Version 4 users to pass unauthenticated, but requires Socks Protocol Version 5 users to

authenticate. Inbound Socks Protocol Version 4 connections are denied and inbound Socks Protocol Version 5 connections are required to authenticate. This is the default profile.

3. The most secure profile requires that all users use Socks Protocol Version 5 clients and provide valid authentications.

To select among the three authentication profiles, modify the configuration file, `explode.cfg`. This file contains an entry `socks5profile=x`, where *x* indicates the authentication profile chosen. The values are as follows:

- permissive
- migration
- strict

Authentication Methods Supported

As already mentioned, the following authentication methods are supported:

1. Username/Password Authentication - A client sends two tokens in a single transmission. These tokens correspond to a username and a password. The Socks server might not send a custom challenge message. Many clients will cache their credentials. This method is suitable for the following firewall authentications:
 - Permit all
 - Deny all
 - NT password
 - Firewall password
 - SDI (with the restriction that the prompt will say *password* instead of *passcode* and that cached credentials will be useless)
 - User-defined authentications that do not require custom challenge messages
2. CRAM Authentication - A client sends a username, then the server sends challenges and receives responses until satisfied. Most clients do not cache responses, because there is no way to know what the challenge will be. This may be cumbersome for some protocols, including HTTP and FTP, because each new connection must authenticate independently. Most HTTP browsers open 8 to 10 connections to do their work, and FTP requires a new connection for each transfer. Therefore, CRAM users will see multiple authentication challenges. This method is suitable for all Firewall authentications.

Socks and Filters

When the Firewall is installed, the Socks server is enabled. However, only a default deny filter rule is provided. You must configure Socks using the configuration client so that Socks clients can use the Socks server.

Socks operates like a proxy in that there are two independent TCP connections established. The filter rules reflect this by requiring two connections. From the client to the Firewall, all traffic is TCP on port 1080. For secure-side clients, *Socks 1/2* is used. For nonsecure-side clients, *Socks In 1/2* is used. From the Firewall to the application server, all traffic takes place on the ports designated by the protocol involved. Use any *proxy 2/2* service appropriate for the protocol and for the secure/nonsecure side involved. For an example of how to set up a Socks service, see "Example of Socks" on page 74.

If you do not create Socks rules using the configuration client, you can build your own by editing the `s5.conf` file. See the *IBM SecureWay Firewall Reference* for more information.

Protocols Supported by Socks Protocol Version 5 Server

The Socks Protocol Version 5 server supports the following TCP and UDP protocols and many more:

- Archie
- Finger
- FTP
- Gopher
- HTTP
- HTTP Proxy
- News
- SNMP
- Telnet
- TFTP
- RealAudio
- RealPlayer
- Whois
- X-Windows

In addition, most e-mail clients are supported. Support for these protocols depends upon their actual implementation.

Configuring the Socks Server Using the Configuration Client

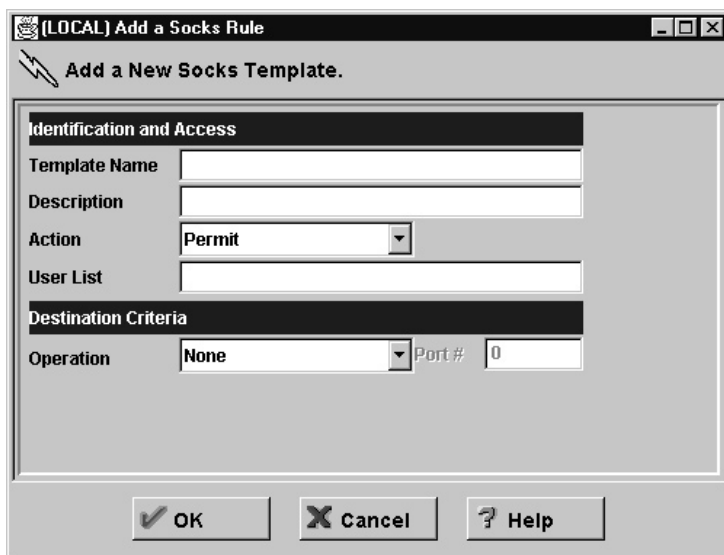
Socks templates are used to create rules that control security through the Socks server. The Socks templates allow you to customize, add to, copy, or delete existing Socks templates. These Socks rules, in turn, can be used in the definitions of connections on the Firewall in the same way filter rules are used.

Add a New Socks Rule

To add a rule to the Socks configuration file using a Socks template provided by the configuration client, select Traffic Control from the configuration client navigation tree. Double-click on the file folder icon to expand the view. Select Connection Templates. Double-click **Socks**. The **Socks** dialog box appears.

1. Double-click **NEW** to add a new Socks template.

The **Add a Socks Rule** dialog box appears, as shown in Figure 27.



Identification and Access	
Template Name	<input type="text"/>
Description	<input type="text"/>
Action	Permit
User List	<input type="text"/>

Destination Criteria	
Operation	None
Port #	0

Figure 27. Add a Socks Rule

2. In the **Template Name** field, enter the name of the Socks entry. This name must be unique and should not contain a pipe symbol(|), a single quote, or apostrophe, character (') or a double quote(") character as these are used as delimiters inside the code. Use of these characters will result in unreliable data.
3. Fill in a description.

4. Click the Action arrow and choose to either permit or deny access from a source to a destination.

When a datagram comes into the Socks server, the server compares the datagram specifications to each rule in the configuration file starting with the first rule until it finds a rule that matches exactly. Then it stops searching and performs the relevant action (either permit or deny access) on that rule. If no match is found, access is denied automatically.

5. In the **User List** field, you can enter a user ID or a list of user IDs to whom this rule applies. If you enter a list, separate the entries with commas. Do not use spaces, tabs, the pipe symbol (|), or double quotes(") in the user list.
 - The user list is limited to 396 characters.
 - User IDs must be IDs of users on the requesting host, not those on the destination host or Socks server host.
 - A user ID can consist of 1 to 8 characters, including:
 - a through z
 - A through Z
 - 0 through 9
 - _ (underscore)
 - A user ID should not contain a pipe symbol (|) or double quote character(").
6. In the **Operation** field, enter the logical operation to be performed on the port number:

eq	Equal to
neq	Not equal to
lt	Less than
gt	Greater than
le	Less than or equal to
ge	Greater than or equal to

When used with Port Number, the logical operation establishes a relationship that must be met. For example, if you enter the Operation **gt** and Port Number 23, then the port number must be greater than 23 for the rule to be invoked.

7. In the **Port #** field, enter the number of a port. The Port Number is used with the Operation to establish a relationship that must be met. For example, if you enter the Operation **gt** and Port Number 23, then the port number must be greater than 23 for the rule to be invoked. If operation and port number are omitted, the rule applies to all destination port numbers.

Use this **Add a Socks Rule** dialog box to permit or deny firewall access to network hosts based on the IP address.

Modify a Socks Rule

1. Double-click on an entry on the **Socks** dialog box.
The **Modify a Socks Rule** dialog box appears.
2. Change the appropriate fields as described in “Add a New Socks Rule” on page 93, and click **OK**.

Delete a Socks Rule

Select an entry from the **Socks** dialog box and click **Delete**. You are asked if you are sure you want to delete this Socks rule. Click **OK** to delete the rule.

Activate Connection Rules

As with the filter rules, you need to activate Socks rules. Click **Connection Activation** on the configuration client navigation tree, select **Regenerate Connection Rules and Activate**, then click **Execute**.

The Firewall copies the rules from the Socks configuration file to the firewall rules and activates the rules. When rules are activated, the new rules are recorded in the firewall log file.

Socks Logging

Socks messages are are written to the firewall log.

Client Considerations for Using the Socks Server

The majority of Web browsers are socksified and you can get socksified stacks for most platforms. Socksified clients for other TCP/IP applications are available from many sources. For a client that implements Socks, refer to that client documentation. For additional information refer to:

http://www.software.ibm.com/security/firewall/about/comp_products/
<http://www.socks.nec.com>

Tuning the Socks Server

The Socks server is configured initially to accept 64 concurrent user connections. If this threshold is exceeded, your Firewall log file will contain message ICA3007, which indicates that the maximum connection threshold has been exceeded. If you need to raise this limit, you may do so by adding the following line to `socks5.header.cfg`:

```
SET SOCKS5_MAXCHILD x
```

where x is the number of concurrent sessions you wish to allow.

However, each session takes approximately 200K of available memory. In very high-usage situations, your Firewall may exhaust its available memory and you may see the Windows dialog indicating that you are almost out of virtual storage. Either close some applications or increase the amount of virtual storage available. Once this dialog appears, various Firewall services could become unstable. Therefore, you should ensure that you do not set the number of concurrent sessions too high.

Note: 500 concurrent threads have been successfully run using 32 MB of real memory and a 100 MB page file.

Socks-Server Chaining

Socks-Server chaining is a feature by which one Socks server can reside behind another Socks server, yet still allow access to the network beyond the outermost Socks server. (It can be thought of as socksifying a Socks server). This is a very useful intranet scenario. Socks-server chaining can also be used with an Internet scenario.

To set up Socks-server chaining with the Socks server, edit the `socks5.header.cfg` file. This file resides in the Firewall's `config` subdirectory. (Alternately, you can edit `s5.conf` directly. See the *IBM SecureWay Firewall Reference* for more information). Add the following:

- A *no proxy* directive, to indicate the subnets to which your Firewall has direct access
- A *socks4* directive, to indicate the subnets that are accessible through a Socks Protocol Version 4 server
- A *socks5* directive, to indicate the subnets that are accessible through a Socks Protocol Version 5 server

For example, consider the following network, as shown in Figure 28 on page 97.

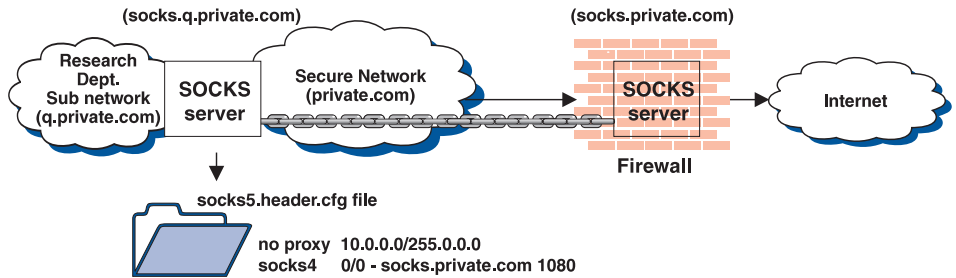


Figure 28. Socks Server Chaining

The Research department has a small private network, *q.private.com*, behind their own firewall. The Research department's subnet is 10.007.007.0/255.255.255.0. The company's private network, *private.com*, contains the entire 10.0.0.0/255.0.0.0 network. The company's Socks Protocol Version 4 server, *socks.private.com*, provides access to the Internet.

On Research's Socks server, *socks.q.private.com*, add the following two lines to `socks5.header.cfg`.

```
noproxy 10.0.0.0/255.0.0.0 - - -
socks4 0/0 - socks.private.com 1080
```

Lastly, add a Traffic Control connection to allow *socks.q.private.com* to communicate with *socks.private.com*. This might have already been done by a more general Service. Add a Connection whose source is the nonsecure interface of the *q.private.com* Firewall, whose destination is *socks.private.com*, and include the *Socks Proxy-Chaining* service. Then reactivate your Traffic Control rules.

Note: Either use the GUI to make modifications or edit the configuration file. But do not use both methods. By using both methods you risk losing your modifications.

Socks Traffic Monitor

Together with the IBM Firewall you receive a Socks Traffic Monitor application. This application allows you to monitor the current Socks traffic on the firewall in real-time. To use this application you must first start the IBM Firewall Socks Monitor Service on your firewall.

Once the service is started you can start the Traffic Monitor application either on the firewall itself or on a remote client. Do not forget to define a connection for this traffic if you want to use it from a remote machine.

For detailed information on the Socks Traffic Monitor, see Chapter 16 of *Guarding the Gates Using the IBM eNetwork Firewall V3.3 for Windows NT*. This book is shipped in PDF format on the IBM SecureWay Firewall product CDROM.

Chapter 12. Configuring Proxy Servers

This chapter contains information about how to configure and use the proxy servers from workstations both inside and outside your secure network.

Introducing HTTP Proxy

The IBM SecureWay Firewall provides a full-featured HTTP proxy implementation based upon technology developed for Web Traffic Express (WTE), the proxy caching component of WebSphere Performance Pack. The HTTP proxy efficiently handles browser requests through the IBM Firewall eliminating the need for a socks server for Web browsing. Users can access useful information on the Internet, without compromising the security of their internal networks and without altering their client environment to implement HTTP proxy.

The HTTP proxy is not a server. The end user cannot load files off of the proxy or put files on the proxy. Also, it is not a caching proxy. Nothing is stored on the firewall on behalf of an HTTP request.

Migration

If you have an existing HTTP proxy configuration file, it will be migrated.

Installation

The HTTP proxy is installed as an optional component of the Firewall installation. If you choose to uninstall this component, which is listed as "IBM Web Traffic Express", you must use the Windows NT "Add/Remove Programs" dialog. The base Firewall is listed as "IBM Firewall". If you use the Firewall's uninstall utility, the HTTP proxy will be uninstalled along with the rest of the product.

Methods Supported

The HTTP proxy supports the following methods, which are different ways of looking at the Internet:

- FTP
- Gopher
- HTTP
- HTTPS
- WAIS

Features

The IBM Firewall offers the following HTTP proxy features:

- Persistent connections
- User authentication
- SSL tunnelling
- Log maintenance
- Simple Network Management Protocol (SNMP)
- URL blocking

Persistent Connections

HTTP 1.0 is a simple request-response protocol. Each request is made over a new connection. After the data transfer is complete, the connection ends and a new connection is created to obtain another resource from the same server. Because of the way TCP works, there can be a delay in establishing each new connection. Therefore, it is helpful to reuse the TCP connections to make multiple HTTP requests over a single connection. You can do this by using persistent connections.

Persistent connections allow a client and a server to signal the close of a TCP connection. This signaling uses a connection header field.

User Authentication

HTTP uses a basic authentication method to authenticate the user to intermediate proxy servers.

SSL Tunneling

The SSL Tunneling Protocol (SSL-Tunneling) allows a Web proxy server to act as a tunnel for SSL enhanced protocols.

The client makes an HTTP request to the proxy and asks for an SSL tunnel. On the HTTP protocol level, the handshake to establish an SSL tunneling connection is simple. It looks like any HTTP request, except that a new "CONNECT" method is used and the parameter is not a full URL, but only a destination hostname and port number, separated by a colon. The port number is always required with "CONNECT" requests.

Log Maintenance

Your server can create two types of logs:

1. Error logs
2. Proxy Access log– used for logging proxy requests

You can tailor these logs to meet your needs.

Simple Network Management Protocol (SNMP)

IBM provides an SNMP management information base (MIB) and SNMP subagent so you can use any network management system, such as TME 10, NetView, TME 10 Distributed Monitoring, or HP OpenView, to monitor your server's health, throughput, and activity. The MIB data describes the Web server being managed, reflects current and recent server status, and provides server performance data.

The Web Traffic Express family provides its own built-in Simple Network Management Protocol (SNMP) subagent, which collects status and throughput information and makes it available to the SNMP Agent. The subagent depends upon having an SNMP agent available with which to communicate. The necessary agent implementation is available from the base AIX operating system. For the Windows NT operating system, it can be downloaded free from: <http://www.support.tivoli.com/sva>.

For more information about SNMP, refer to the *IBM WebSphere Performance Pack Web Traffic Express for Multiplatforms Webmaster's Guide: Version 2.0 for WebSphere Cache Manager*

URL Blocking

WTE provides basic URL-blocking features, which are useful for preventing access to certain undesirable web sites, for example. You can enable this feature by editing the `ibmpoxy.conf` file directly. This feature is not available from the GUI.

Browser Configuration

The client browser must be configured to connect to the port that the HTTP proxy is listening on.

If using HTTPS, point to the HTTP proxy on the IBM Firewall for security proxy also.

If you want to represent your Internet Explorer browser as an HTTP/1.1 browser to the proxy, do the following:

- Open the *View* pull-down.
- Select *Internet Options*.
- Select the *Advance Tab*.
- Scroll down to the HTTP 1.1 settings and set the switches to on.

Configuring Particular Features of HTTP Proxy

You can configure the HTTP proxy using the five-page notebook control that is part of the GUI, as shown in Figure 29.

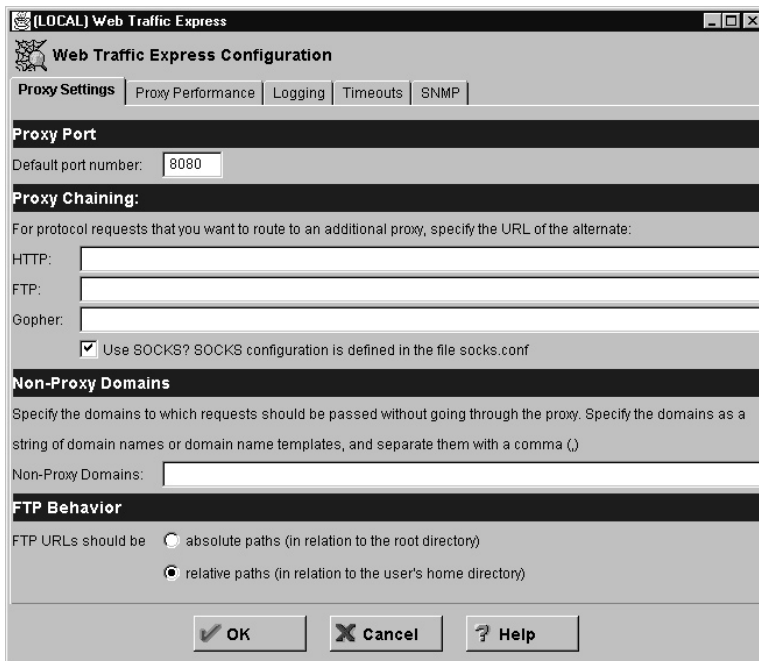


Figure 29. Proxy Notebook Control

Many configuration options are available, however you may choose not to modify anything and just use the default values given.

Certain function is available to the administrator by editing the configuration file, `ibmpoxy.conf` directly. The configuration file resides in `/etc` on AIX and in `%SYSTEMROOT%` on Windows NT.

Proxy Settings

On the **Proxy Settings** tab enter the following:

1. Enter the proxy port number the server should listen to for requests. The standard port number is 80. Other port numbers less than 1024 are reserved for other TCP/IP applications and should not be used. Common ports used for proxy Web servers are 8080 and 8008. The default port number is 8080.

2. If the proxy server is part of a chain of proxies, use proxy chaining to identify the name of another proxy that this server should contact for HTTP, FTP, or Gopher requests. You must identify a full URL including the trailing slash.
3. Check this box to instruct the proxy to use the Socks configuration file to determine the type of connection to make. The Socks configuration is defined in the `socks.conf` file.
4. If you are using proxy chaining you can identify the domains that you want the server to directly connect to rather than going through the proxy. Specify the domain as a string of domain names or domain name templates. Separate each entry in the string with a comma. Do not put any spaces in the string.
5. Check the box to specify whether the path information in FTP URLs should be interpreted as being relative to the logged-in user's working directory or as being relative to the root directory.

If absolute path is specified, then the file described in the URL is interpreted relative to the FTP server's root file system. Given URL: `ftp://user@server.com/etc/ibmproxy.conf`, the proxy will retrieve:
`/etc/ibmproxy.conf`

If relative path is specified, then the file described in the URL is interpreted relative to the user's home directory on the FTP server. Given URL: `ftp://user@server.com/etc/ibmproxy.conf`, the proxy will retrieve:
`/u/user/etc/ibmproxy.conf`

Proxy Performance

The IBM Firewall HTTP proxy supports persistent connections between a client and the proxy and between a proxy and the server. The *maximum persistent requests* condition and the *persistent connection timeout* condition control how long that connection will exist. Should one of these conditions arise, the socket connection between the proxy and the client will close. If the *maximum persistent requests* condition and the *persistent connection timeout* condition are not met, the connection will remain open and it is the client's responsibility to determine when a request is complete.

If determined incorrectly, this could result in a display indicating traffic on the connection when there is none. An example of this is the animated icon of a browser that continually runs even though the complete page has been loaded. Click **Stop** to halt the animation.

On the **Proxy Performance** tab enter the following:

1. Specify the maximum number of threads that you want to have active at one time. If the maximum is reached, the server holds new requests until another request finishes and threads become available. Generally, the more

power your machine has, the higher the value you should use. If your machine starts to spend too much time on overhead tasks, such as swapping memory, try reducing this value.

2. Enter the number of listen backlog client connections you want the server to carry before sending connection refused messages to clients. This number depends upon the number of requests that your server can process in a few seconds and should not be set higher than the number the server can process before the clients timeout and abort the connection from their end.

The range is 1–5.

3. Enter the buffer size for dynamic data generated by the server. Dynamic data is output from CGI programs, server-side includes, and API programs. The value can be specified in either kilobytes or megabytes.
4. Choose whether you want your server to look up the host name of requesting clients.

The value you use affects the following things about how your server works:

- The performance of your server. Using the default value of OFF improves the performance and response time of the server because it does not use resources to perform the host name lookup.
 - The information your server records about clients when writing to log files.
 - OFF - Clients identified by IP address
 - ON - Clients identified by host name
 - Whether you can use host names on address templates in protection setups, server group files, and ACL files.
 - OFF - Cannot use host names on address templates; must use IP addresses
 - ON - Can use host names on address templates; cannot use IP addresses
5. Choose whether you want to allow persistent connections. A persistent connection reduces latency for users and reduces the CPU load on the proxy server while requiring more resources. More threads, and therefore more memory on the proxy server are required for a persistent connection. Persistent connections must not be used on a multi-level proxy server setup if any of the proxies is not HTTP/1.1 compliant.
 6. Specify the maximum number of requests the server should receive on a persistent connection. When determining this number, be sure to consider the number of images used in your pages. Each image requires a separate request.

Logging

On the **Logging** tab enter the following:

1. Enter a directory path and file name where you want the server to log access statistics that pertain to proxy requests. By default, the server writes an entry to this log each time it acts as a proxy for a client request.
The server starts a new log file each day at midnight if it is running. Otherwise, the server starts a new log file the first time you start it on a given day. When creating the file, the server uses the file name you identify and appends a date suffix or extension. The date suffix or extension is in the format `Mmmddyyyy`, where `Mmm` is the first three letters of the month; `dd` is the day of the month; and `yyyy` is the year.
2. Enter a directory path and file name where you want the server to log internal errors.
The server starts a new log file each day at midnight if it is running. Otherwise, the server starts a new log file the first time you start it on a given day. When creating the file, the server uses the file name you identify and appends a date suffix. The date suffix is in the format `Mmmddyyyy`, where `Mmm` is the first three letters of the month; `dd` is the day of the month; and `yyyy` is the year.
3. Select the log archiving method from the dropdown box. The archiving options work on all logs with global settings and either it purges the logs, compresses them, or does nothing.
4. If you select compress to identify when to compress the log:
 - Specify when to compress the log. When the log is older than the number of days you entered, it will be compressed.
 - Specify when to delete the log. When the log is older than the number of days you entered, it will be deleted.
 - Enter a compress command. The command identifies the compression utility used to compact the logs and passes parameters to that utility.
5. If you select purge as the log archiving method then:
 - Specify when to delete the log. When the log is older than the number of days you entered, it will be deleted.
 - Specify how large the log file can grow (in megabytes) before it is purged. If you specify 0, there is no size limit for the log.

Timeouts

On the **Timeouts** tab enter the following:

1. Enter the time allowed for a client to send a request after making a connection to the server. A client first connects to the server and then sends a request. If the client does not send a request within the amount of

time specified, the server drops the connection. Specify the time value in any combination of hours, minutes, and seconds.

2. Enter the maximum time allowed for your server to send output to a client. The time limit applies to requests for local files and requests for which the server is acting as a proxy. The time limit does not apply for requests that start a local CGI program.

If the server does not send the complete response within the amount of time specified, the server drops the connection. Specify the time value in any combination of hours, minutes, and seconds.

3. Specify the time after which a long-running (but not idle) connection is terminated. Specify the time value in any combination of hours, minutes, and seconds.
4. Specify the amount of time the server should wait between client requests before cancelling a persistent connection.

The server uses a different timeout, the input timeout, to determine how long to wait for the client to send the first request after the connection is established.

After the server sends its first response, it uses the persistent timeout to determine how long it should wait for each subsequent request before cancelling the persistent connection. Specify the time value in any combination of hours, minutes, and seconds.

SNMP

On the **SNMP** tab do the following:

1. Check whether you want to enable SNMP. SNMP is a network management system program that runs continuously and is used to monitor, reflect the status, and control a network.
2. Use the SNMP community name to define the password between the webserver DPI subagent and the SNMP agent. The SNMP community name authorizes a user to view the status of the network. Public is the default community name.
3. Use the webmaster's email address for the Web server administrator to receive problem reports. The default webmaster's email address is webmaster.

The network management system uses SNMP GET commands to look at MIB values on other machines. It then can notify you if specified threshold values are exceeded. You can affect server performance by modifying configuration data for a server, to proactively tune or fix server problems before they become server outages.

User Authentication

HTTP uses a basic authentication method to authenticate the user to intermediate proxy servers.

The basic authentication method uses a request header, a response status code, and a response header. The authentication challenge returns from the proxy with the 407 status and the Proxy-Authenticate: header. The client will reissue the request, attaching the printable-encoded username and password in the Proxy-Authorization: header.

On the HTTP proxy configuration panel, set the proxy to **authenticate all**. Use the **User Administration** dialog box to set the users for **secure http** to something other than **deny**. For more information on authentication, see “User Authentication Methods” on page 120.

SSL Tunneling

SSL tunneling for HTTP Secure Connection to other servers is supported. The client makes an HTTP request to the proxy and asks for an SSL tunnel. On the HTTP protocol level, the handshake to establish an SSL tunneling connection is simple. It looks like any HTTP request, except that a new “CONNECT” method is used and the parameter is not a full URL, but only a destination hostname and port number, separated by a colon. The port number is always required with “CONNECT” requests.

The IBM Firewall acts as a gateway. The tunnel goes from the client through the firewall to the server. Use the standard port 443 for HTTP Secure Connection: <https://www.ibm.com:443>. Also, use the predefined service HTTPS proxy out 2/2.

If using HTTPS, point to the HTTP proxy on the IBM Firewall for security proxy also.

For more information, see “Example of Proxy HTTP” on page 73.

Log Maintenance

On the Log Files panel you have a choice of two log facilities:

- Proxy Access log
- Error log

For each logging facility you must provide a fully-qualified filename root. This filename root will be appended-to with a suffix indicating the date. Each day at midnight (local time), the preceding day's log files are closed and the following day's log files are created.

Compression and Purging

With the log maintenance options, you can specify how to handle the accumulation of daily logs for days past. You can choose whether to:

- Compress and then remove logs
- Remove logs after they reach a certain age or a collective size
- Run your own program at midnight each night to handle old logs

Note that the collective size is the total size for each log type.

To compress the logs, set Log Archiving method to compress. The IBM Firewall compresses the logs that are older than the value (in days) defined by the CompressAge directive into a common zip file. If the value for CompressAge is less than or equal to 0, then no files are compressed. Today's logs are not compressed. The compressed file is stored in the path defined by the Compress command directive.

After compressing the logs, you can delete logs older than the value defined by the CompressDeleteAge to recover space used by the logs. You can have the logs removed automatically, based on the age of the log using the PurgeAge directive, the collective size of the logs (using the PurgeSize directive), or both.

Log files are not deleted if the value entered for CompressDeleteAge is less than the value of CompressAge.

Integration with Firewall Logging

A logging plugin is provided, which generates the Firewall log message ICA2099, identical to the current Firewall's log message.

Reducing the Volume of Log Information

In addition to the options presented on the panel, you can choose to exclude certain types of information from the Proxy Access log. By editing the configuration file, `ibmproxy.conf`, the Proxy Access log can filter out log records based upon the following criteria:

- URL
- Browser type
- Method

- MIME type
- Return code

Proxy Log Reporting

The logs produced by the proxy adhere to the Common Log Format, which is widely recognized in the web server and web proxy industry. To analyze the log files, you can use any report analysis tool that accepts standardlog files.

For example:

- WebTrends for Firewalls and VPNs manages, monitors, and reports on firewall activity and proxy servers to help users understand network traffic and security issues. It analyzes leading firewall and proxy server log files, and generates customizable reports showing security violations, bandwidth usage, email usage, employee productivity, and more. Filters can drill down on specific information by users, departments, actions, or protocol.

For more information, see <http://www.webtrends.com/>.

- ProxyReport, a proxy server log file analyzer, installs on the proxy server. Users can run and view reports from any web browser on any computer. Use ProxyReport to:

- Monitor and analyze employee access to the Internet
- Monitor and analyze the performance of your proxy server
- Produce management reports

For more information, see <http://www.netrics.com/proxyreport>.

- ProxyReporter for IBM Web Traffic Express helps departmental managers and network administrators track and analyze employee Internet activity. ProxyReporter presents how Internet resources are being used and helps any organization police its Internet Use Policy.

For more information, see

http://www.wavecrestcomputing.com/searchEngine/_prIBM.html.

Firewall Report Utilities

The http proxy activity is recorded in the firewall log. You can use the report utility functions to assist you in generating usage reports for the http proxy.

Using the report utilities and the firewall log, you can create a Session table, which contains proxy session start/stop information from the session.tbl file. See “Report Utilities” on page 193 for more information. See the *Report Utilities* chapter of the *IBM SecureWay Firewall Reference* for information about the Session table.

URL Blocking

To implement URL blocking, add one or more *Fail* directives to the Mapping Rules section of `ibmpoxy.conf`. Mapping rules are processed in order. The first rule encountered is the only one which is applied. Therefore, *Fail* directives should appear before any *Pass*, *Exec*, *Map*, *Redirect* or *Proxy* directives.

A *Fail* directive takes a single parameter—a request template. This template is a simple wildcard-matching template (not a full-blown regular expression). This template is applied against the URL provided in the request— if the requested URL matches the pattern defined by the template, the directive is applied. For example:

- Fail `https://*` will fail all HTTPS requests.
- Fail `ftp://*` will fail all FTP requests.
- Fail `*.class` will fail all requests for objects whose name ends with `".class"`. However, this will not block all java applications, because it is common for class files to be contained within a `".jar"` file. URL blocking only refers to the object's name, not to its actual content.
- Fail `*dirtysite.com*` will fail all requests to `"dirtysite.com"`. Note that this will also fail requests such as `http://info.bulletinboard.com/announcements/dirtysite.com.html` or `http://search.engine.com/search?dirtysite.com`.

FTP

1. Use the FTP proxy to access the firewall host. (We will use `ftp_gw.domain.net.com` as the host name for the firewall).

```
ftp ftp_gw.domain.net.com
```

The proxy server will ask for your user name:

```
login:
```

2. Enter your user name as authorized to use the Firewall:

```
login: jane_doe
```

The server validates your identity depending on the authentication scheme selected when your user name was added to the Firewall (see "Adding a User to the IBM Firewall" on page 115).

After you are authenticated, the proxy server displays an FTP command prompt.

```
ftp>
```

Use the `quote` and `site` FTP commands to connect to the foreign host:

```
ftp> quote site forhost.network.outside.com
```

The foreign host will now ask for a user name and password for you to connect. This is probably a different user name and password from those you used to FTP to the Firewall.

The default timeout value for login is 60 seconds and for idle proxy is 7200 seconds. To change the default timeout values see “Overriding Timeout Values in FTP and Telnet Proxies” on page 113.

Transparent FTP

You can ftp transparently through the Firewall. Transparent proxies require no firewall authentication, therefore users of transparent proxies do not have to be defined as firewall proxy users. Transparent proxies are only allowed from the secure side of the firewall going out to the nonsecure side of the firewall. In order for transparent proxy to work, you have to select it on the Security Policy configuration client panel.

1. Use ftp to access the firewall host. (We will use ftp_gw.domain.net.com as the host name for the firewall.)

```
ftp ftp_gw.domain.net.com
```

2. The proxy server will ask for your user name:

```
USER:
```

3. Enter your user name at the nonsecure network:

```
USER: username@remote_site_host_name
```

4. You are then prompted by the target host for your password of the user name entered in the previous step.

```
password:
```

5. Enter your password.

The default timeout value for login is 60 seconds and for idle proxy is 7200 seconds (two hours). To change the default timeout values see “Overriding Timeout Values in FTP and Telnet Proxies” on page 113.

Telnet

Use the Telnet proxy to login to the firewall proxy server. You can use either the host name or Internet address. Then, after your credentials are authenticated, you use the Telnet command at the Firewall to log in to the intended host. For example, let’s use Telnet from inside the secure network,

through the Firewall with the host name of `telnet_gw`, to access your ultimate destination, `forhost.network.outside.com`.

1. To start the process, use Telnet to access the firewall host. (We will use `telnet_gw.domain.net.com` as the host name for the Firewall.)

```
telnet telnet_gw.domain.net.com
```

2. The proxy server will ask for your user name:

```
login:
```

3. Enter your user name as authorized to use the Firewall:

```
login: jane_doe
```

The server validates your identity depending on the authentication scheme selected when your user name was added to the Firewall, see “Adding a User to the IBM Firewall” on page 115 for more information.

You will be using the `oneact` shell. With the IBM Firewall proxy Telnet daemon, all communications pass through the firewall.

If you are using the `oneact` shell, after you are authenticated, the proxy server displays:

```
ENTER DESIRED HOST:
```

Type

```
telnet forhost.network.outside.com
```

The foreign host asks for your user name and password, as you are known on that host. These might be different from the user name and password that you used on the firewall proxy server.

The default timeout value for login is 60 seconds and for idle proxy is 7200 seconds. To change the default timeout values see “Overriding Timeout Values in FTP and Telnet Proxies” on page 113.

Transparent Telnet

You can telnet transparently through the Firewall. Transparent proxies require no firewall authentication, therefore users of transparent proxies do not have to be defined as firewall proxy users. Transparent proxies are only allowed from the secure side of the Firewall going out to the nonsecure side of the Firewall. In order for transparent proxy to work, you have to select it on the Security Policy configuration client panel.

1. Use Telnet to access the firewall host. (We will use `ftp_gw.domain.net.com` as our host name.)

```
telnet telnet_gw.domain.net.com
```

2. The proxy server will ask for your user name:
Login:
3. Enter your user name at the nonsecure network:
Login@remote_host

The foreign host asks for your user name and password, as you are known on that host. These might be different from the user name and password that you used on the firewall proxy server.

The default timeout value for login is 60 seconds and for idle proxy is 7200 seconds. To change the default timeout values see “Overriding Timeout Values in FTP and Telnet Proxies”.

Overriding Timeout Values in FTP and Telnet Proxies

Both FTP and Telnet have timeout values for logging in and idle waits. By default, there must be session activity at least once every 60 seconds during login and user authentication. This is known as the loginTimeout.

Once the login has completed successfully, there must be activity on the session at least once every 7200 seconds or the session is disconnected.

You can override these defaults by creating an fwTimeout.cfg file in the ROOTDIR\config directory by specifying new timeout values in seconds. The fwTimeout.cfg file should have the following format.

```
telnet
proxyTimeout = 7200
loginTimeout = 60
```

```
ftp
proxyTimeout = 7200
loginTimeout = 60
```

Chapter 13. Administering Users at the Firewall

This chapter describes how to do the daily administrative tasks with the IBM Firewall, including:

- Adding users to the IBM Firewall so that they can access hosts outside your protected network
- Changing the attributes of the users who access the firewall
- Deleting users who no longer need access outside your network

Do not edit the configuration files directly; if you do, your IBM Firewall user attributes will not be set up correctly. Do all IBM Firewall administration using the configuration client dialogs or command line.

Adding a User to the IBM Firewall

The IBM Firewall defines three types of users and stores information about them in two different user databases.

Types of Users

The IBM Firewall divides users into three categories:

Proxy Users

Use firewall services, such as the HTTP proxy service to access Web sites on the Internet from within a corporate network. Proxy users are able to use services through the firewall but do not have access to the firewall machine and cannot perform local logins to the Firewall machine.

Firewall Administrators

Can use the Firewall proxy services, but they can also configure the Firewall by using the configuration client and by logging on to the Firewall from a remote host. Like proxy users, firewall administrators cannot perform local logins to the Windows NT operating system after the Firewall has been installed.

Firewall administrators can create and modify definitions for proxy users but they cannot create or modify the definitions of other firewall administrators.

Primary Firewall Administrators

Have the same capabilities as firewall administrators. They can also

perform local logins to the Firewall machine. Primary firewall administrators can create and modify definitions for other firewall administrators.

Types of Databases

There are two types of user databases.

Firewall User database

Contains firewall-related attributes for each proxy user and administrator. Includes attributes such as the user's firewall password and password rules and which authentication methods should be used to authenticate the user for each service.

If the proxy user is not defined in the firewall user database and the user tries to use firewall proxy services, the default user record, `fwdfusr` will be used to define the attributes and authentication schemes used to validate the user.

Primary firewall administrators cannot be defined in the firewall user database. Use the default firewall administrator record, `fwdfadm`, to assign attributes to administrators.

Like proxy users, if firewall administrators are also defined in the Windows NT user database, their Windows NT logon password will be used when the user requests services that must be authenticated using Windows NT logon passwords.

Windows NT User database

Contains the Windows NT logon passwords for users. In general, proxy users do not have to be defined in the Windows NT user database unless they are going to be authenticated using their Windows NT logon password.

If other authentication methods are going to be used to authenticate proxy users, they do not have to be defined in the Windows NT User database.

Primary firewall administrators are synonymous with Windows NT users that are members of the Windows NT administrators group and must be defined in the Windows NT user database.

Using the Configuration Client to Add a User

Adding a user to the IBM Firewall gives them access to the external network.

1. From the configuration client navigation tree, select **Users**. The **User Administration** dialog box appears.
2. Select **New** from the **User Administration** dialog box and click **Open**. The **Add User** dialog box appears, as shown in Figure 30 on page 117.

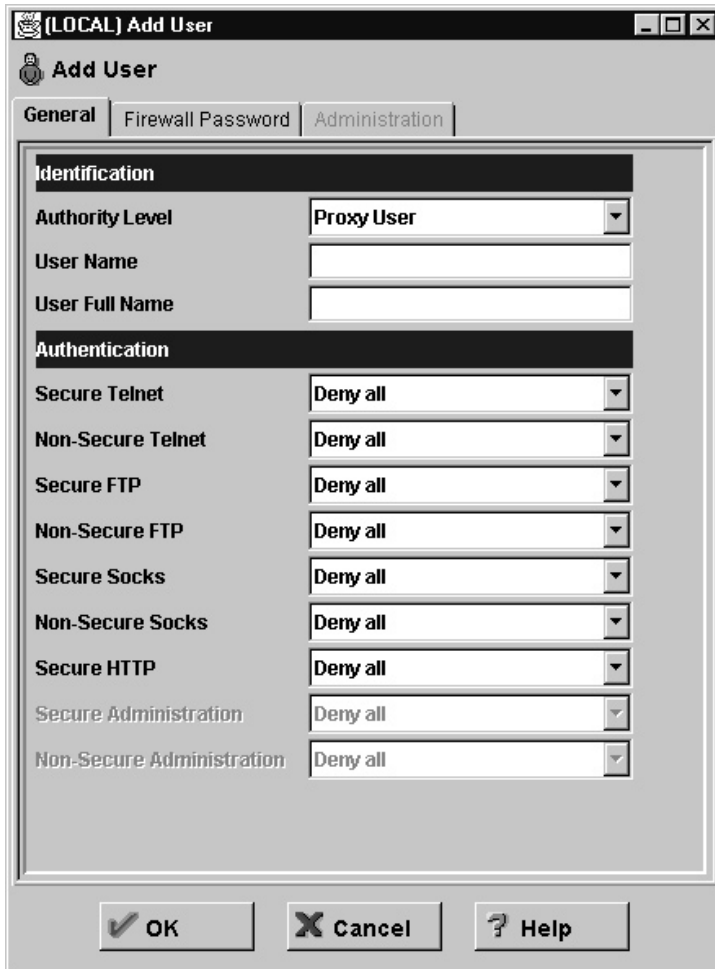


Figure 30. Add User

3. Provide this information:

Authority Level

Specifies the authority level for this user. Click the **Authority Level** arrow to select user type.

Socks/Proxy User

The user being defined is for both Socks server access and proxy access. The user has no administration authority. This level is the default.

Firewall Administrator

Has all of the attributes of a user, but an administrator can also log in to the Firewall and perform administrative

tasks. An administrator has additional attributes that define what administrative functions he or she is allowed to perform. A firewall administrator can create firewall users but cannot create other firewall administrators. Firewall administrators cannot log in locally to the Firewall machine. They must access the configuration server from a remote machine.

Primary Firewall Administrator

The primary firewall administrator is allowed to log in locally to the Firewall machine. The primary firewall administrator has full access to all administrative functions and also can create other firewall administrators except for primary firewall administrators.

The primary firewall administrator is defined by creating a user in the NT database and making that user a member of the NT administrators group. Modify the `fwdfadm` record to define the attributes for the primary firewall administrator.

User Name

Specifies the name for this user. This is the user name with which this user will log into the telnet or FTP server on the IBM Firewall. This is not necessarily the user's TCP/IP user name or host name, but they can be the same.

A user name can consist of from 1 to 20 characters, including:

- a through z
- A through Z
- 0 through 9
- _ (the underscore)

User names are not case sensitive.

The Firewall comes with two preinstalled users:

- a. Default user or `fwdfuser`. If a user is not defined in the firewall database, `fwdfuser` is used to determine the user's firewall attributes, such as which authentication methods to use when authenticating the user.

At installation, when the `fwdfuser` is created, all authentication methods are set to deny all. The permission for `fwdfuser` controls how the firewall processes undefined usernames.

The administrator can view `fwdfuser` or change the assigned authentication method using the configuration client or the command line. However, `fwdfuser` cannot be deleted and must

always exist at the firewall. In addition, firewall password is not a valid authentication type for `fwdfuser`.

- b. Default Primary Firewall Administrator, `fwdfadm`, defines the firewall attributes for all primary firewall administrators. Because primary firewall administrators do not have user records of their own in the firewall database, this record is used to define the authentication methods used to authenticate primary firewall administrators.

At installation, all the authentication methods for `fwdfadm` are set to *deny all* except for the secure and nonsecure administration authentication methods, which are set to NT logon password. The primary firewall administrators can view and modify this record, but it cannot be deleted. In addition, firewall password is not a valid authentication type for `fwdfadm`.

User Full Name

Specifies a description of the user.

The following fields refer to authentication methods. Click the arrows to select from the list of authentication methods. They are explained in “User Authentication Methods” on page 120.

Secure Telnet

Indicates whether this user’s identity, when logging in from the secure network, must be authenticated by some means.

Nonsecure Telnet

Indicates whether this user’s identity, when logging in from the nonsecure network, must be authenticated by some means.

Secure FTP

Specifies the level of authentication this user needs to use FTP to access the Firewall from the secure network.

Nonsecure FTP

Specifies the level of authentication this user needs to use FTP to access the Firewall from the nonsecure network.

Secure Socks

Specifies the Socks V5 authentication method for Socks client connections coming from the secure side of the firewall. Click the arrow to select from a list of choices. They are explained in “User Authentication Methods” on page 120.

Non-Secure Socks

Specifies the Socks V5 authentication method for Socks client connections coming from the nonsecure side of the firewall. Click

the arrow to select from a list of choices. They are explained in “User Authentication Methods” on page 120.

Note: Even if *permit all* is selected, if the socks server is set to authenticate users, it will still require a user name and will present a prompt to the user. The user need not provide a password at the prompt. To suppress the prompt, adjust your authentication profile as described in “Chapter 11. Configuring the Socks Server” on page 89.

Secure HTTP

Specifies a user ID/password type of authentication on outbound HTTP proxy requests. Click the arrow to select from a list of choices. They are explained in “User Authentication Methods” on page 120.

The browser prompts for user ID and password so if you are using SDI, fill in a passcode at the password prompt.

User-supplied must recognize that Socks/password cannot support interactive dialogs and behave accordingly.

Secure Administration

Specifies the authentication method used to log on from the configuration client through a secure interface. Note that when you log on locally (by choosing local on the logon panel) you are always in a secure environment, so this is the authentication method you would use.

Nonsecure Administration

Specifies the authentication method used to log on from the configuration client through a nonsecure interface.

User Authentication Methods

The choices for user authentication are:

Deny All

The user is denied access.

Permit All

No authentication is required. The server does not try to authenticate the user; but it proceeds with a command prompt so that the user can access a foreign host.

NT Logon Password

NT logon password is less secure than the firewall password. However, if users are already defined in a Windows NT domain, they can use the Windows NT logon password so they do not need multiple passwords.

If the user chooses this method of authentication, their user ID and password will be validated against the local Windows NT user database. If the Firewall is configured to trust other Windows NT servers, these trusted servers will be searched for user definitions.

Before trust relationships can be set up between the Windows NT Firewall and trusted Windows NT servers, a connection must be set up to allow TCP/IP communication traffic between the two machines.

Set up this connection using the following predefined services:

1. Domain Controller Authentication - which allows the use of the Domain Controller for user authentication
2. NetBT Name Services broadcasts - which allows NetBIOS over TCP/IP Name Services broadcasts

Use the Windows NT configuration utilities to define the trust relationships.

SecurID Card

Authentication is done using a Security Dynamics SecurID card or pinpad card. The PIN must be set before using this authentication method with the IBM Firewall. For FTP, the SDI new PIN mode and next token mode are not supported.

To use the Security Dynamics ACE/Server to authenticate users, you must install the ACE/Server server code on a non-firewall host inside the secure network. On the ACE/Server do the following:

1. Create a client (your firewall). Select Add Client. Enter the Name (IP address will be filled in automatically based on DNS). Set Client Type to Single-Transaction Comm Server. Set Encryption Type to DES. Click OK.
2. Create a Group. Select Add Group. Enter a Name. Click on Client Activations and select the firewall machine. Click on Add Client. Exit.
3. Create Users. Select Add User. Enter the Name for the Firewall User and their Default Login. Select Local User. Click on Assign Token and Yes. Refer to the ACE/Server documentation for more information on installing tokens. Select a token for the user. Click on Group Membership and Join Group.
4. Edit the token. Highlight the token on the User panel, and click on Edit Assigned Token. Resynchronize Token and then Set PIN to Next Tokencode.
5. Edit the ACE/Server Configuration Management. Make sure the Encryption Type is DES and set any other login options such as # of retries.

Note: It is helpful to install the ACE/Agent client on the ACE/Server machine to test authentication. Be sure to set up another client for this machine and activate the user's group.

You must also copy the `sdconf.rec` file from the ACE/Server to the firewall in the `\\Winnt\system32` directory and then reboot the machine.

Next, configure firewall users for SDI authentication. Click the **User** panel of the configuration client and set the authentication to SecurID Card.

Also, on the firewall, you must create a connection from the secure adapter to the ACE/Server host using the predefined service, SDI Authentication.

The proxy server asks for your PASSCODE (which will not be displayed) before letting you proceed.

Enter PASSCODE:

At this point, enter your 4-digit SecurID PIN code followed by a comma, (the comma is optional) and then the code from your SecurID card. For example, to log in as user NEWUSER with an assigned PIN of 1234, when your SecurID card shows the code 179091, you would enter:

```
login: NEWUSER
Enter PASSCODE: 1234,179091
```

If the SecurID card is in new PIN mode, you have to set the PIN before using this authentication method with the IBM Firewall.

User-Supplied Authentication 1, 2, and 3

Authentication is supplied by the user. You can install up to three user-supplied authentication methods on the Firewall. For information on how to create and compile a subroutine for user-supplied authentication, refer to the *IBM SecureWay Firewall Reference*.

Firewall Password

The firewall password allows more secure passwords and password rules than the Windows NT logon password so this is the recommended choice for passwords.

Note: Give Telnet access to those users who will need to change their password and instruct them to Telnet to the Firewall and change their password when required. Remember, only the Telnet protocol will prompt the user to change their password.

Require User to Change

Click Yes or No to indicate whether the user is required to change their password the next time they are authenticated.

Lock Password

Click Yes or No to indicate whether the password is locked. This is set to Yes when the maximum number of failed logins is exceeded or when the password has not been used for the number of weeks specified in Maximum Time Before Lockout.

The administrator can set this field to yes to prevent a user from using password authentication.

Notes:

1. Passwords are case-sensitive. If you enter a user's password in mixed-case, the user must then enter the password identically. If you have workstations that work in uppercase only, enter passwords for those users in uppercase.
2. The operating system allows you to define password rules. These password rules apply when a user changes his or her password but not when an administrator makes password changes. Password rules are:

Warning Days Before Expiration (days)

Number of days before a password expires in which a warning will be given to the user informing them of the impending password expiration.

Maximum Weeks Before Expiration

Number of weeks before the user is required to change the password.

Maximum Weeks Before Lockout

Number of weeks after the password has expired that the user can change the password. The administrator must reset the password after a lockout.

Maximum Login Retries Allowed

Maximum number of failed login attempts before the password is locked.

Passwords Before Reuse

Number of passwords stored in the password history list. The password cannot be changed to any password that is currently in the history list. This parameter is only valid if Weeks Before Password Reuse is zero.

Weeks Before Password Reuse

Number of weeks passwords are kept in the password

history list. The password cannot be changed to any password that is currently in the history list.

Minimum Length

Minimum number of characters in a password.

Minimum Alphabetic Characters

Minimum number of alphabetic characters in a password.

Minimum Other Characters

Minimum number of non-alphabetic characters in a password.

Maximum Repeated Characters

Maximum number of times any single character can be repeated in the password.

Minimum Different Characters

Minimum number of different characters in the password.

Click the **Firewall Password** tab to customize these values for each user, as shown in Figure 31 on page 125.

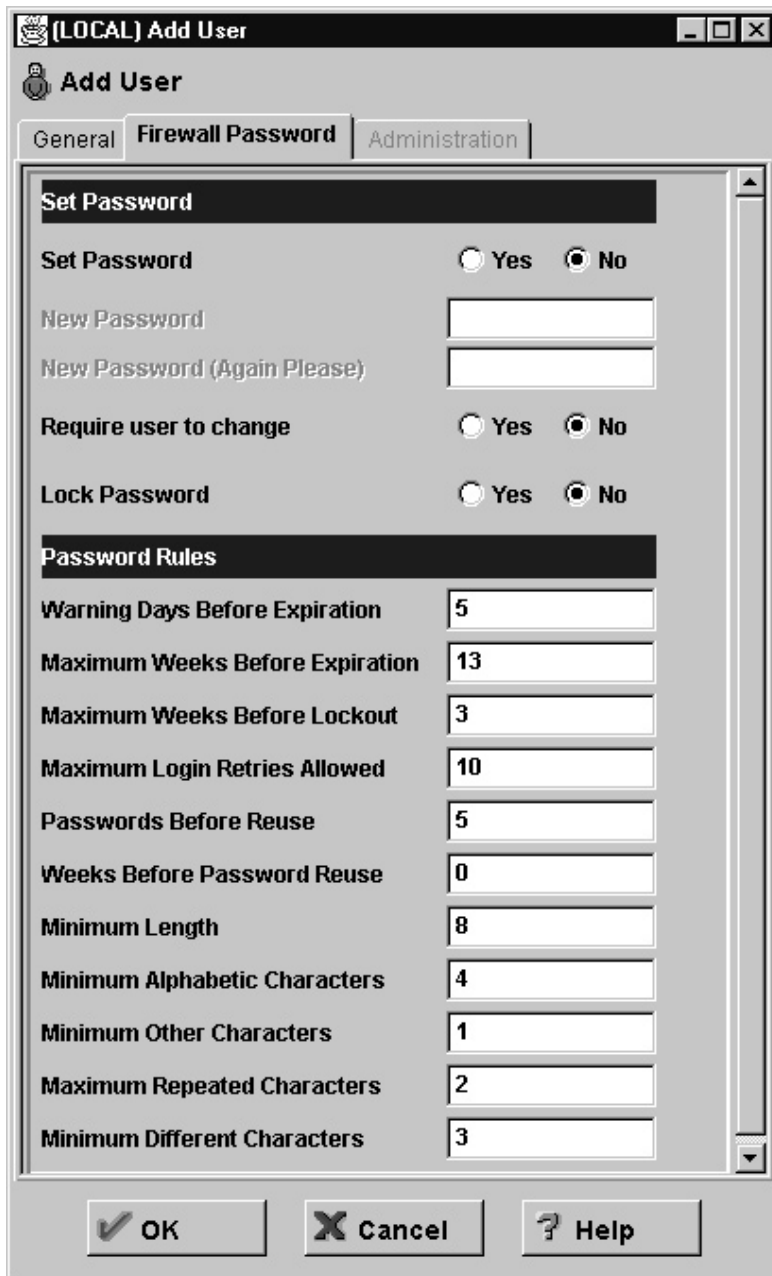


Figure 31. Firewall Password Tab

Changing a User's Access

After you add a user to the Firewall, you can change that user's security attributes from the **Modify User** dialog box.

1. Select the user you want to change from the **Users** dialog box and click **Open**.
2. When the **Modify User** dialog box appears, change the appropriate fields. See "Adding a User to the IBM Firewall" on page 115 for a list of user attributes that you can change.
3. When you have made the changes, click **OK**.

Deleting a User from the IBM Firewall

Note: Do not delete the users `fwdfuser` or `fwdfadm`.

To delete a user, click **Delete** on the **User's List** panel.

Administrator Authority Level by Function

Only the *primary firewall administrator* can create and modify administrators and determine which firewall functions they will have authority over. For example, you can limit a particular administrator to just having the authority to perform the Users and Log Monitor functions.

On the **Add User** dialog box, select Firewall Administrator for the **Authority Level** field. See "Adding a User to the IBM Firewall" on page 115 for more details on completing the **Add User** dialog box.

Then, select the **Administrator** tab at the top of the **Add User** dialog box. Select which functions the administrator is authorized to use.

Chapter 14. Creating a Virtual Private Network

The IBM Firewall provides support for manually configured virtual private network (VPN) tunnels. A Virtual Private Network is an extension of an enterprise's private intranet across a backbone network, which typically will be a public backbone such as the Internet. A VPN allows you to create a secure connection to protect your data while it is in transit over the backbone. The VPN tunnel uses the open IPSec security standards to protect your data from modification or disclosure while it is travelling between firewalls. Your data will flow within a VPN tunnel, which can provide data origin authentication, confidentiality, and integrity checking on every packet. VPN protocols can keep your data private, hiding it from any eavesdroppers on the public network. Packet filtering in the firewall can be used with VPN technologies to further protect your intranets from unwanted intrusions.

Manual Tunnels

VPNs can be created by manually configuring VPN tunnels between pairs of IBM Firewalls or between an IBM Firewall and any other device (client, router, server, or firewall) that supports the latest open IPSec standards. IBM Firewall supports the encryption algorithms including DES, 3DES, and CDMF and authentication algorithms including HMAC-MD5 and HMAC-SHA. Filters for your VPN tunnels can be customized through the creation of connection rules or by choosing a default set of dynamic filter rules.

The tunnel-defined policy specifies that the data (original IP packets) be either:

- Encrypted
- Authenticated
- Encrypted and authenticated

The user determines the tunnel policy or level of protection based on security requirements. A different policy can be used for different IP protocols, for example, Telnet may be different from FTP. The concept of a tunnel carrying encrypted or authenticated data or both is integrated with the IP filtering rules.

Filters are used to permit or deny traffic and direct datagrams into or out of specific VPN tunnels. The IBM Firewall offers two types of filters for VPNs: static and dynamic. Static filter rules are manually configured and must be explicitly activated for them to take effect. Dynamic filter rules are implicitly activated when a VPN tunnel is activated.

Currently the dynamic filter rules have very coarse granularity: they will direct all traffic between a given pair of endpoints into a specific tunnel. The static filter rules offer you the ability to have a much finer granularity. For example, you can define filter rules that will allow one set of applications to communicate over a given tunnel, while constraining a different set of applications to communicate over a different tunnel. Thus even if two tunnels have the same endpoints, they can support different policies. For example, one tunnel might use DES encryption, while the other uses 3DES. Examples of additional qualifiers that can be used to create static filter rules are protocols, ports, traffic, and direction.

VPN requires that each of the two parties (tunnel endpoints) have shared secret keys. The tunnel definitions, including the values of the shared secret keys, are created by one firewall endpoint and manually distributed to the other endpoint before the tunnel can be activated.

Supported authentication algorithms are Hashed Message Authentication Code (HMAC) using Message Digest 5 (MD5) or Secure Hash Algorithm (SHA).

The supported encryption algorithms are:

- Commercial Data Masking Facility (CDMF)
- Data Encryption Standard (DES)
- Triple DES (3DES)

Encryption is a process of scrambling cleartext data so that it is very hard to tell from the encrypted data (or ciphertext) what the original data was. Encryption algorithms typically convert cleartext into what look like random bit patterns through a fairly complex mathematical process. The data is recoverable only because the ciphertext is not really random but, in fact, was created using a key. If the key is not known, then the process of analyzing the data and recovering cleartext can be difficult. One common cryptanalytic method is known as the brute force method, where all possible keys are used for decryption until meaningful cleartext is obtained. If the key is short, the number of possible keys is small, and brute force can be a reasonably powerful attack strategy. If the key is long, the number of possible keys can be far too many to reasonably find the correct one in a short amount of time. CDMF has a 40-bit key length, while DES keys are 56 bits long, and 3DES keys are 168 bits long.

Notes:

1. CDMF keys are really 64 bits long, but many of the bits are duplicated in a known fashion so that the effective key length is 40 bits.
2. DES keys are converted from 56 bits to 64 bits, but again in a known process, so that the key length is effectively 56 bits.

3. Triple DES is actually performed with three DES keys, each of which is 56 bits long, so that the effective key length is 168 bits.
4. One of the weaknesses of a simple cryptographic system like CDMF, DES, 3DES and most others, is that the same cleartext yields the same ciphertext with the same key. So, if a ciphertext message has been compromised once (that is, the cleartext is known), then until the key is changed, the ciphertext is compromised. This weakness is overcome with a method known as cipher block chaining. All the encryption algorithms used in the IBM Firewall use cipher block chaining (CBC), so the algorithms are conventionally denoted DES_CBC and 3DES_CBC.

Because an operational IP tunnel requires administration definitions at both of the tunnel endpoints, those endpoints and the administration tasks associated with IP tunnel operations are defined in terms of **tunnel owner** and **tunnel partner**.

Tunnel Type

The IBM Firewall allows you to create a manually configured VPN tunnel that uses the latest IPSec Authentication Header (AH) and Encapsulating Security Protocol (ESP) to protect your traffic while it is in transit over a nonsecure backbone. The IBM Firewall V4R1 can establish a VPN tunnel with another IBM Firewall V4R1 or any box (for example, other firewalls, routers, or hosts) that supports manual configuration of the latest AH and/or ESP protocols.

Figure 32 is an illustration of a tunnel and a VPN.




Figure 32. Tunnel, All IP Traffic between Two Secure Networks. FW₁ and FW₂ represent nonsecure interface IP address and mask. SN₁ and SN₂ represent any host in the secure network. The shaded area of the picture represents a VPN.

IP Tunnel Configuration and Activation

To configure and activate tunnels, use the configuration client to take appropriate steps depending upon where you are in the configuration.

For the tunnel owner:

- Define your tunnel. See “Add a Tunnel” on page 131.
- If you are using static filter rules, add connections to allow the required firewall-to-firewall communication.

- Export a tunnel from the tunnel owner to the tunnel partner.
- Activate the tunnel.
- Activate the filter rules if you are using static filters.

For the tunnel partner:

- Import (load) the tunnel.
- If you are using static filter rules, define the connections.
- Activate the tunnel.
- Activate the filter rules if you are using static filters.

If you are using dynamic filters and have a nested tunnel configuration, activate the outer tunnel before you activate the inner tunnel. Deactivate the inner tunnel before you deactivate the outer tunnel.

Note: Do all tunnel administration using the IBM Firewall configuration client.

Tunnel Configuration with Endpoints in the Same Subnet

If you have a tunnel with endpoints in the same subnet, you must create routing table entries for the clients on the secure side of the partner firewall, which indicate routing through the partner firewall. This will cause the TCP/IP stack to put the correct destination MAC address in the MAC headers. Otherwise, any time you have VPN tunnel traffic, the firewall will get ICMP Type=5 redirect messages from the router (one for every packet sent).

Example of Virtual Private Networks Using Static Filter Rules

To establish a Virtual Private Network requires an intricate configuration. In this case, the packets being sent between the client and the host are encapsulated for their journey between the two firewalls. For this reason, each packet passes through the filter mechanism twice: once in its encapsulated form and once in the clear. On each iteration, the packet looks completely different because the header is rewritten, and therefore requires a different connection to permit its passage.

The client, in the secure network, will be sending packets addressed to the remote host. These packets will be permitted by the Service *VPN Traffic 1/2*. Next, the same packet, still addressed to the remote host from the client in the secure network, will be encrypted and/or encapsulated based upon the tunnel being used by the service *VPN Traffic 2/2*. (It is recommended to copy this service once for each tunnel being used. Each copy would reference a single tunnel ID, and any connections to that VPN would include the appropriate

copy of this service). Once the packet has been encrypted and/or encapsulated, the Firewall sends the encapsulated packet to the remote firewall directly, where the packet will be decapsulated and/or decrypted and sent to its destination.

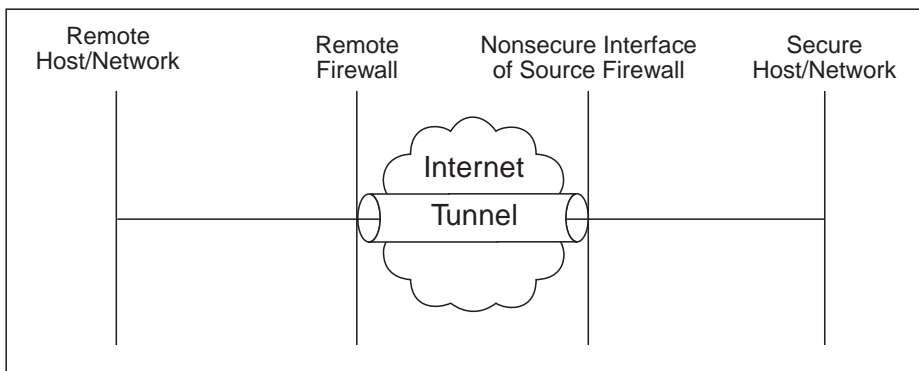


Figure 33. Virtual Private Networks

Such a configuration requires the following connections:

Table 6. Virtual Private Networks

Source Object	Destination Object	Services Required
Secure Host/Network	Remote Host/Network	<ul style="list-style-type: none"> • VPN traffic 1/2 • VPN traffic 2/2
NonSecure Interface of the Source Firewall	Remote Firewall	VPN encapsulation

Configuring Tunnels Using the Configuration Client

This section describes how to use the configuration client to configure your tunnels on the firewall.

Select **Traffic Control** from the configuration client navigation tree. Double-click the file folder icon to expand the view. Select **Virtual Private Network**.

From the **Virtual Private Network Administration** dialog box, you can open, copy, delete, import, export, activate, and deactivate a tunnel.

Add a Tunnel

1. Select **NEW** from the **Tunnels** dialog box and click **Open**.

A dialog asks you to specify the values required for a tunnel context ID specification, as shown in Figure 34.

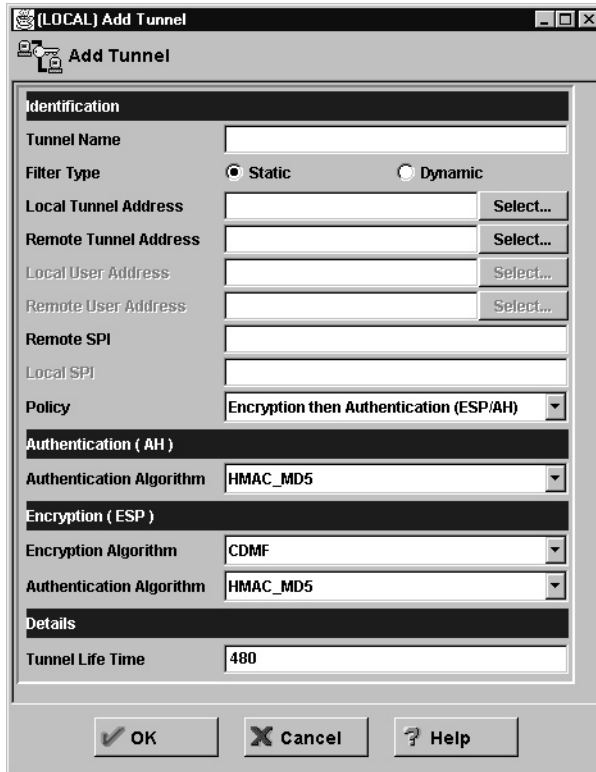


Figure 34. Add a Tunnel

2. Enter the following values:

Tunnel Name

Enter the name of the tunnel.

Filter Type

Select either **static** or **dynamic**. If you select **static**, you must create the tunnels filter rules yourself. If you select **dynamic**, the Firewall will generate dynamic filters each time the tunnel is activated. Selecting dynamic filters allows all supported protocols on any port through the specified tunnel.

If you select the dynamic filters type, you are required to fill in the local and remote user address fields.

For special consideration when using NAT and tunnels together, see “IPSec Tunnel with Dynamic Filters” on page 155.

Local Tunnel Address

IP address of the local firewall interface to be used by the tunnel. Click **Select** to get the Interface list. Select an interface and click **Apply** or enter an IP address directly. The local tunnel address will be added to the Tunnels screen.

Remote Tunnel Address

Click **Select** to select from a list of network objects or create a new one. Or enter an IP address directly. Click **OK**. The address of that network object is entered in the remote tunnel address field.

Local User Address

If you have selected the dynamic filters type, click **Select** to select from a list of network objects or create a new one. The local user address and the net mask are specified by selecting a network object.

Remote User Address

If you have selected the dynamic filters type, click **Select** to select from a list of network objects or create a new one. The remote user address and the net mask are specified by selecting a network object.

Remote SPI

Specifies the security parameter index (SPI) value the tunnel owner will insert into outbound IPSec-protected datagrams that it sends to the tunnel partner. It is usually agreed upon by you and the tunnel partner. This value can be entered in decimal format and must be greater than 255.

Local SPI

Firewall security parameter index is assigned by the firewall when you add a tunnel. You cannot set or change this value. All SPIs are 32-bit random numbers. It will use this value to locate the security agreement properties for inbound IPSec datagrams addressed to itself.

Policy Allows you to enter a combination of encryption and authentication values: AH only, ESP only, or ESP/AH. Click the arrow to select a particular policy. Click **OK** and your selection is added to the Tunnels Definition screen.

Authentication Algorithm (AH)

Authentication Header (AH) authentication algorithm is used for IP packet authentication. HMAC_MD5 or HMAC_SHA are the types available.

Encryption Algorithm

The Encapsulating Security Payload (ESP) algorithm specifies the algorithm used for IP packet encryption. Specify either CDMF,

DES_CBC, 3DES_CBC, or none. If you specify none, you cannot also specify none for **Authentication Algorithm (ESP)**.

Authentication Algorithm (ESP)

The Encapsulating Security Payload authentication algorithm is used for IP packet authentication. Choose from HMAC_MD5, HMAC_SHA, or none. If you specify none, you cannot also specify none for **Encryption Algorithm**.

Tunnel Life Time

Specifies the time in minutes that a tunnel will be operational. The default is 480 (8 hours) and the maximum time allowed is 99999. If you specify 0, the tunnel will not time out; the lifetime is unlimited.

3. Click **OK** and your entries are added to the Tunnels screen.

Modify a Tunnel

You cannot modify an active tunnel; you have to deactivate the tunnel first.

If you modify an encryption or authentication algorithm, you have to export and import again to your tunnel partner.

1. Select a tunnel from the **Tunnels** dialog box and click **Open**.
2. Modify the desired fields on the **Modify Tunnel** dialog box and click **OK**.

Delete a Tunnel

You cannot delete an active tunnel; you have to deactivate the tunnel first.

1. Select the tunnel you want to delete from the **Tunnels** dialog box and click **Delete**.

The configuration client asks you to confirm your request.

2. Click **Yes** to confirm the delete.

The configuration client confirms your request.

You will receive a warning message at filter activation time if rules are encountered which reference tunnel IDs that do not exist.

Note: If you delete a tunnel and then add it again, you have to re-export it.

Export Tunnel Definition Files

As a tunnel owner, after you have defined a set of tunnel definitions, you will export one or more of these definitions to a tunnel partner. The export function writes the tunnel definition to an export file. If the user is using the AIX[®] operating system or an IBM Firewall as the tunnel partner, you can use

that file to import to the tunnel partner. Otherwise, follow the instructions in “Creating a Tunnel to a Product Other Than the IBM Firewall or the AIX Operating System” on page 139.

1. Select tunnels from the VPN Administration panel and click **Export**.
2. Verify the list of tunnel IDs.
3. Enter a directory name in the field.

This directory must have already been created. The directory is where the file containing tunnel information is temporarily placed for export to a partner.

4. Click **OK**.

The tunnel IDs that you have selected are added to the **Export Tunnel Definition Files** panel.

5. When you have completed this operation, the directory contains the file that needs to be moved to your tunnel partner’s machine. If a directory name of `c:\tmp` was used, the name for the export file would be:

```
c:\tmp\ipsec_tun_manu.exp
```

6. Transfer the file to the tunnel partner’s machine by copying the file to a diskette. Note however, if you are exporting the file from a Windows NT Firewall to an AIX server the file format and name of the export file will be unreadable by the AIX server because the long filename is not properly read by the dosutils.

To overcome this problem:

- a. Make sure that you have DOS utilities installed on the AIX server.
- b. Export the file from the Windows NT firewall
- c. Place the diskette in an AIX machine
- d. Issue the following DOS command: `dosread -a 8.3 filename unix long filename.`

An example would be:

```
dosread -a IPSEC_1.EXP ipsec_tun_manu.exp
```

Import Tunnel Definition Files

Obtain the export file from your tunnel partner. Then:

1. Select **Traffic Control, Virtual Private Network**, and click **Import**.

A dialog box appears.

2. Enter the name of the directory where you have placed the file that you have obtained from the tunnel partner.
3. Click **Select**.
4. Select the desired tunnel and click **Apply**. (One or more items can be selected).
5. Click **OK** after making all selections.

6. Click **OK**.

Tunnel Activation Status

Use this procedure to activate or deactivate a tunnels. When you activate a tunnel, dynamic filter rules are automatically activated and the IBM Firewall enables the use of that tunnel. However, if you have static filter rules, you must activate them.

Activate a Tunnel

Select tunnels from the Tunnels dialog box and click **Activate**.

Deactivate a Tunnel

To stop communication at a tunnel, select tunnels from the **Tunnels** dialog box and click **Deactivate**.

Setting Up Static Filter Rules for a VPN

To set up static filter rules for a VPN, do the following:

1. From the configuration client navigation tree, select **Connection Setup**. The **Connections List** dialog box appears.
2. Double click **NEW**. The **Add a Connection** dialog box displays.
3. Under the heading **Identification**:
 - Enter the name of this connection, for example, *Encapsulation for My Tunnel*.
 - Enter a description of this connection that will be helpful to you in identifying it.
 - In the source field, enter the nonsecure IP address of the firewall machine you are entering this data into.
 - Click **Select**. The **Select Network Object** dialog box displays.
 - Click on the network object representing the nonsecure IP address (interface) of the firewall and click **OK**.

If a network object does not exist, you need to create one. See “Using the Configuration Client to Define Network Objects” on page 32.
 - In the destination field, enter the nonsecure IP address of your tunnel partner at the other end of the VPN.
 - Click **Select**. The **Select Network Object** dialog box displays.
 - Click on the network object representing the nonsecure IP address (interface) of the Firewall and click **OK**.

If a network object does not exist, you need to create one. See “Using the Configuration Client to Define Network Objects” on page 32.

4. Under the heading **Connection Services**:
 - Click **Select**. The **Select a Service from the List** dialog box displays.
 - In the **Search** field type **VPN**. Click **Find**.
 - Click on the Service: *VPN encapsulation* and then click **OK**.
The **Add a Connection** panel redisplay.
 - Click **OK**.
Your newly created connection now displays on the **Connection List** panel. Note that this connection permits encrypted data between the source and destination firewalls. See FW1 and FW2 in Figure 32 on page 129.

You need a second connection to give a secure host or network the ability to pass data to another secure host or network. See SN1 and SN2 in Figure 32 on page 129.

1. From the configuration client navigation tree, select **Connection Setup**. The **Connections List** dialog box appears.
2. Double-click **NEW**. The **Add a Connection** dialog box displays.
3. Under the heading **Identification**:
 - Enter the name of the connection, for example, *VPN traffic for My Tunnel*.
 - Enter a description of this connection that will be helpful to you in identifying it.
 - In the source field, enter the secure IP address of a host or a network on the secure side of the source firewall.

Note: If you are using network address translation, make sure you enter the private untranslated address of the host or network in the source field.

- Click **Select**. The **Select Network Object** dialog box displays.
 - Click on the network object representing the secure IP address (interface) of the host or network and click **OK**.
If a network object does not exist, you need to create one. See “Using the Configuration Client to Define Network Objects” on page 32.
 - In the destination field, enter the secure IP address of a host or a secure network on the secure side of the partner firewall.
 - Click **Select**. The **Select Network Object** dialog box displays.
 - Click on the network object representing the secure IP address (interface) of the host or network and click **OK**.
If a network object does not exist, you need to create one. See “Using the Configuration Client to Define Network Objects” on page 32.
4. Under the heading **Connection Services**:

- Click **Select**. The **Select a Service from the List** dialog box displays.
 - In the **Search** field type **VPN**. Click **Find**.
 - Click on the Service: *VPN traffic 1/2* and then click **Apply**.
 - Click on the Service: *VPN traffic 2/2* and then click **Copy**.
5. Under the heading **Identification**, enter the the name of this service in the service name field. For example, *VPN Traffic 2/2 for My Tunnel*.
 6. Under the heading **Service Override Values**:
 - In the **Override Log Control** field, you can choose whether or not to log this service to your log file.
 - In the **Override Frag. Control** field, you can choose whether or not to allow packets to be fragmented.
 - In the **Override Tunnel ID** field, click **Select**. The **Select a Tunnel** screen displays.
 - Select the tunnel ID that will be used between the two secure clients or networks previously defined in this second connection.
 - Click **OK**.
 7. The **Select a Service from the List** screen redisplay. Click on the service you just created and then click **OK**.
 8. The **Add a Connection** screen redisplay. Click **OK**.
 9. Reactivate filter rules.

Setting Up Tunnels Using Dynamic Filters

To set up a tunnel using dynamic filters:

1. On the **Add Tunnel** dialog box, select **Filter Type Dynamic**.
2. Click **Select** in the **Local User Address** field to select from a list of network objects or create a new one. By selecting a network object, you are specifying the secure local host address or subnet.
3. Click **Select** in the **Remote User Address** field to select from a list of network objects or create a new one. By selecting a network object, you are specifying the destination host address or subnet.
4. Export the tunnel definition to your tunnel partner. See “Export Tunnel Definition Files” on page 134.
5. Activate tunnels. See “Tunnel Activation Status” on page 136.

On the partner firewall:

1. Obtain the export file from your tunnel partner. See “Import Tunnel Definition Files” on page 135.

- Note:** If the tunnel owner is using network address translation, change the remote user address to the translated NAT address. See “IPSec Tunnel with Dynamic Filters” on page 155 for more information.
2. Activate tunnels. See “Tunnel Activation Status” on page 136.

Firewall Interoperability

Because the IBM Firewall supports standard IPSec headers, it can establish a manual tunnel connection with any competitor firewall, router, host, another IBM Firewall, AIX 4.3.0 or 4.3.1, that also supports the standard IPSec headers.

How to Use the IBM Firewall with the AIX Operating System

To establish a tunnel between an IBM Firewall and the AIX operating system 4.3.1:

1. Create a tunnel on the IBM Firewall.
2. Export the tunnel definition into a file following the instructions in “Export Tunnel Definition Files” on page 134.
3. Run the **conv_export_file** utility on the file to convert it to a format that AIX can understand. See “How to Use the conv_export_file Utility”.
4. Go to the AIX operating system to import the file in order to set up a tunnel definition.

Note: We recommend that you create the tunnel on the IBM Firewall and export it to the AIX operating system. This is because the IBM Firewall only supports encr/auth.

How to Use the conv_export_file Utility

Use the **conv_export_file** utility to convert from the IBM SecureWay Firewall V4R1 file format to the AIX 4.3 file format. The command syntax is:

```
conv_export_file [dir=dddd]
```

dir=dddd

Specifies the directory of the location of the export file to be converted. It defaults to the current directory.

Creating a Tunnel to a Product Other Than the IBM Firewall or the AIX Operating System

If you want to create a tunnel between an IBM Firewall and a VPN endpoint that is not the AIX operating system or another IBM Firewall, do the following:

1. Create the tunnel definition on the IBM Firewall.

2. Export the tunnel definition into an export file.
3. Use the information in the export file to configure the tunnel partner.

The following table describes each line in the **export file**, including what it is and any restrictions on it. The **tunnel owner** is the IBM Firewall. The **tunnel partner** is the other end of the tunnel. Each tunnel exported to the export file will have its own set of information in the file.

Keep in mind that the tunnel owner and tunnel partner need to be reversed on another firewall.

Field in Export File	Restrictions
String indicating start of new tunnel definition	Starts with "#"
IP version number	IBM SecureWay Firewall V4R1 only supports version 4
IP address of tunnel endpoint of tunnel owner	IP address is in dotted decimal format
IP address of tunnel endpoint of tunnel partner	IP address is in dotted decimal format
Tunnel identifier used internally by tunnel owner	
Security Parameter Index (SPI) AH used by the tunnel partner	Must be greater than 255
Security Parameter Index (SPI) ESP used by the tunnel partner	Must be greater than 255
Security Parameter Index (SPI) AH used by the tunnel owner	Must be greater than 255
Security Parameter Index (SPI) ESP used by the tunnel owner	Must be greater than 255
Encryption algorithm used by the tunnel partner	Either: <ul style="list-style-type: none"> • CDMF - Commerical Data Masking Facility • DES_CBC - Data Encryption Standard encryption (56 bit) • 3DES_CBC - Data Encryption Standard encryption (168 bit) • None - No encryption being used
Length of key used by tunnel partner encryption algorithm	Length of key in bytes
Key used by tunnel partner for encryption algorithm	Values in hex with 0x preceding it

Field in Export File	Restrictions
Encryption algorithm used by the tunnel owner	Either: <ul style="list-style-type: none"> • CDMF - Commerical Data Masking Facility • DES_CBC - Data Encryption Standard encryption (56 bit) • 3DES_CBC - Data Encryption Standard encryption (168 bit) • None - No encryption being used
Length of key used by tunnel owner encryption algorithm	Length of key in bytes
Key used by tunnel owner for encryption algorithm	Values in hex with 0x preceding it
Authentication algorithm used by the tunnel partner for AH	Either: <ul style="list-style-type: none"> • HMAC_MD5 - Hashed Message Authentication code using Message Digest 5 • HMAC_SHA - Hashed Message Authentication code using Secure Hash Algorithm • None - AH protocol not being used
Length of key used by tunnel partner for AH authentication algorithm	Length of key in bytes
Key used by tunnel partner for AH authentication algorithm	Values in hex with 0x preceding it
Authentication algorithm used by the tunnel owner for AH	Either: <ul style="list-style-type: none"> • HMAC_MD5 - Hashed Message Authentication code using Message Digest 5 • HMAC_SHA - Hashed Message Authentication code using Secure Hash Algorithm • None - AH protocol not being used
Length of key used by tunnel owner for AH authentication algorithm	Length of key in bytes
Key used by tunnel owner for AH authentication algorithm	Values in hex with 0x preceding it
Unused	Always 0
Time (in seconds) that tunnel will be operational	Must be a positive integer. Use 0 to indicate the tunnel will remain operational indefinitely

Field in Export File	Restrictions
Encapsulation mode for ESP protocol	IBM SecureWay Firewall V4R1 only supports tunnel mode
Encapsulation mode for AH protocol	IBM SecureWay Firewall V4R1 only supports tunnel mode
Policy - specifies which protocols to use for this tunnel	First 2 characters of string will be either: <ul style="list-style-type: none"> • ae - Use both AH and ESP protocols • ex - Use only ESP protocol • ax - Use only AH protocol
Replay prevention indicator	Either: <ul style="list-style-type: none"> • 0 - Do not enforce replay prevention • 1 - Enforce replay prevention
Header structure indicator. Determines which set of RFCs to follow when generating structure of protocol headers.	Value is always 1. Windows NT Firewall only supports headers format described in RFCs issued in September 1998.
Authentication algorithm used by the tunnel partner for ESP	Either: <ul style="list-style-type: none"> • HMAC_MD5 - Hashed Message Authentication code using Message Digest 5 • HMAC_SHA - Hashed Message Authentication code using Secure Hash Algorithm • None - Authentication not used by ESP protocol
Length of key used by tunnel partner for ESP authentication algorithm	Length of key in bytes
Key used by tunnel partner for ESP authentication algorithm	Values in hex with 0x preceding it
Authentication algorithm used by the tunnel owner for ESP	Either: <ul style="list-style-type: none"> • HMAC_MD5 - Hashed Message Authentication code using Message Digest 5 • HMAC_SHA - Hashed Message Authentication code using Secure Hash Algorithm • None - Authentication not used by ESP protocol
Length of key used by tunnel owner for ESP authentication algorithm	Length of key in bytes
Key used by tunnel owner for ESP authentication algorithm	Values in hex with 0x preceding it

Field in Export File	Restrictions
Unused	Always 0
Unused	Always "-"
Unused	Always "-"
Tunnel name used by tunnel owner	
Type of filters used internally by tunnel owner	Either: <ul style="list-style-type: none"> • 0 - static filters • 1 - dynamic filters
IP address of machine allowed to send/receive traffic through tunnel. This machine is protected by the tunnel owner.	Valid only when dynamic filters are used.
Mask used in combination with previous field to determine a series of machines allowed to send/receive traffic through the tunnel on the tunnel owner side.	Valid only when dynamic filters are used.
IP address of machine allowed to send/receive traffic through tunnel. This machine is protected by the tunnel partner.	Valid only when dynamic filters are used.
Mask used in combination with previous field to determine a series of machines allowed to send/receive traffic through the tunnel on the tunnel partner side.	Valid only when dynamic filters are used.

Unique SPI Values

When you import a tunnel definition into an IBM Firewall, you may get a message indicating that a duplicate SPI was found. An SPI is a Security Parameter Index. Two SPIs are used by a tunnel: one is used by the tunnel owner, and one is used by the tunnel partner. The SPI used by the tunnel owner must be unique on a IBM SecureWay Firewall V4R1.

When you create a tunnel, the SPI used by the tunnel owner is generated for you and guaranteed to be unique. You must specify an SPI for the tunnel partner. When the tunnel definition is imported into the tunnel partner, the tunnel partner's SPI is checked for uniqueness. If the value was not unique, it will be regenerated. You will get a message indicating that you need to go back to the tunnel owner and change the value you specified for the tunnel partner's SPI, to this new value. If you do not do this, the SPIs will not match, and the tunnel will not work.

Chapter 15. Network Address Translation

Network address translation (NAT) was originally designed as a solution to the IP address depletion problem threatening the development and use of the Internet. The NAT designers realized that private networks could use IP addresses already in use by other private networks, as long as packet traffic remained within that private network. It was not until the packets traveled out of that private network that globally unique IP addresses were required for correct routing.

The original NAT designers envisioned placing NAT hosts at the borders between private and non-private networks. As packets passed through the NAT host from the private network to the non-private network, the packet's source IP address would be changed from a non-unique IP address to a globally unique IP address. The mapping of the two IP addresses would be remembered in NAT tables while the session was active. As packets for the session traveled back into the private network from the non-private network, the NAT host would use the destination IP address to look-up the mapping in its tables and change the destination IP address back from the globally unique IP address to the private network host's IP address. An advantage of this solution is that the translation is transparent to the application running on endpoint hosts of the session. Neither the client nor server knows the translation is taking place and so, no additional endpoint configuration is necessary.

Because at any given time, the number of active sessions, from inside the private network to outside the private network is a relatively small subset of all the hosts on the private network, a private network's hosts could be successfully represented by a relatively small number of globally unique IP addresses. This leads to less demand for globally unique IP addresses.

A side effect of this solution is that IP addresses of private network hosts are actually hidden from hosts outside the private network. Thus, NAT began to be thought of as a security tool because it can be used to introduce a level of obscurity between protected networks and the rest of the world.

The IBM Firewall NAT Implementation

The IBM Firewall NAT implementation provides two types of IP address translation: static and dynamic.

Static Mapping

Static translation is a mapping of a private network IP address to a globally unique IP address that is set by the NAT administrator and exists until a subsequent NAT configuration changes it. The static mapping does not expire or timeout through lack of activity. Only packets from that single private network host are translated with the static, globally unique IP address. No other private network hosts have their packets translated with the same static IP address. Because there is no look-up involved in this one-to-one mapping, sessions can be established in either direction: private network to non-private network or non-private network to private network. An administrator indicates which private network IP addresses are to be statically mapped through NAT **Map** configuration statements.

Dynamic Mapping

A dynamic mapping is established by NAT on a per session basis. As a new outbound session is started, NAT dynamically allocates the resources necessary to translate packets that flow on that session. As the session becomes idle, NAT frees up the resources which can then be used to translate packets for a different session. Understanding dynamic translation requires that one recognize a distinction between session direction and packet direction. Translation resources are only allocated when a new session is initiated in a direction from the private network to the non-private network. Once the session is established, packets flowing in either direction over that session are translated using the NAT IP address mapping for that session. Packets flowing from the non-private network to the private network will only be translated by NAT if a mapping already exists for that packet. If the mapping does not already exist, NAT will do one of two things:

1. If the packet's destination IP address is one of NAT's globally unique IP addresses, the packet is dropped. This is because these packets must be translated to be delivered correctly, and without the mapping it cannot be translated.
2. If the packet's destination IP address is not one of NAT's globally unique IP addresses, NAT does not drop the packet but instead lets the packet pass unaltered. NAT passes the responsibility of ultimately permitting or denying this packet on to the packet filtering code.

Dynamic translation can be implemented in one of several different ways. "Many-to-many" is the informal name for a type of dynamic translation that makes use of a pool of globally-unique IP addresses that are reserved for use by NAT. The number of concurrent sessions supported at any given point in time is limited to the number of IP addresses in the reserve pool.

"Many-to-one" is another type of dynamic translation. "Many-to-one" translation uses one globally unique IP address along with a NAT-allocated (and therefore NAT-unique) port number to perform translation. Theoretically,

the number of concurrent sessions supported at any given point in time by the many-to-one implementation is bounded only by the size of the NAT-allocated port number ((2**16)-1024). Practically though, because this number is so high with many-to-one, the number of concurrent sessions supported at any given point of time is actually bounded by the resources available on the NAT host.

The IBM Firewall NAT implementation uses the "many-to-one" algorithm for dynamic translations because it is easier to configure, it can support a greater number of concurrent sessions, and it is less costly for the customer. An administrator configures the NAT many-to-one IP address through the **many-to-one** configuration statement.

Example of Static and Dynamic Translation

Figure 35 illustrates basic NAT operation in an IBM Firewall environment.

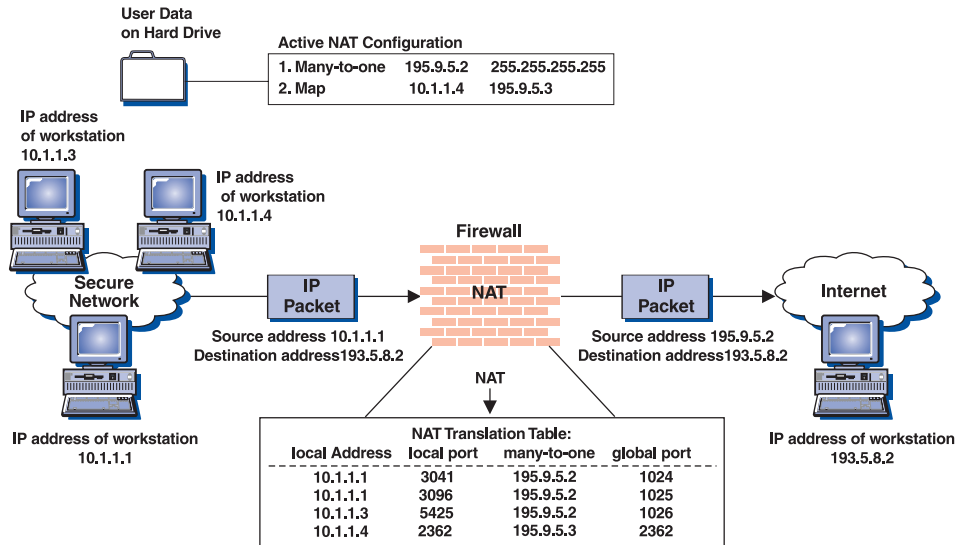


Figure 35. Network Address Translation

In this figure, the NAT many-to-one IP address is 195.9.5.2. There is a NAT Map IP address configured and it statically maps 10.1.1.4 to 195.9.5.3. The private network is called Secure Network and the non-private network is the Internet.

For dynamic translation, TCP and UDP packets generated from the secure network hosts have their source IP addresses (10.1.1.1, 10.1.1.3) replaced by the many-to-one IP address. Their source port numbers (3041, 3096, 5425) are

replaced by NAT-allocated port numbers (1024, 1025, 1026). These mappings are kept in the NAT translation table in kernel memory. To the server operating on the Internet host at 193.5.8.2, all this traffic appears to be coming from different ports at the same IP address.

For static translation, TCP and UDP packets generated from Secure Network host 10.1.1.4 have their source IP address replaced by the Map address. Notice the source port number is not replaced by a NAT-allocated source port number during static translation. Because the IP addresses map one to one, the port number does not need to be modified to distinguish between secure network hosts using this Map IP address. Also, because the mapping is one-to-one, a client executing on the Internet host at 193.5.8.2 can establish an inbound session with a server on the secure network host at 10.1.1.4, provided the TCP or UDP packets are destined with the NAT Map IP address (195.9.5.3). In this way, the IBM Firewall administrator can allow, for example, NAT-translated access to a company Web server or FTP server, by Internet hosts.

Two additional IBM Firewall NAT configuration statements give the administrator more control over exactly which private network IP addresses are dynamically translated and which are not.

The **Translate** statement is used by the administrator to restrict the set of dynamically translated private network IP addresses. If a many-to-one configuration statement exists by itself, **all** private network IP addresses are translated. The Translate statement can be used to modify this behavior by allowing the administrator to explicitly state which private network IP addresses should be translated. If at least one Translate statement exists, all private network IP addresses not within a Translate set are, by default, not translated.

The **Exclude** statement is used by the administrator to define sets of private network IP addresses that are to be excluded from NAT dynamic translation. If a many-to-one configuration statement and an Exclude configuration statement exist, all packets from private network IP addresses within the Exclude set are passed untranslated. All packets from private network IP addresses not within the Exclude set are translated.

For a full explanation of the IBM Firewall NAT configuration statements and their parameters, see “Configuring Network Address Translation Using the Configuration Client” on page 156.

More about Packet Changes Made by IBM Firewall NAT

NAT only performs translation on TCP and UDP packets that are routed. Packets which are local to the Firewall’s secure network are not translated by

NAT. NAT passes these packets unaltered. Note that this means that NAT will not modify packets that are bound for the Socks daemon (or any proxy daemon) operating on the Firewall host because the packets that are handled by these daemons will be local traffic to and from the Firewall. Note that this also means NAT requires your Firewall host to be a router. IP forwarding must be enabled for NAT translation to happen. IP forwarding is enabled by default during IBM Firewall install processing on AIX. On Windows NT, you have to enable it manually. Because the IBM Firewall NAT design depends on distinguishing between local and routed IP traffic, you cannot specify a NAT IP address (static or dynamic) that is the same IP address as one of the Firewall host's network adapters.

NAT performs IP address translation, and (possibly) port translation in IP and protocol headers only. In general, applications that communicate IP addresses and port numbers in packet payloads will not have that data translated by NAT. The exception to this is FTP. Because FTP is so common, special-case code exists in NAT to detect the FTP PORT command in the payload and perform NAT translation on the IP address and port number that accompanies it. After issuing an outbound FTP PASV command, outbound data connections that are established by the secure client work without special-case handling by NAT. This is because the IP address and port number are nonsecure data communicated from the nonsecure server to the secure client. NAT does not support translation of an inbound FTP data connection that was attempted as a result of a nonsecure client issuing an FTP PASV command to a secure server (through static MAP). Other protocols that rely on IP address and port data communicated in the packet payload, such as DNS, CU-SeeMe, RealAudio or SNMP, will not undergo a network address translation. Other solutions to safely transmit these protocols through the Firewall must be employed (such as Socks or proxies).

NAT performs IP address translation on ICMP Request/Reply and ICMP error packets as well as TCP and UDP packets. This means that applications that rely on ICMP packets, such as PING, traceroute on AIX, tracert on NT, and applications that employ PATH MTU Discovery, will work through IBM Firewall hosts that have NAT configured and activated. Unlike TCP and UDP, NAT may translate ICMP packets that are local to the Firewall, if NAT determines it is necessary, based on the current active configuration.

After an intentional or unintentional shutdown of the Firewall host, whatever NAT configuration was active before the shutdown will be automatically reactivated at the next reboot. However, NAT's data structures for keeping track of active dynamic translations are built in kernel heap storage and as such are volatile. The dynamic mappings are not preserved elsewhere. This means that an intentional or unintentional shutdown of the Firewall host will cause all of NAT's current dynamic mappings to be lost. After reboot, you may see a flurry of NAT error messages in the log indicating NAT could not

successfully perform a translation. This is due to packets continuing to flow on the sessions that NAT was dynamically translating before the shutdown. This NAT error logging activity should quiesce as the now broken sessions timeout. NAT static mappings are re-established automatically at the next reboot along with the rest of the NAT configuration statements.

Unlike AIX, on Windows NT, NAT has no way to determine which adapter an outbound packet has come in on. Because it cannot test for this, NAT cannot avoid processing packets twice that are routed from one nonsecure adapter to another. This will corrupt the packet and invalidate the session. Because of this limitation, NAT is not supported in configurations with more than one nonsecure adapter defined.

IBM Firewall NAT and Routing

Before endeavoring to make use of the IBM Firewall NAT function, it is important to take a moment to think about network routing. Because, depending on the packet's direction, NAT modifies a packet's source or destination IP address, and because network routers look at the packet's IP address to determine where to route the packet, you must make some change to your network routing before activating NAT. If you do not make the appropriate routing changes when using NAT, packets are going to get dropped and sessions are going to hang. What routing changes you need to make depend on the IP addresses you are going to define as the many-to-one IP address or as static Map IP addresses.

Refer to Figure 35 on page 147 for the following discussion on NAT and network routing.

Recall that the NAT many-to-one IP address active on the Firewall in the figure is 195.9.5.2. It is not shown in the figure, but let us also say the Firewall host's nonsecure adapter IP address is 195.9.5.1 with a subnet mask of 255.255.255.0. Let us say, for example, that the company depicted in the figure is accessing the Internet through an Internet Service Provider. Finally, let us say the Internet Service Provider's (ISP's) router, which is defined as the default router for the Firewall, has an IP address of 195.9.5.177 and a subnet mask of 255.255.255.0. As packets flow outbound on a session from 10.1.1.1 to 193.5.8.2, NAT changes their source IP address from 10.1.1.1 to 195.9.5.2. The many-to-one algorithm will also change the source port number but that does not impact how the packet gets routed. When the packet is responded to by the server on 193.5.8.2, the response packet created by the server will be destined with the NAT IP address, 195.9.5.2. The firewall administrator must make whatever changes are necessary to get a packet with 195.9.5.2 as a destination IP address back to the Firewall host so NAT can translate the 195.9.5.2 destination IP address back to the proper Secure Network IP address, 10.1.1.1.

So how do you know what changes are required? It all depends on what kind of IP addresses you are using for your NAT IP addresses.

If your NAT IP addresses, whether static or dynamic, are in the same subnet as your Firewall nonsecure adapter's IP address, all that is required is for you to configure your Firewall host to respond to ARP requests for the NAT addresses. In our example, this is because once the ISP's router discovers the destination IP address (195.9.5.2) is local to the subnet, it will broadcast an ARP request to the subnet to discover the MAC address for the corresponding IP address. No host in the subnet will respond to that request because it is not an actual adapter's IP address; it is in use by NAT. So, by configuring the Firewall to respond to ARP requests for the NAT IP addresses with the MAC address of its nonsecure adapter, the ISP's router will properly route the packet to the Firewall.

To configure the Firewall host to make the arp response, go to Control Panel-> Network -> Protocols -> TCP/IP. Select the nonsecure adapter's IP address and then click on advanced. Under IP Addresses, click "Add", and then add the address you are aliasing. In this case, this would be the many-to-one NAT address. You also need to supply the subnet mask for this address.

If the IP addresses you are using for your NAT addresses are not in the same subnet as your Firewall nonsecure adapter's IP address, a static route must be activated on the ISP's router so that it will forward packets destined with your NAT IP addresses to the Firewall host.

In this case it is not required to configure your Firewall to respond to ARP requests for the NAT IP addresses because the static route will cause the ISP's router to ARP for the specified gateway's IP address instead. For our example, this would be the case if the ISP had given the company 195.9.6.2 to use as its NAT many-to-one IP address instead of 195.9.5.2. As the response packet for the session made its way back to the ISP's router, the router would discover the packet was not destined for a local subnet, but by consulting its routing table, it would discover this packet should be forwarded to the gateway that is the Firewall.

The static route to make that happen on the ISP's router in our example would be:

```
route add 195.9.6.2 195.9.5.1
```

IBM Firewall NAT's Position in the Packet-Processing Sequence

Figure 36 on page 152 indicates the position of NAT packet-processing relative to the rest of the IBM Firewall packet-processing functions.

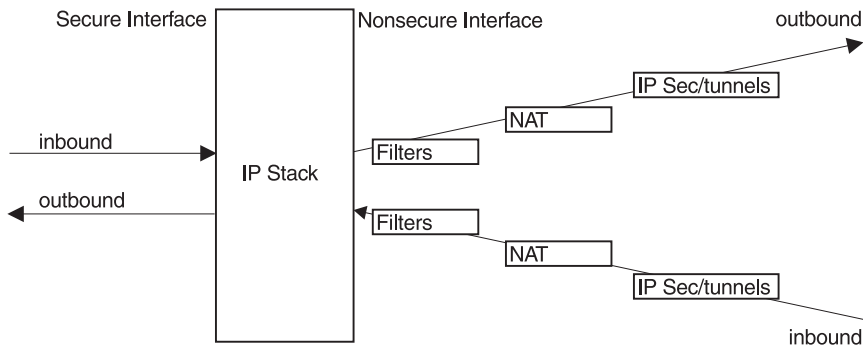


Figure 36. Relationship between NAT, Filters, and Tunnels

When a packet is traveling through the Firewall outbound from the secure network to the nonsecure network, filtering occurs first. If the packet passes the filters, NAT performs its translation (if necessary), and finally if the packet will be traveling through an IPsec Tunnel the encryption for the tunnel occurs. For a packet inbound from the nonsecure network to the secure network, if it came through an IPsec Tunnel, decryption happens first, then NAT translation (if necessary) and finally, filtering will either permit or deny the inbound packet. Figure 36 depicts the conceptual relationship. In actuality, IPsec tunnels do some packet filtering as well to determine, for example, if an inbound encapsulated packet should be allowed into the Firewall to be decrypted. The filters box in Figure 36 represents the filtering that occurs on a decapsulated, decrypted packet.

Note that Figure 36 illustrates that all filtering done by the IBM Firewall is based on secure-side information. You should never configure your filter connections with a NAT IP address as an endpoint. Filter Connections should only be configured with IP addresses for actual network adapters as endpoints (either on the secure or nonsecure side of the Firewall).

IBM Firewall NAT and Filters

After you have completed the NAT configuration, create the filter rules (if you have not already done so) that will allow packets from your secure hosts to flow outbound and the response packets to flow inbound through the Firewall. As stated above, your filter rules should not contain any NAT IP addresses. NAT makes no assumptions based on the NAT configuration statements you have activated and so does not perform any filter connection activation automatically for you.

Review “Chapter 8. Controlling Traffic through the Firewall” on page 63 and use the predefined services that are for direct connections. Examples of predefined services that are for direct connections are:

- HTTP direct out
- Telnet direct out

See “Building Connections Using Predefined Services” on page 65 for more information.

If you want a service to come directly into your network, you will have to create one. See “Using the Configuration Client to Create Services” on page 86 for information on how to do this.

IBM Firewall NAT and IPSec Tunnels

The following scenarios for the interaction between NAT and IPSec Tunnels are described:

- IPSec tunnel with static filters
- IPSec tunnel with dynamic filters
- IPSec tunnel through the Firewall host

IPSec Tunnel with Static Filters

The functional relationship between NAT, filters, and tunnels can be confusing and warrants further clarification. Figure 37 on page 154 illustrates an example interaction between NAT, filters, and tunnels.

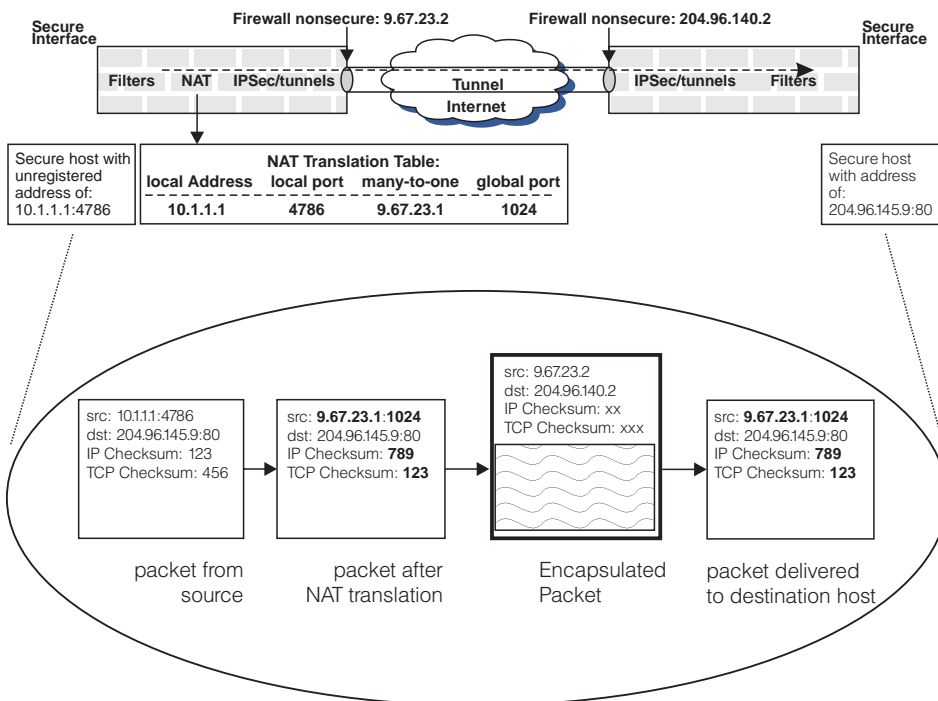


Figure 37. Example Interaction between NAT, Filters and Tunnels

Assume an IPSec tunnel is manually established between Firewalls 9.67.23.2 and 204.96.140.2. NAT is active only at the 9.67.23.2 firewall because this secure network uses unregistered addresses. The secure network at the other end of the tunnel is not using NAT. In addition to illustrating basic NAT translation (the bold fields in the second packet from the left illustrate the fields in the packet that are modified during outbound IP address translation), Figure 37 also illustrates that the translated packet from the host is encapsulated in a packet whose IP addresses are not translated.

In general, filtering is applied to outbound packets prior to NAT and to inbound packets after NAT translation. Therefore the filter rules are based on untranslated addresses. When IPSec tunnels are activated, the filter rules at the firewall that has active network address translation, are still based on untranslated addresses. At the other end of the tunnel (assuming network address translation is not also active at this firewall), the filter rules for inbound packets must be based on NAT-translated source and destination addresses (for the inbound and outbound cases respectively). If NAT is active at both ends of the tunnel the discussion above applies in both directions.

Using the scenario illustrated in Figure 37 as an example, and assuming that the goal is to allow secure host 10.1.1.1 to communicate with secure host

204.96.145.9 over a tunnel, the firewall protecting 10.1.1.1 must have a filter rule permitting 10.1.1.1 to communicate with 204.96.145.9 over a tunnel. At the remote firewall, protecting the destination host, a filter rule is required permitting communication between 9.67.23.1 and 204.96.145.9 through the tunnel.

If you create static filter rules to permit this tunnel traffic with the NAT IP address, be sure to only allow traffic with the NAT IP address *that also came through the tunnel*. Do this by specifying the tunnel's tunnel id on the new static filter rules.

IPSec Tunnel with Dynamic Filters

Tunnels can be created with the dynamic filters feature. This means that the IP filters necessary to allow packets through the tunnel to flow into and out of the secure side of your network are automatically generated in addition to the tunnel definition information.

When a tunnel definition with dynamic filters is activated on the Firewall, the corresponding dynamic filter rules are activated as well and when the tunnel is deactivated the corresponding dynamic filter rules are deactivated. The tunnel definition is exported on the tunnel-owning firewall and imported on the tunnel-partner firewall. The dynamic filters are imported along with the tunnel definition. This feature saves the firewall administrator time and configuration problems when using tunnels.

Be aware that if the tunnel-owning firewall also has NAT configured and active, the dynamic filters imported and activated with the tunnel definition on the tunnel-partner firewall do not include filters to allow the NAT-translated IP addresses to flow into and out of the tunnel partner's secure network. In other words, the IPSec tunnel dynamic filter feature is not aware of NAT or the configured NAT IP addresses.

To modify this and allow NAT's translated packets to flow through a tunnel created with dynamic filters, the partner firewall must modify the imported tunnel definition. The secure host IP address of the tunnel-owning side must be replaced with the tunnel-owning Firewall's NAT address.

IPSec Tunnel Through the Firewall Host

Recall earlier we said that IBM Firewall NAT only performs translation on TCP, UDP, and ICMP packets. For all other protocols NAT passes the packet unaltered. This includes AH and ESP protocols for IPSec Tunnels. The significance is that IPSec Tunnels that are defined across an IBM Firewall host that has NAT enabled will have packets flowing through that tunnel that are not translated by NAT, even if those packets are routed through the Firewall

host and originated from a secure host whose IP address is within NAT's translate set. Because NAT is not altering the AH or ESP protocol header, the IPSec Tunnel will not drop the packet due to (what the tunnel would interpret as) header corruption. This does mean though that the secure host on the IPSec Tunnel endpoint will have its IP address exposed in the tunnel header on the nonsecure network.

IBM Firewall NAT Log Messages

An informational log record is created for each allocation and deallocation of a NAT dynamic (many-to-one) mapping. Allocation occurs when the session is activated and deallocation occurs when the session has become idle. An informational log record is created only once at NAT configuration activation time for each static mapping (map) defined in the new configuration. There are no NAT log records created for each packet being translated through either a static or dynamic mapping. Informational log records are also created at NAT configuration activation and configuration deactivation time.

NAT error log records are created if NAT determines it is required to translate an inbound or outbound packet and cannot for some reason. In error cases, along with the creation of the NAT error log record, the packet is normally dropped.

NAT logging is activated through the NAT Activation panel of the Firewall GUI. NAT logging is independent of filter logging. It can be enabled and disabled without affecting filter logging.

Configuring Network Address Translation Using the Configuration Client

1. From the configuration client navigation tree, double-click the NAT folder icon to expand the view.
2. Double-click **Setup** to create NAT configuration data.

The **Network Address Translation List** dialog box appears, as shown in Figure 38 on page 157.

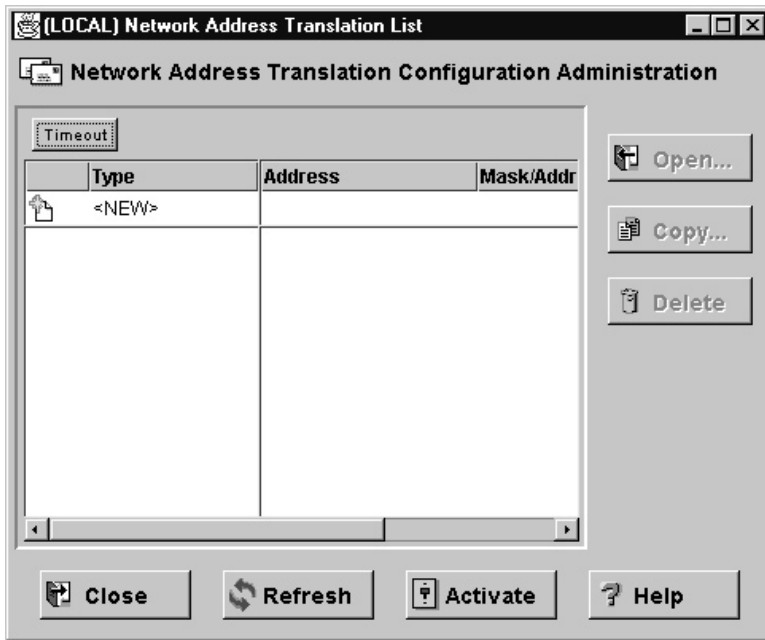


Figure 38. Network Address Translation List

3. The dialog box displays any existing NAT configuration statements. You can modify or delete these or create new NAT configuration statements.

Add NAT Entry

1. Select **New** from the **Network Address Translation List** and click **Open** to add new entries to the NAT configuration file.
The **Add NAT Entry** dialog box appears.
2. From the **Add NAT Entry** dialog box, click the arrow in the Type of NAT field and select from the following:
 - **Many-To-One**: Specifies a globally unique IP address to be used by NAT for dynamic translations.
 - **Translate**: Specifies a range of secure IP addresses that require dynamic translation.
 - **Exclude**: Specifies a range of secure IP addresses that should be excluded from dynamic translation.
 - **Map**: Defines a one-to-one secure-to-registered IP address static translation.

Many-To-One Registered IP Address

A many-to-one registered IP address entry dynamically translates an outbound packet's secure address and port number to a registered IP address and a NAT-unique port number. Many (up to 64512) secure IP addresses can share one registered IP address. Thus you can hide many secure IP addresses with one registered IP address. (You will need at least one additional registered Internet address for the Firewall's nonsecure address).

If you selected Many-To-One from the Add NAT Entry dialog box, enter the following values:

Registered IP Address

The single globally-unique (registered) IP address that NAT should use to perform all dynamic translations. For example, if you are using NAT on a host system between your secure network and the Internet and if your Internet access is provided by an Internet Service Provider (ISP), you would obtain this IP address from your ISP along with the IP address for your Firewall's nonsecure adapter. Enter this address in dotted-decimal form.

Alternately, you can choose a network object by clicking **Select** to get the **Select Network Object** dialog box. Select a network object and click **OK**. The network object's IP address is added to the registered IP address field on the **Add NAT Entry** dialog box.

Timeout Value

The number of minutes a dynamic address translation mapping should remain inactive before NAT will consider it an idle mapping and remove it. This timeout value applies to all mappings created for dynamically translated TCP and UDP packets. The timeout value for ICMP mappings is not configurable. Due to the nature of ICMP traffic, the timeout value of mappings created for ICMP packets has been hardcoded to 5 minutes.

The default is 15 minutes. The range of values is 5 through 45. NAT translation tables can grow very large, especially in high load environments dominated by HTTP traffic. As the table grows, the processing time associated with address translation increases. In high load situations, a timeout value of 5 minutes is recommended to keep the translation table size low.

Translate Secured IP Address

A translate secured IP address entry defines a set of secure network addresses that requires NAT to perform dynamic IP address translation. A translate statement augments the many-to-one dynamic translation statement. Once a

translate set is configured and activated, packets from secure hosts not in a translate set will pass without NAT translation.

If you selected Translate from the Add NAT Entry dialog box, enter the following values:

Secured IP Address

A dotted-decimal IP address that identifies a range of secure IP addresses that require network address translation.

Secured IP Address Mask

A mask, similar to a subnet mask that specifies the bits in the secure IP address used to identify a range of IP addresses. Bits in these masks that are set to 0 indicate that bit positions that have 0 or 1 are included in the range of IP addresses. So, specifying 255.255.255.255 in the mask indicates that only one secure IP address is included in this translation entry, whereas a mask of 255.255.255.0 indicates class C IP addresses require address translation.

Alternately you can choose a network object by clicking **Select** to get the **Select Network Object** dialog box. Select a network object and click **OK**. The network object's IP address and mask are added to the secured IP address and IP address mask fields on the **Add NAT Entry** dialog box.

Exclude Secured Network Address

An exclude secure IP address entry defines a set of secure network addresses that does not require NAT to perform IP address translation. By default, NAT performs address translation on all secure IP addresses in the translate secured IP address set.

If you selected Exclude from the **Add NAT Entry** dialog box, enter the following values:

Secured IP Address

A dotted-decimal IP address that identifies a range of secure IP addresses that should be excluded from network address translation.

Secured IP Address Mask

A mask, like a subnet mask that specifies the bits in the secured IP address used to identify a range of IP addresses. Bits in these masks that are set to 0 indicate that bit positions that have 0 or 1 are included in the range of IP addresses. So, specifying 255.255.255.255 in the mask indicates that only one secure IP address is specified in this entry, whereas a mask of 255.255.255.0 indicates class C IP addresses are excluded from address translation.

Alternately you can choose a network object by clicking **Select** to get the **Select Network Object** dialog box. Select a network object and click **OK**. The network object's IP address and mask are added to the secured IP address and IP address mask fields on the **Add NAT Entry** dialog box.

Map Secured Network Address

A map secured IP address entry defines a one-to-one mapping from a secure IP address to a registered IP address. This one-to-one IP address mapping allows external application clients, such as FTP or telnet clients, to set up TCP sessions with server machines that reside within the secured network. The secure IP addresses in the map secure IP address entries can overlap the secure IP address space specified by the many-to-one and/or translate secure IP address entries. If there is overlap, the static translation takes precedence over the dynamic translation.

If you selected Map from the **Add NAT Configuration** dialog box, enter the following values:

Secured IP Address

A dotted-decimal IP address that should be translated into a specified registered IP address.

Choose a network object by clicking **Select** to get the **Select Network Object** dialog box. Select a network object and click **OK**. The network object's secure IP address is added to the Secured IP Address field on the **Add NAT Configuration** dialog box. Or, type a value directly into the field if you have not previously created a network object.

Registered IP Address field

A dotted-decimal IP address into which a specified secure IP address should be translated.

You can choose a network object by clicking **Select** to get the **Select Network Object** dialog box. Select a network object and click **OK**. The network object's registered IP address is added to the Registered IP Address field on the **Add NAT Entry** dialog box.

Change NAT Entry

Select an existing NAT entry from the **NAT Configuration Administration** dialog box and click **Open** to change the NAT configuration entry.

After you change the NAT configuration, click **Activate**, then choose **Activate/Update Configuration** and click **Execute**.

Delete NAT Entry

1. Select an existing NAT entry from the **NAT Configuration Administration** dialog box and click **Delete** to remove a Network Translation entry from the NAT configuration file.
A confirmation dialog box appears.
2. Select Yes or No.

NAT Activation

1. From the configuration client navigation tree, double-click the NAT file folder icon to expand the view.
2. Double-click **Activation** and a dialog box similar to the one shown in Figure 39 on page 162 appears.

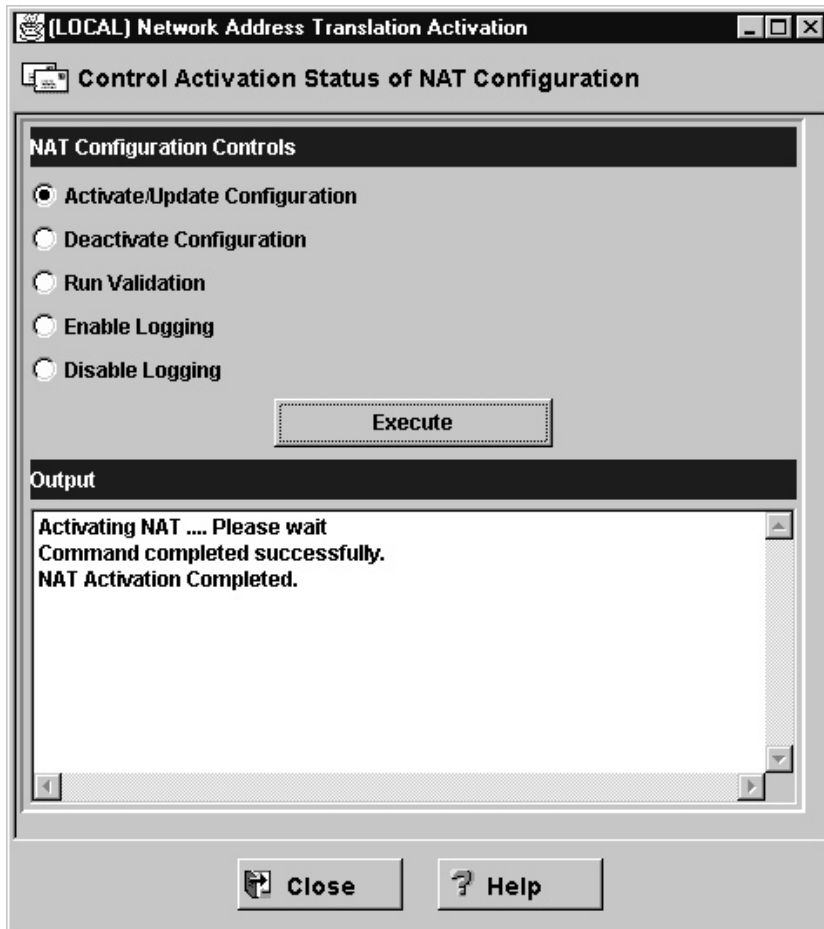


Figure 39. NAT Activation

3. You can select any of the following and then click **Execute**:
 - Activate/Update Configuration
 - Deactivate Configuration
 - Run Validation
 - Enable Logging
 - Disable Logging

NAT Configuration Examples

To illustrate the steps involved in getting NAT configurations to work, we discuss some basic examples below. These examples do not cover all the uses of NAT but give you two common configurations:

1. **Many-to-one configuration:** Allow users in the secure network to use web and telnet services in the nonsecure network, using a NAT many-to-one configuration.
2. **Map configuration:** Allow users in the nonsecure network to access a web server that is in the secure network using a NAT map configuration.

Case 1: Many-To-One Configuration

In this example, a firewall administrator wishes to allow users in the secure network to use web and telnet servers on the nonsecure network. Figure 40 illustrates the network layout in this example.

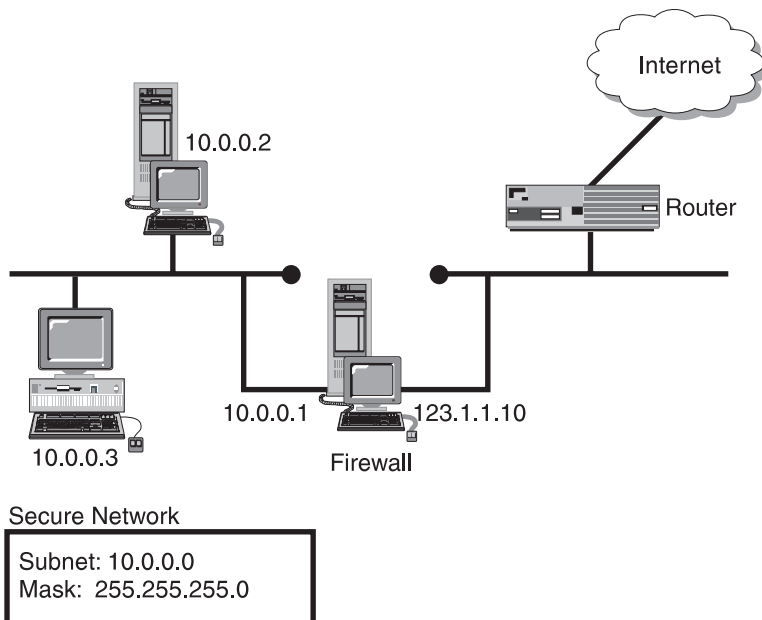


Figure 40. Network Layout

To set up NAT for this scenario, the administrator needs to complete three main tasks:

1. Obtain address and set up routing.
2. Set up connection rules on the firewall.
3. Configure many-to-one NAT configuration.

1. Obtain Address and Set Up Routing

The first task is to obtain a registered IP address that will serve as the one external address that the many internal addresses can be translated into. Also, routing must be set up for this address. When nonsecure hosts attempt to

send packets back to this external address, the routing must be configured so that these packets are sent to the nonsecure interface of the firewall. To accomplish this, the administrator should complete the following steps:

1. Obtain a registered external IP address. For the purposes of this example, we will pretend that the external address is as follows:

IP Address: 123.1.1.11

Mask: 255.255.255.0

2. Establish a route to the Firewall's nonsecure network interface. There are two alternative methods to accomplish this:
 - a. **Add static route to router table:** For performance reasons, this is the preferred method. The administrator needs to access the router tables of the router and add a static route that will route traffic destined for the registered IP address (123.1.1.11) to the firewall's nonsecure interface (123.1.1.10).
 - b. **Add the registered IP Address as an alias to the nonsecure interface. This alternative may be warranted if the administrator is not at liberty to configure router(s) that are outside of the firewall.**

Note: You can only use this method if the firewall's nonsecure address and the 'one' external address are in the same subnet.

If the administrator in this example chose to use this method, he would follow these steps:

- 1) Starting at the NT Start menu, access the appropriate dialog by clicking on the following sequence of options:

Settings
Control Panel
Network
Protocols
TCP/IP Protocol
Properties
IP Address (tab)

- 2) Highlight the firewall's nonsecure adapter in the Adapter window, and click the Advanced... box.
- 3) In the IP Address box, click Add... and then enter
IP Address: 123.1.1.11
Mask: 255.255.255.0
- 4) Close and reboot the firewall.

2. Set Up Connection Rules on the Firewall

To allow the secure users to send traffic through the Firewall, appropriate connection rules must be in place. These rules can be set up by following these steps:

1. Add a connection to the connection setup with the following attributes defined:

Source	Network object that has the following characteristics: IP Address: 10.0.0.0 Mask: 255.255.255.0
Destination	Network Object called "The World"
Services	<ul style="list-style-type: none"> • HTTP direct out • Telnet direct out

2. Regenerate connection rules and activate.

3. Configure Many-to-One NAT Configuration

At this point the administrator is ready to set up NAT on the firewall. To accomplish this, the administrator will access the NAT setup list from the main Configuration Client tree and proceed through the following steps:

1. Add a NAT configuration with the following attributes defined:

Type	Many-to-One
Registered IP Address	123.1.1.11

2. Add a second NAT configuration with the following attributes defined:

Type	Translate
Secured IP Address	10.0.0.0
Secured IP Mask	255.255.255.0

3. Click on the NAT activation panel and execute the Activate/Update Configuration option.

Case 2: Map Configuration

For this example, the firewall administrator wishes to allow users from the Internet to access a web server located inside the secure network. The web server is located at 10.0.0.2. To allow nonsecure users to access this web server, the administrator needs to complete three main tasks:

1. Obtain address and set up routing.
2. Set up connection rules on the firewall.
3. Configure NAT Map configuration.

1. Obtain Address and Set up Routing

Like the first example, the first task is to obtain a registered IP address that will serve as the external address that will be mapped to the internal address. Also, routing must be set up for this address. When nonsecure hosts attempt to send packets to this external address, the routing must be configured so

that these packets are sent to the nonsecure interface of the firewall. To accomplish this, the administrator should complete the following steps:

1. Obtain a registered external IP address. For the purposes of this example, we will pretend that the external address is as follows:

IP Address: 123.1.1.12

Mask: 255.255.255.0

2. Establish a route to the Firewall's nonsecure network interface. We will discuss two alternative methods to accomplish this:
 - a. **Add static route to router table:** For performance reasons, this is the preferred method. The administrator needs to access the router tables of the router and add a static route that will route traffic destined for the registered IP address (123.1.1.12) to the firewall's nonsecure interface (123.1.1.10).
 - b. **Add the registered IP Address as an alias to the nonsecure interface.** This alternative may be warranted if the administrator is not at liberty to configure router(s) that are outside of the firewall.

Note: You can only use this method if the firewall's nonsecure address and the external registered address are in the same subnet. If the administrator in this example chose to use this method, he would follow these steps:

- 1) Starting at the NT Start menu, access the appropriate dialog by clicking on the following sequence of options:

```
Settings
  Control Panel
    Network
      Protocols
        TCP/IP Protocol
          Properties
            IP Address (tab)
```

- 2) Highlight the firewall's nonsecure adapter in the Adapter window, and click the Advanced... box.
- 3) In the IP Address box, click Add... and then enter

IP Address: 123.1.1.12
Mask: 255.255.255.0
- 4) Close and reboot the firewall.

2. Set Up Connection Rules on the Firewall

To allow the nonsecure users to send traffic through the firewall, appropriate connections rules must be in place. These rules can be set up by following these steps:

1. Add a connection to the connection setup that has the following attributes defined:

Source	Network object called "The World"
Destination	Network object that has the following characteristics: IP Address: 10.0.0.2 Mask: 255.255.255.255
Service	HTTP direct in (note that this service is not a pre-defined service shipped with the firewall. To create this service, see the sub-section below).

To create the *HTTP direct in* service:

- a. Copy the 'HTTP 1/2' rule template so that you create a new rule template with the following changes:
Name: HTTP 1/2 (for direct in)
Interface: Nonsecure
 - b. Copy the 'HTTP 2/2' rule template so that you create a new rule template with the following changes:
Name: HTTP 2/2 (for direct in)
Interface: Secure
 - c. Copy the 'HTTP Ack 1/2' rule template so that you create a new rule template with the following changes:
Name: HTTP Ack 1/2 (for direct in)
Interface: Secure
 - d. Copy the 'HTTP Ack 2/2' rule template so that you create a new rule template with the following changes:
Name: HTTP Ack 2/2 (for direct in)
Interface: Nonsecure
 - e. Create a new service with the following parameters:
Name: HTTP direct in
Rule Object: HTTP 1/2 (for direct in)
 HTTP 2/2 (for direct in)
 HTTP Ack 1/2 (for direct in)
 HTTP Ack 2/2 (for direct in)
 - f. Change the flow parameter for both the 'HTTP Ack 1/2' and 'HTTP Ack 2/2' rule objects so that the direction is from right to left. See Figure 41 on page 168 to see what the rule objects look like after the flow parameters have been changed.
2. Regenerate connection rules and activate.

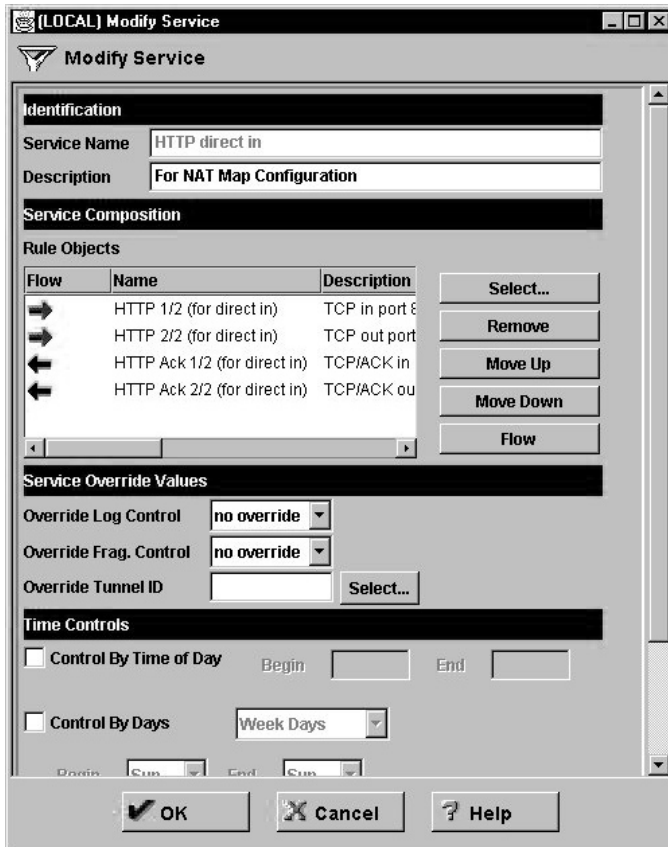


Figure 41. Network Address Translation

3. Configure NAT Map Configuration

At this point the administrator is ready to set up NAT on the firewall. To accomplish this, the administrator will access the NAT setup list from the main Configuration Client tree and proceed through the following steps:

1. Add a NAT configuration with the following attributes defined:

Type	Map
Secured IP Address	10.0.0.2
Registered IP Address	123.1.1.12
2. Click on the NAT activation panel and execute the 'Activate/Update Configuration' option.

NAT Configuration Statements

The following table summarizes NAT's behavior for various combinations of configuration rules. The configuration client or the command line interface (`fwnat cmd=list`) will show the current NAT configuration. If for example, the output of an `fwnat cmd=list` shows:

1. MANY-TO-ONE 204.114.22.3 255.255.255.255 15
2. MAP 10.1.1.1 9.37.51.211
3. MAP 10.1.1.2 9.37.51.212

Then the corresponding NAT behavior description in the table is in the row that has an Active NAT Entry of **MANY-TO-ONE and MAP**.

Active NAT Entry	What NAT Code Does
None	When there are no active NAT configuration file entries, NAT is not active and no secure addresses are translated.
MANY-TO-ONE only	All secure source addresses are translated in all outbound packets.
TRANSLATE only	Outbound packets with secure source addresses matching the TRANSLATE entry are discarded because there is no MANY-TO-ONE entry specifying available external addresses. Outbound packets that don't match the TRANSLATE are allowed through without translation.
EXCLUDE only	Outbound packets with secure source addresses matching the EXCLUDE entry are allowed through without translation. Outbound packets that do not match the EXCLUDE are discarded.
MAP only	Packets with secure source or destination addresses matching the MAP entry are translated. Outbound packets that do not match the MAP are allowed through without translation.
MANY-TO-ONE and TRANSLATE	Outbound packets with secure source addresses matching the TRANSLATE entry are translated because there is a MANY-TO-ONE entry specifying available external addresses. Outbound packets that do not match the TRANSLATE are allowed through without translation.

Active NAT Entry	What NAT Code Does
MANY-TO-ONE and EXCLUDE	Outbound packets with secure source addresses matching the EXCLUDE entry are allowed through without translation. Outbound packets that do not match the EXCLUDE are translated.
MANY-TO-ONE and MAP	Packets with secure source or destination addresses matching the MAP entry are translated. Outbound packets that do not match the MAP entry are translated.
EXCLUDE and TRANSLATE	Outbound packets with secure source addresses matching the EXCLUDE entry are allowed through without translation. Outbound packets with secure source addresses matching the TRANSLATE entry are discarded because there is no MANY-TO-ONE entry specifying available external addresses. Outbound packets that do not match either entry are allowed through without translation even though they do not match the EXCLUDE entry.
EXCLUDE and MAP	Outbound packets with secure source addresses matching the EXCLUDE entry are allowed through without translation. Packets with secure source or destination addresses matching the MAP entry are translated. Outbound packets that do not match either entry are allowed through without translation.
MAP and TRANSLATE	Packets with secure source or destination addresses matching the MAP entry are translated. Outbound packets with secure source addresses matching the TRANSLATE entry are discarded because there is no MANY-TO-ONE entry specifying available external addresses. Outbound packets that do not match either entry are allowed through without translation.

Active NAT Entry	What NAT Code Does
MANY-TO-ONE, TRANSLATE, and EXCLUDE	Outbound packets with secure source addresses matching the TRANSLATE entry are translated because there is a MANY-TO-ONE entry specifying available external addresses. Outbound packets with secure source addresses matching the EXCLUDE are allowed through without translation. Outbound packets that do not match either entry are allowed through without translation even though they do not match the EXCLUDE entry.
MANY-TO-ONE, TRANSLATE, and MAP	Outbound packets with secure source addresses matching the TRANSLATE entry are translated because there is a MANY-TO-ONE entry specifying available external addresses. Packets with secure source or destination addresses matching the MAP are translated. Outbound packets that do not match either entry are allowed through without translation.
EXCLUDE, TRANSLATE, and MAP	Outbound packets with secure source addresses matching the EXCLUDE entry are allowed through without translation. Outbound packets with secure source addresses matching the TRANSLATE are discarded because there is no MANY-TO-ONE entry specifying available external addresses. Packets with secure source or destination addresses matching the MAP are translated. Outbound packets that do not match any entry are allowed through without translation.
MANY-TO-ONE, EXCLUDE, and MAP	Outbound packets with secure source addresses matching the EXCLUDE entry are allowed through without translation. Packets with secure source or destination addresses matching the MAP are translated. Outbound packets that do not match either entry are translated.

Active NAT Entry	What NAT Code Does
MANY-TO-ONE, TRANSLATE, EXCLUDE, and MAP	<p>Outbound packets with secure source addresses matching the TRANSLATE entry are translated because there is a MANY-TO-ONE entry specifying available external addresses. Outbound packets with secure source addresses matching the EXCLUDE are allowed through without translation. Packets with secure source or destination addresses matching the MAP are translated. Outbound packets that do not match any entry are allowed through without translation.</p>

Chapter 16. Remote Access Service

Windows NT Remote Access Service (RAS) provides network connections over dial-up, ISDN, or X.25 media using Point-to-Point Protocol (PPP). NDISWAN is a networking driver which is provided as part of RAS and converts the underlying PPP data to resemble Ethernet LAN data.

With some configurations in the Firewall, RAS connections can be used by the Firewall just like normal LAN interfaces.

RAS connections may be configured to permit the following:

- Only connections that are initiated at the Firewall machine (dial-out)
- Only incoming connections (dial-in)
- Both dial-out and dial-in connections

This feature may provide additional security by denying incoming connections from unauthorized callers. If incoming connections are permitted, RAS can use a number of authentication methods to ensure that the caller is authorized to make the connection.

While RAS connections may be used for any purpose, this support was added to permit a modem connection to an ISP (and therefore, the Internet) to be used as a nonsecure interface. RAS can also be used for lightly-used (or backup) connections to either secure or nonsecure networks.

Installing Remote Access Service

Previous versions of the Firewall removed RAS as part of the hardening process. Therefore, if you are migrating from a previous version and hardening was done (which is the normal case), you must reinstall the Windows NT operating system in order to recover the services that were removed.

If you are not migrating from a previous version of the Firewall or the old Firewall was not hardened, you can install RAS on your system. Go to Control Panel-> Network -> Services -> Add and select Remote Access Server. You need to supply the Windows NT installation CDROM when prompted. Reboot your machine following the RAS installation.

You may install RAS either before or after you have installed the Firewall. However, because there are several configuration requirements associated

with RAS, installing RAS first is recommended. Before installing the Firewall, you should test the RAS connectivity to ensure that it is functioning correctly.

If you have applied Service Packs or hotfixes to your system, you must re-apply them after RAS is installed.

Configuring Remote Access System

Follow these steps to ensure that RAS and the Firewall work together:

1. Ensure that the RAS configuration allows for connection access using TCP/IP only. NetBeui and IPX access should not be enabled. Click **Network** on the **Remote Access Setup** panel to perform this configuration.
2. If RAS is configured to allow dial-in, then you must configure the IP address assignment to use a static IP address pool. The use of DHCP to dynamically assign IP addresses is not supported. Also, for a user to be able to dial in to the RAS server, the user record (from the User Management application) must have the dial-in capability activated.
3. If RAS is configured to allow only dial-out, then the phonebook entry for the associated modem must specify the static IP address, which is assigned by the administrators of the system that will be contacted. Dynamic IP address assignment is not supported. The IP address is specified in the TCP/IP parameters section of the phonebook entry.

Configuring the Firewall

1. The RAS interface will appear in the list of interfaces (displayed through System Administration -> Interfaces from the IBM Firewall configuration client). Initially, the interface is considered to be nonsecure. You may change the type just as you would do for a normal LAN interface.
2. Define filter rules, VPN tunnels, and so forth involving the RAS interface in the same way that you use normal LAN interfaces.

Using Remote Access Service with the Firewall

In order to pass data across the RAS connection, the connection must be established between the Firewall machine and the RAS partner machine.

- If only dial-out connections are permitted for the Firewall, you must initiate the connection to the partner machine either through the Modem icon in the control panel or the rasphone command from a command window.
- If only dial-in connections are permitted, you must initiate the connection from the partner machine.

If you intend to use Network Address Translation through a RAS interface, you must obtain a second (fake) IP address from your ISP, see “Chapter 15. Network Address Translation” on page 145. Unlike normal LAN interfaces, it is not possible to assign an alias IP address to a RAS interface and so you must ask your ISP to add an explicit host route statement to their router for the fake NAT IP address. This route statement must indicate that traffic to the fake NAT IP address is to be routed through the RAS IP address. For example, if your ISP has assigned the IP address 10.120.3.46 to your RAS interface and reserved the IP address 10.120.3.47 as your fake NAT address, a route statement similar to this one must be used on the ISP’s router:

```
route add 10.120.3.47 10.120.3.46
```

Chapter 17. Monitoring the Firewall Logging

This chapter describes how to monitor the logging of alerts in real time. An alert is generated when a configured threshold is violated.

The IBM Firewall monitors the messages sent to the firewall log for potential crisis situations, based upon user-defined thresholds. In the event of a threshold violation, the Firewall delivers an alert, in a manner specified by the firewall administrator.

Threshold Definitions

A threshold consists of count and time parameters — if a count (number of specific logon failures (events)) is exceeded in the specified time (minutes), the threshold has been violated and an alert message is generated. Log monitor recognizes four types of thresholds:

1. Total authentication failures - total number of authentication failures by any user ID or host
2. Authentication failures against any user ID
3. Authentication failures originating from any host
4. Occurrences of a message tag in the log

All thresholds can be configured using the configuration client or the command line interface. Any changes to the threshold definitions are picked up automatically by the IBM Firewall.

Note: Any logon failure event that triggers an alert will be removed from the list of logon failure events tracked by the monitor. If you want to see a user ID or host failure alert, we recommend that you set the total number of authentication failures to a high threshold.

Alert Messages

When a threshold has been reached, the IBM Firewall generates an alert message. Delivery of the alert message can take any of the following four forms:

1. Entry in a log file:
 - Through the alert log facility configurable through the configuration client or the command line.
 - In the firewall log

2. Send an e-mail message to a list of users
3. Pager, as configured. See “Pager Notification Support” on page 179.
4. Execution of a user-defined command, with the alert message as the first parameter

The alert message contains information relevant to the particular threshold violation. For example:

```
ICA0001e: ALERT – 20 authentication failures.  
ICA0002e: ALERT – 10 authentication failures for user root.  
ICA0003e: ALERT – 15 authentication failures from host 56.67.78.89  
ICA0004e: ALERT – Tag ICA1234e with 3 log entries.
```

Alert messages and other messages originated by the Log Monitor are not monitored.

Configuring Log Monitor Using the Configuration Client

This section describes how to use the configuration client to configure the real-time log monitor. Select System Logs from the configuration client navigation tree. Double-click the file folder icon to expand the view. Click **Log Monitor Thresholds**.

From the **Log Monitor Threshold Administration** dialog box, you can add, change, or delete a threshold definition.

Add Log Monitor

To add a threshold definition, select **NEW** from the **Log Monitor Threshold Administration** dialog box and click **Open**. The **Add Log Monitor** dialog box appears. Fill in the following fields:

1. Click the **Class type** arrow to choose from the list of class types. Class types are:
 - Mail notification
 - Execute command
 - Per User Authentication Failure Threshold
 - Total Authentication Failure Threshold
 - Per Host Authentication Failure Threshold
 - Message Threshold
2. If you selected class type: Mail Notification, enter an e-mail address. You can define multiple mail notification classes.
All threshold violation messages are sent to the specified e-mail address.
3. If you selected class type: Execute Command, fill in a command filename.
The log monitor will execute this command with the alert message as its first parameter. You can only define one execute command class.

4. If you selected class type: Message Threshold, fill in a message tag, a standard tag from the IBM Firewall log messages that you want to be monitored.
5. If you selected one of the threshold classes, fill in the threshold count field. The threshold count is the maximum number of failed events allowed within the specified time period.
6. If you selected one of the threshold classes, fill in the threshold time field. The threshold time is the number of minutes beginning with the first occurrence of an event.
7. If you selected one of the threshold classes, click Yes or No to indicate whether you want pager notification to be active.
8. Filling in a comment is optional.
9. Click **OK**.

Change a Threshold Definition

To change a threshold definition, select the item to be changed from the **Log Monitor Threshold Administration** dialog box and click **Open**. The **Change Log Monitor** dialog box appears.

1. Enter the changes you want for the threshold count and threshold time fields.
The threshold count is the maximum number of failed authentication messages to be detected within the specified time period. The threshold time is the number of minutes beginning with the first occurrence of a message.
2. Click **OK**.

Delete a Threshold Definition

To delete a threshold definition, select the item to be deleted from the **Log Monitor Thresholds** dialog box and click **Delete**. You will be asked to confirm the deletion. Click **Yes** to confirm. Note that delete does not mean delete from the log file. It means delete the definition.

Pager Notification Support

The Firewall can page a system administrator by sending a message to the administrator's beeper when there are intrusion alerts on the Firewall. To set up pager notification support, you need to configure the following three pager components.

1. Command Customization - This component must be created and modified using the configuration client. It sets defaults for the pager command, which is used by the log monitor and can be used from the command line.

This component will contain a unique entry that defines the pager environment. See “Command Customization” on page 181 for more information on defining and customizing this component.

2. Carrier Administration - You must define a suitable carrier before connecting your modem. This component contains a list of default carriers used in the U.S. If the carrier you are using is not one of these, then add your carrier in this component. See “Carrier Administration” on page 183 for more information.

Validate the existing phone numbers for the carriers by getting these numbers from your carriers. When talking with your carriers, be sure to get the carrier’s modem phone number and other settings that are valid for the particular service you have purchased.

3. Modem Administration - Before connecting your modem, you must create suitable modem definitions. These definitions will contain all relevant modem information that pager notification support will use. This component contains a list of modems that you can choose from. You can add to this list, however some modems might not be compatible with your carrier’s support. See “Modem Administration” on page 184 for information on maintaining modem definitions.

Note: IBM Firewall supports the Tele-AlphaNumeric Protocol (TAP) communications protocol for pager notification support.

What Carriers and Modems are Supported

The carriers database file contains a list of the carriers and related transmission parameters. You can add other carriers. Some of the parameters, besides the carrier’s name and modem phone number are:

- The maximum message length for an alphanumeric pager and the maximum digits for a numeric pager
- The baud rate, parity, data and stop bits length

Before using a particular carrier, make sure that the carrier uses the TAP protocol.

The pager code comes with default modem definitions. These are:

- IBM MOD 448 14400 bps
- IBM 5853 2400 bps
- IBM 7852 28800 bps
- IBM 7855 2400 bps
- Generic Hayes compatible
- US Robotics Courier 9600 bps
- Zoom V.34

The paging facility might not work with all modem/carrier combinations even though they may be "Hayes compatible" and support the TAP protocol. Modems designated V.34bis have proven the most successful. Older, V.32bis modems have worked with some services provided by a carrier but not with other services, even if provided by the same carrier.

Default Configuration Files Supplied with the Firewall

There are several default configuration files that come with the Firewall. For example, sample templates are provided for carriers, modems, and pager configurations. These are provided as samples only and likely require modifications for international users.

Configuring Pager Notification Support

Pager Setup is used to configure the command customization file and to maintain carriers and modems. If you are using a pager, you must use Pager Setup to customize your pager environment before using Log Monitor.

Before starting, you need to get the correct modem phone numbers, pager ID, and modem parameters from your carrier.

To configure pager notification support, select System Administration from the configuration client navigation tree. Double-click the file folder icon to expand the view. Select **System Logs**. Double-click the file folder icon to expand the view. Select **Pager Setup**.

Command Customization

When you select **Pager Setup** you can select a carrier and modem to use and write a pager message.

Command Customization Settings

When you select **Pager Setup** from the navigation tree you get a **Pager Setup** dialog box with Command Customization Settings similar to the dialog box shown in Figure 42 on page 182.

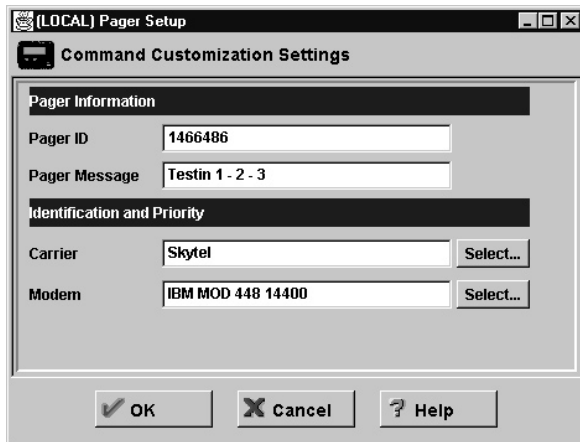


Figure 42. Pager Setup

Type or select values in the entry fields to be added.

1. Enter the pager ID. This is usually a unique PIN assigned to your pager by your carrier company.
2. Enter the pager message. This is a string containing the default message the user wants to send. For numeric pagers, this must be a number only. For alphanumeric pagers, this can be a text message. Do not exceed the maximum message length specified in your carrier setup or your message might be truncated. Do not use a colon (:). If you do, it will be replaced by a blank space character.
3. If there is no carrier name, click **Select** to define a carrier. You will get the **Pager Carrier Administration** dialog box. See “Carrier Administration” on page 183 for details on how to fill in this panel.
4. If there is no modem name, click **Select** to define the modem. You will get the **Pager Modem Administration** dialog box. See “Modem Administration” on page 184 for details on how to fill in this panel.
5. Click **OK**.

Change Command Customization

When you select Pager Setup from the navigation tree you get the **Pager Setup** dialog box with Command Customization Settings.

1. Type or select values in the entry fields to modify the values of the existing customization entry fields.
2. Click **OK**.

Delete Command Customization

1. You can delete an entry on the **Pager Carrier Administration** dialog box or the **Pager Modem Administration** dialog box by selecting an item from the list and double-clicking **Delete**.

You will be asked to confirm the deletion.

2. Click **Yes** to confirm the deletion or **No** to return to the **Pager Setup** dialog box.

If no customization entry exists, then pager notification support will not be able to send a page.

Carrier Administration

From the **Pager Setup** dialog box, go to the carrier name field and click **Select**. You get a **Pager Carrier Administration** dialog box similar to the one shown in Figure 43.

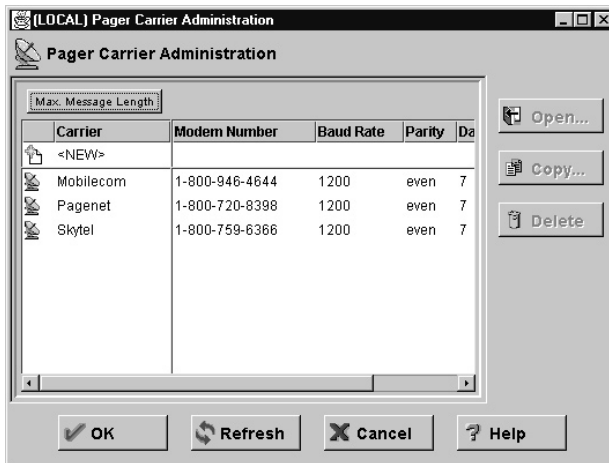


Figure 43. Pager Carrier Administration

Add a Carrier

To add a new carrier select **NEW** on the **Pager Carrier Administration** dialog box and click **Open**. Type or select values in the appropriate entry fields:

1. Enter the carrier name. This can be anything as long as it is unique and provides enough information for you to recognize which carrier it is.
2. Enter the carrier phone number, which is the phone number for a modem at the carrier company, as opposed to their voice paging or other service number. It must be the right modem number for regional or national coverage and for a numeric or alpha pager, as required by the paging device and the service you have contracted.

3. Enter TAP for paging method, the only value allowed.
4. Enter the password if the carrier allows or requires one.
5. Enter the maximum message length for an alphanumeric pager and the maximum digits for a numeric pager.
6. Enter the baud rate. Click the arrow and choose a value from the list.
7. Click **Even**, **Odd**, or **None** for the parity field.
8. Choose the default data bits; click either 7 or 8.
9. Choose the default stop bits; click either 1 or 2.
10. Click **OK**.

Change Carrier

1. Select the carrier you want to change from the **Pager Carrier Administration** dialog box and click **Open**.
2. Refer to “Add a Carrier” on page 183 for an explanation of the fields you can change. The carrier name itself cannot be changed. This field will be disabled.
3. Make your desired changes.
4. Click **OK**.

Delete Carrier

1. Select the carrier you want to delete from the **Pager Carrier Administration** dialog box and click **Delete**.
2. You will be asked to confirm the deletion. Click **Yes** to confirm.

Note: The carrier database must always contain at least one carrier. If no carriers are defined, then pager notification support will fail.

Modem Administration

Your modem manual will contain relevant information about how to initialize your modem. You might need to coordinate modem settings with your carrier. In general, only Hayes-compatible modems, that use the standard modem commands are supported.

From the **Pager Setup** dialog box, go to the modem name field and click **Select**. You get a **Pager Modem Administration** dialog box similar to the one shown in Figure 44 on page 185.

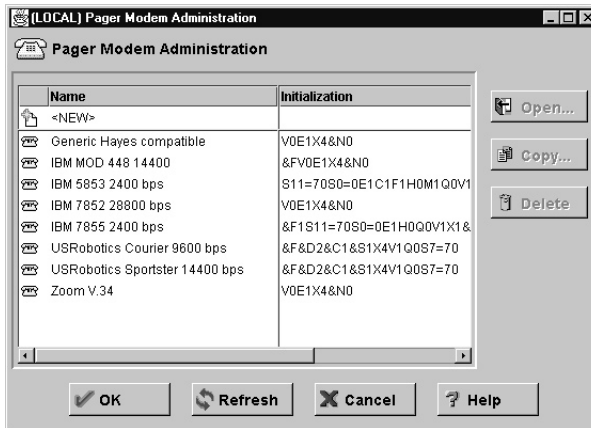


Figure 44. Pager Modem Administration

You can add, change, or delete various modems using this dialog box.

Add a Modem

To add a new modem definition file, select **NEW** from the **Pager Modem Administration** dialog box and click **Open**. On the **Add Modem** dialog box, type or select values in the entry fields.

1. Enter the modem name. This can be anything as long as it is unique among the other definitions and provides enough information for you to recognize which modem it is.
2. Enter the COM Port number, which defines the serial COM Port to which the modem is attached. Enter a number less than 10. While the modem must be hardware-configured to this port, it must not be defined to Windows NT; doing so will cause the pager functions to be denied access to the port. If the modem does not match the hardware settings, the pager code will retry for a long time and eventually fail.
3. Enter the initialization string, which should define the modem as a data modem with an echo on X level4 and a fixed baud rate defined by the local site. Do not include the AT command. The pager function will put it at the beginning of the initialization string.
4. Enter the outside line prefix. This is the number you dial to get outside of your company.
5. Click **OK**.

Change Modem

1. Select a modem name from the **Pager Modem Administration** dialog box and click **Open** to change a modem definition file.

On the **Change Modem** dialog box you will see a list of fields you can change for the modem definition. Refer to “Add a Modem” on page 185 for explanations of these fields.

2. Click **OK**.

Delete Modem

1. Select a modem name from the **Pager Modem Administration** dialog box and click **Delete** to delete a modem definition file.
2. You will be asked to confirm the deletion. Click **Yes** to confirm.

Pager Notification Logging

The pager notification process uses the firewall log utility to write output logs. All pager messages and errors are written to the general firewall syslog facility. For more information on how to set up and use your firewall log files, see “Chapter 18. Managing Log and Archive Files” on page 189.

Testing Pager Setup

You can verify your pager setup by using the pager command. See the *IBM SecureWay Firewall Reference* for details. It is strongly recommended that you use the pager command any time you define or change the setup to be sure your system, modem, carrier, and paging devices all talk with each other correctly and that pages can actually be sent and received.

Execute Commands

You can specify a program that is invoked each time an alert threshold is reached. To specify a program:

1. Click **Log Monitor Administration** and then double-click **NEW**.
The **Add Log Monitor** dialog box appears.
2. In the **Class Type** drop-down box, select **Execute Command**. This enables the **Command Filename** field of the panel.
3. In the **Command Filename** field, enter the fully-qualified pathname of the program you want to invoke when an alert threshold is reached.

The working directory for the commands executed by the log monitor is `\winnt\system32`. Because the command shell is launched from a system process, only system environment variables are set. User environment variables are not set. In general, the launched program should use fully qualified file names instead of relying on path variables.

The Firewall will pass the full Alert message as the first parameter of the program as follows:

Total Authentication Failure Alerts: ICA0001e
Per User Authentication Failure Alerts: ICA0002e
Per Host Authentication Failure Alerts: ICA0003e
Message Threshold Alerts: ICA0004e

See the *IBM SecureWay Firewall Reference* for a complete description of these messages.

Chapter 18. Managing Log and Archive Files

This chapter describes how to use the log facilities through the configuration client. As users try to access hosts through the various IBM Firewall servers, the IBM Firewall writes entries in the log file maintained by the IBM Firewall logging service.

The IBM Firewall can generate large volumes of logging information depending on how you configure your firewall. Log entries can come from a variety of places such as socks and static filters. Additionally, log files can be written to at a variety of severity levels; for example, *debug*, *information*, or *error*. This chapter also tells you how to use the log management and log archive management facilities to manage the size of your log and archive files.

Log File Creation Using the Configuration Client

You can use the configuration client for log management and log archive management. It is assumed that your available disk space is sufficient to contain all the log information. The Firewall generates routine debug and error information to the `firewall log` facility. Only the primary firewall administrator has access to the `firewall log` facility. Alert messages go to the alert log facility. Administrative audit log information goes to the audit log facility.

For report utilities to function properly, it is important that only firewall log messages appear in their input files. No other facility should be directed to the same file as `firewall log` so set firewall logging accordingly.

If you want to see alerts on the main configuration client panel, you have to direct your alerts to a file designated as an alert log facility. Nothing else should be designated for that file.

The following priority levels are cumulative with *debug* capturing the most information. *Critical* captures only the most severe firewall events.

- Debug
- Information
- Warning
- Error
- Critical

It is suggested that you begin with the *information* level until your firewall procedures are stable. Then you can change to *warning* or *error* to reduce the logging activity and the size of the system log.

The priority levels do not correspond precisely to the message tag suffix (*i,e,w,s..*). You might need to experiment to determine how to *shut off* certain messages.

Add Log Facilities

From the configuration client navigation tree, double-click the System Administration file folder icon to expand the view. Double-click the System Logs file folder icon to expand the view. Select Log Facilities. The **Log Facilities** dialog box appears displaying the set of log facilities currently enabled.

1. Select **NEW** from the **Log Facilities** dialog box and click **Open** to add a syslog entry to those currently enabled.

The **Add Log Facilities** dialog box appears, as shown in Figure 45.

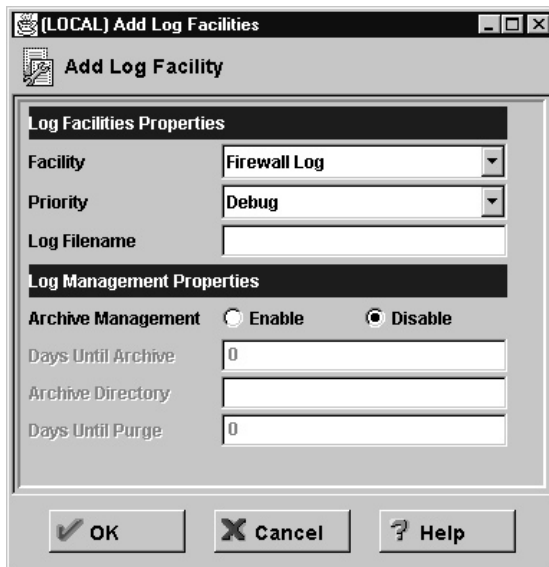


Figure 45. Add Log Facilities

2. The log facility determines the type and source of information that gets logged. Click the **Facility** arrow to select one of the following log facilities:
 - Firewall log - general firewall logs, including filter logging

- Alert log - log monitor daemon status and threshold violation warnings used to populate the Alerts Display
3. Click the **Priority** arrow to choose the priority. The logging priorities are listed in order of increasing severity. The priority you select will be the minimum severity level to be logged because prior levels are cumulative.
 4. Fill in the log filename. The log filename must have an absolute path (beginning with a drive, colon, and a backslash \) and the path to the file must exist.
 5. Archive management only works with the Firewall log. When enabled, the log file size can be reduced on a periodic basis. Enabling archive management means that you set parameters upon which the `fwlogmgmt` command depends. See “Archiving Logs Using the Configuration Client” on page 192.
 6. In the **Days Until Archive** field, enter the number of full days, until record(s) in the active log should be archived. The value must be zero or greater. Archival will occur when an `fwlogmgmt -l` command is issued and records in the active log meet the specified criteria. For example, with a 2 in this field, active log entries that are more than 2 days old will be moved to an archive file.
 7. Enter the name of an archive directory. The IBM Firewall provides an archiving function, which uses the directory you named. However, you can write a Log Archiver plug-in to replace the IBM Firewall archiving function. See the *IBM SecureWay Firewall Reference* for information on the Log Archiver Plug-in software development kit.
 8. In the **Days Until Purge** field, enter the number of full days until an archived log file should be deleted from the archive directory. The value must be zero or greater. Purge will occur when an `fwlogmgmt -a` command is issued and an archived file exists that is older than the specified number of days.
 9. Click **OK**.

Change Log Facilities

1. Select the firewall logging entry you want to change from the **Log Facilities** dialog box and click **Open**.
The **Change Log Facilities** dialog box will appear.
2. Change the desired fields. See “Add Log Facilities” on page 190 for an explanation of the fields.
3. Click **OK**.

Delete Log Facilities

1. Select a firewall logging entry from those currently enabled on the **Log Facilities** dialog box and click **Delete**.
The **Delete Warning** panel appears.

2. Click **OK** if you want to continue with the delete. Click **Cancel** if you change your mind.

Archiving Logs Using the Configuration Client

The archival process:

- Moves qualifying records from the active log to a separate file
- Compacts the resulting file with the Windows NT compact command

To reduce the size of your active log file:

1. Run the `fwlogmgmt -l` command from the command line periodically, or
2. Set up the `fwlogmgmt -l` command as an NT Schedule Service.

To keep the number of archived files to a minimum:

1. Run the `fwlogmgmt -a` command from the command line periodically, or
2. Set up the `fwlogmgmt -a` command as an NT Scheduled Service.

Qualifying records and files are determined by the values specified in the log facilities definitions, as described in “Add Log Facilities” on page 190.

The most efficient and convenient means of running the log management process is to set it up as a Windows NT Schedule Service event. Start the Windows NT Schedule Service by using the Services object on the control panel. Click Startup and set the startup to automatic. When the Windows NT machine starts, this service will start also. Click Start to start the service and continue.

To schedule `fwlogmgmt` events, you must issue a command string at the command line. For example, to schedule log management to archive active log records at 2:00 AM every day, type:

```
at 2:00 /every:M,T,W,Th,F,S,Su fwlogmgmt -l
```

Or, to schedule the deletion of old archive files every Sunday at 9:00 AM, type:

```
at 9:00 /every:Su fwlogmgmt -a
```

Plug-in DLL

See the *IBM SecureWay Firewall Reference* for information on the log archiver plug-in DLL that you can use to replace the Firewall default DLL.

Log Management Outputs

The log management facility does some preliminary integrity checks before proceeding with any log management activities. If any problems are found, diagnostics are sent to the firewall log facility when you run the `fwlogmgmt` command from the command line.

The admin audit (local0) log facility is subject to different archival rules than other facilities. All log facilities require that archival be enabled in order to be archived. However, firewall (local4) and alert (local1) log records are only archived if their dates exceed criteria specified in the facilities definition at the time the archive process is run; whereas the *entire* audit log file will be archived each time.

Report Utilities

You can use the report utility functions to assist you in generating reports from current or archived log files. Report utilities generate tabulated files of administrative information that are organized and formatted for easy mapping to relational database tables. These tables help the firewall administrator to analyze:

- General usage of the Firewall
- Errors in the firewall process
- Attempts at unauthorized access to the secured network

Using the utilities and the firewall log, the administrator can create a regular text file of the messages. Additionally, tabulated files can be generated and imported into tables in a relational database system, such as the DB2[®] family of products. The administrator can then use the Structured Query Language (SQL) to query the data and generate reports.

Report Utilities are installed as part of the Firewall installation. They can also be separately installed and run on a non-firewall host. The configuration client can be used to run them on a firewall. On a non-firewall machine, use the command line.

For report utilities to function properly, it is important that only `firewall log` messages appear in their input files. No other facility should be directed to the same file as `firewall log`, so set firewall logging accordingly.

Do not try to use report utilities on any log files prior to the IBM Firewall for AIX V3R1. You can however, use report utilities to process log files from the IBM Firewall for AIX V3R1 or later. You can also use them to process the AIX `su` log. See the *IBM SecureWay Firewall Reference* for more detailed information on report utilities.

Running Report Utilities Using the Configuration Client

From the configuration client navigation tree, double-click the System Administration file folder icon to expand the view. Double-click the System Logs file folder icon to expand the view. Select **Report Utilities**. The **Report Utilities** dialog box appears, as shown in Figure 46.

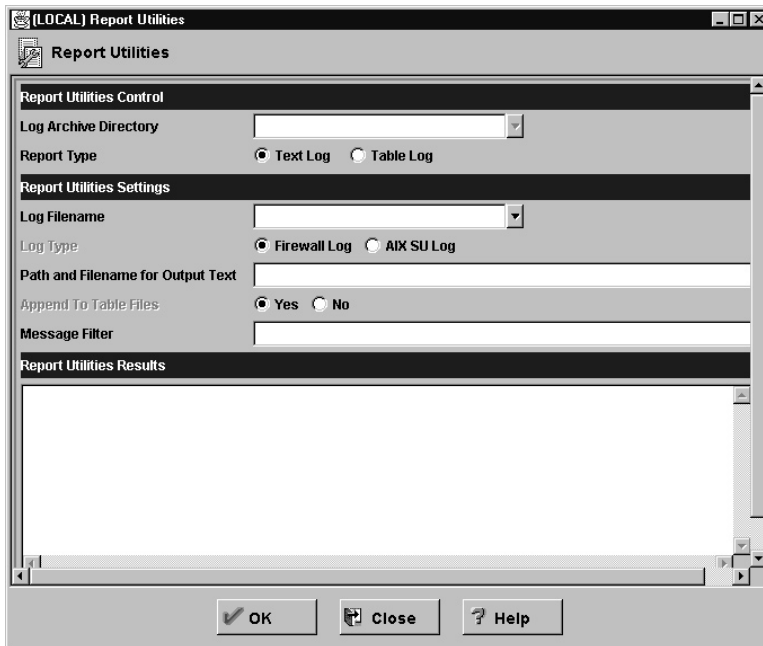


Figure 46. Report Utilities

1. For the default archiver provided with the IBM Firewall, the log archive directory is the directory that contains compressed log files. In the log archive directory field, enter the directory that you specified in the archive directory field on the **Log Facilities** dialog box. Enter the absolute path name to the archive directory. If you want to view a log file that is not archived, leave this field blank.
2. Select the **Report Type**. To produce the expanded log message text, select **Text Log**. To create tabulated files for DB2 usage, select **Table Log**. If you import the resulting files into DB2, you can perform SQL queries on the log data. Refer to the *IBM SecureWay Firewall Reference* for more information.
3. The log filename is any one of the compressed archived log files or other valid firewall log, or the name of an AIX su log file. If you made an entry in the log archive directory field, you can click the **Log Filename** arrow to choose which log to work with. If you do not enter a log archive

in step 1, the log file name you enter here must be the name of a valid, uncompressed firewall log file or an AIX su log file. You must specify a full path.

4. Select the **log type**, either **firewall** or **AIX su**.
5. Enter the **Path and Filename for Output Text**.
6. Select **Yes** to append the results of a table log request to existing tabulated files or **No** to replace the existing files.
7. This field allows you to select certain types of messages to be placed in the output text file. The contents of this field are treated as parameters that are placed into a standard Windows NT Find command. For example, if you type "ICA0" into the field (you must include the quotes), it is as though you are running the following command:

```
fwlogtxt < my.log | find "ICA0"
```

Here are some example entries that you can place into this field and the results:

FILTER	RESULT
"ICA0"	Lists log monitor threshold alert messages
"ICA3"	Lists Socks-related messages (#ICA3000 - 3999)
"ICA2010"	Only lists occurrences of the ICA2010 message
/V "ICA3"	Lists all messages except Socks messages
/C "ICA001"	Counts the number of ICA0001 messages

8. Clicking **OK** produces the requested file(s) in the specified output directory on the firewall machine.
9. The Report Utilities Results area shows any error message from the report utility that was run. To view the log text resulting from a Text Log report type, click **Log Viewer** on the main Firewall configuration client panel, and enter the fully-qualified output file name. The .tbl files resulting from a Table Log report type can be loaded into a database as described in the *IBM SecureWay Firewall Reference*.

Appendix. Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
500 Columbus Avenue
Thornwood, NY 10594
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:
INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make

improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Department TL3B/ Building 062
P.O. Box 12195
3039 Cornwallis
Research Triangle Park, NC 27709-2195
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

This product includes software developed by the University of California, Berkeley and its contributors.

Trademarks

The following terms are trademarks of the IBM corporation in the United States or other countries or both:

- AIX
- DB2
- SecureWay
- IBM
- OS/2
- VisualAge

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States and/or other countries.

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States and/or other countries.

Other company, product, and service names may be trademarks or service marks of others.

Bibliography

For additional information about security on the Internet, visit the IBM SecureWay Firewall home page at:
<http://www.ibm.com/software/security/firewall>.

Information in IBM Publications

Other IBM sources of information on firewalls, Internet security, and general security topics are listed here.

Firewall Topics

The following documents are available on the IBM Firewall CD-ROM and the IBM SecureWay Firewall home page.

- *IBM SecureWay Firewall User's Guide*, GC31-8658
- *IBM SecureWay Firewall Reference*, SC31-8659
- *Guarding the Gates Using the IBM SecureWay Firewall for NT 3.3*, SG24-5209

Internet and World Wide Web Topics

- *A Comprehensive Guide to Virtual Private Networks, Volume 1: IBM Firewall, Server and Client Solutions*, SG24-5201
- *A Guide to the Internet Connection Servers*, SG24-4805
- *Accessing CICS Business Applications from the World Wide Web*, SG24-4547
- *Accessing OS/390 OpenEdition MVS from the Internet*, SG24-4721
- *Accessing the Internet*, SG24-2597
- *Building the Infrastructure for the Internet*, SG24-4824
- *Cool Title about the AS/400 and Internet*, SG24-4815

- *The Domino Defense: Security in Lotus Notes and the Internet*, SG24-4848
- *Examples of Using MQSeries on WWW*, SG24-4882
- *How to Secure the Internet Connection Server for MVS/ESA*, SG24-4803
- *Lotus Domino Server Release 4.5 on AIX Systems: Installation, Customization, and Administration*, SG24-4694
- *Netscape Proxy Server*, SK2T-7444
- *Running CICS Transactions through the Web: The CICS Internet Gateway to VSE/ESA*, SG24-4799
- *Safe Surfing: How to Build a Secure World Wide Web Connection*, SG24-4564
- *Teach Yourself CGI Programming with PERL in a Week*, SR23-7343
- *Using the Information Super Highway*, GG24-2499
- *World Wide Web Access to DB2*, SG24-4716

General Security Topics

- *The Basics of IP Network Design*, SG24-2580
- *Elements of Security: AIX V4.1*, GG24-4433
- *Enterprise-Wide Security Architecture and Solutions Presentation Guide*, SG24-4579
- *HACMP/6000 Customization Examples*, SG24-4498
- *IBM Global Network (IGN) Security Policy*, GC34-2206
- *IBM Security Architecture: Securing the Open Client/Server Distributed Enterprise*, SC24-8135
- *IBM Systems Monitor: Anatomy of a Smart Agent*, SG24-4398
- *Security Overview of Open Systems Networking*, GG24-3815

- *Systems Monitor for AIX User's Guide*, SC31-8173
- *TCP/IP Tutorial and Technical Overview*, GG24-3376

Information in Industry Publications

These industry publications pertain to TCP/IP and UNIX:

- Albitz, Paul, and Cricket Liu. *DNS and BIND*. Sebastopol, CA: O'Reilly and Associates, 1998. (ISBN: 1-56592-236-0)
- Costales, Brian with Eric Allman. *Sendmail*. O'Reilly and Associates, Inc. (ISBN: 1-56592-222-0)
- Hunt, Craig. *TCP/IP Network Administration*. O'Reilly and Associates, Inc. (ISBN: 0-937175-82-X)
- Nemeth, Snyder, et al. *UNIX System Administration Handbook*. Prentice Hall. (ISBN: 0-13-151051-7)

This industry publication pertains to Windows NT:

Cowart, Robert. *Windows NT Server 4.0 Administrator's Bible*. IDG Books Worldwide, 1996. (ISBN: 0764580094)

These industry publications pertain to firewalls and security on the Internet:

- Ahuja, Vijay. *Network and Internet Security*. Boston: Academic Press Professional, 1996. (ISBN: 0120455951)
- Ahuja, Vijay. *Secure Commerce on the Internet*. Boston: Academic Press Professional, 1997. (ISBN: 0120455978)
- Anderson, Bart, et al. *The Waite Group's UNIX Communications and the Internet*. Indianapolis, IN: Sams Pub., 1995. (ISBN: 0672305372)

- Atkins, Derek, et al. *Internet Security: Professional Reference*. Indianapolis, IN: New Riders Publishing, 1996. (ISBN: 1562055577)
- Chapman, D. Brent, and Elizabeth D. Zwicky. *Building Internet Firewalls*. Sebastopol, CA: O'Reilly and Associates, 1995. (ISBN: 1565921240)
- Cheswick, Willam R., and Steven M. Bellovin. *Firewalls and Internet Security*. New York: Addison-Wesley, 1994. (ISBN: 0201633574)
- Cooper, Frederic J., et al. *Implementing Internet Security*. Indianapolis, IN: New Riders Publishing, 1995. (ISBN: 1562054716)
- Curry, David. *UNIX System Security: Guide for Users and Systems Administrators*. Sebastopol, CA: O'Reilly and Associates, 1994. (ISBN: 0201563274)
- Garfinkel, Simson, and Gene Spafford. *Practical UNIX Security*. Sebastopol, CA: O'Reilly and Associates, 1991. (ISBN: 0937175722)
- Garfinkel, Simson, and Gene Spafford. *Practical UNIX and Internet Security*. Sebastopol, CA: O'Reilly and Associates, 1996. (ISBN: 1565921488)
- Hare, Chris, and Karanjit Siyan. *Internet Firewalls and Network Security*. Indianapolis, IN: New Riders Publishing, 1996. (ISBN: 1562056328)
- Randall, Neil. *Teach Yourself the Internet in a Week*. Indianapolis, IN: Sams.Net, 1995. (ISBN: 0672307359)
- Stallings, William. *Internet Security Handbook*. Foster City, CA: IDG Books, 1995. (ISBN: 0077092546)
- Stevens, W. Richard. *TCP/IP Illustrated*. Reading, MA: Addison-Wesley, 1994. (ISBN: 0201634953)

Index

Numerics

3DES 128

A

ACE Server 5
activate socks rules 95
activate tunnel 136
activation, connection 67
add tunnel 131
address depletion problem 6
address translation, network 6, 145
administration 115
administrator authority level 126
AH 129
alert log 22, 189
alert message 177
alert records, view 22
archive files 189
archive management, log 189
audit log 189
authentication, Socks 90
authentication, user 5, 120
authentication profiles, Socks 90
authority level, administrator 126

B

basic configuration steps 27
bibliography 199
blanket policies for firewall, set 30
build a connection 65

C

carriers 180
CDMF 128
chaining, Socks-server 96
change user's security attributes 126
checklist, planning 9
client, configuration 15
clients, socksified 5, 95
Commerical Data Masking Facility 128
components, pager 179
concepts, firewall 1
configuration, default filter 70
configuration client 13, 15, 63
configuration server 13
configuration steps, basic 27
configure DNS 36
configure filters 63

configure Socks server 92
connection, build 65
connection activation 67
connections, order 67
conv_export_file utility 139
CRAM 91

D

Data Encryption Standard 128
deactivate tunnel 136
default filter configuration 70
default network object 31
default set of services 64, 83
define filter rules and services 77
delete rule 82
delete tunnel 134
depletion problem, address 6
DES 128
DNS 35
Domain Name Service 35
domain name services, configure 36
dynamic filters 128

E

Encapsulating Security Protocol 129
encryption algorithms
 3DES 128
 CDMF 128
 Commerical Data Masking Facility 128
 Data Encryption Standard 128
 DES 128
 Hashed Message Authentication Code 128
 HMAC 128
 Triple DES 128
ESP 129
exclude secure IP address 159
expert filters 2
export tunnel 134

F

facility, syslog 186
File Transfer Protocol (FTP) 89
filter configuration, default 70
filter rules and services, define 77
filters, configure 63
filters, dynamic 128
filters, static 2
Firewall, IBM 1
Firewall, logon 17

firewall log 23, 189, 193
FTP 89
FTP proxy 110
fwdfadm 119
fwdfuser 118
fwlogmgmt -a command 192
fwlogmgmt -l command 192
fwlogmgmt command 193

G

gateways, SMTP 49
general security policy 29
generate tabulated files 193
graphical user interface 13, 15
group, network object 33
group, network objects 65
group of network objects 33

H

Hashed Message Authentication Code 128
hiding internal IP addresses 6
HMAC 128
HTTP proxy viii, 99

I

IBM Firewall 1
IBM Firewall tools 2
import tunnel 135
interface, graphical user 13, 15
interfaces 29
interfaces, network
 nonsecure 29
 secure 29
internal IP addresses, hiding 6
interoperability, firewall 139
IP rule, modify 82
IP tunnel 129
IPSec Authentication Header 129

L

log archive management 189
log facilities 189
log monitor, real-time 178
log on to the IBM Firewall 17
log viewer 22, 23
logging, Socks 95
logon, remote 17

M

mail servers, secure 49
management, log archive 189

- manually configured VPN 129
- many-to-one registered address 158
- map secured IP address 160
- MIME 6
- modem administration 184
- modify an IP rule 82
- modify tunnel 134
- Multipurpose Internet Mail Extensions (MIME) 6

N

- NAT 6, 145
- navigation tree 20
- network address translation 145
- network address translation (NAT) 6
- network interfaces
 - nonsecure 29
 - secure 29
- network object group 65
- network objects 65
 - default 31
 - group 31
- network topology 9
- notification support, pager 181

O

- objects, network 31, 65
- order connections 67

P

- pager components 179
- pager notification support 181
- pager setup 181
- partition network 9
- planning checklist 9
- planning worksheets 10
- Point-to-Point Protocol 173
- policy, tunnel 127
- PPP 173
- protocols
 - AH 129
 - ESP 129
- proxies, transparent 111
- proxy, HTTP viii, 99
- proxy, Telnet 111
- proxy services 3

R

- real-time log monitor 178
- references 199
- remote administration 15
- remote logon 17
- report utility functions 109, 193
- rule, delete 82
- rule templates 77

S

- Secure Mail Proxy 6, 49
- secure mail servers 49
- secure network interface 29
- SecurID token 5
- security attributes, change user's 126
- Security Dynamics 5
- Security Parameter Index 143
- security policy, general 29
- security strategy 2
- server, Socks 5
- servers, secure mail 49
- Service, Domain Name 35
- services, default set 64, 83
- services, proxy 3
- set blanket policies for firewall 30
- set of services, default 64, 83
- setup, pager 181
- Simple Mail Transfer Protocol (SMTP) 6
- SMTP 6
- SMTP gateways 49
- Socks 4
- Socks authentication profiles 90
- Socks logging 95
- Socks rules, activate 95
- Socks server 5, 89
- Socks Server, configure 92
- Socks templates 93
- socksified clients 5, 95
- SPI 143
- steps, basic configuration 27
- syslog facility 186

T

- tabulated files, generate 193
- TCP 89
- TCP-based applications 6
- Telnet 89
- Telnet proxy 111
- templates, rule 77
- templates, Socks 93
- token SecurID 5
- tools, IBM Firewall 2
- toplogy, network 9
- traffic monitor, socks 97
- translate secured IP address 158
- translation, network address 6, 145
- Transmission Control Protocol (TCP) 89
- transparent proxies 111
- Triple DES 128
- tuning, Socks-server 95
- tunnel, IP 129

- tunnel policy 127

U

- UDP-based applications 6
- URLs 199
- user authentication 120
- user interface, graphical 13, 15
- user's security attributes, change 126
- username/password authentication 91
- utility, conv_export_file 139

V

- view alert records 22
- virtual private network 7
- VPN 7, 129
- VPN example 130

W

- Web page 199
- worksheets, planning 10

Readers' Comments — We'd Like to Hear from You

IBM SecureWay® Firewall for Windows NT®
User's Guide

Publication No. GC31-8658-03

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Complete	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to find	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to understand	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Well organized	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Applicable to your tasks	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? Yes No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Name

Address

Company or Organization

Phone No.



Fold and Tape

Please do not staple

Fold and Tape




NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM Corporation
Information Development
Department CGMD / Bldg 500
P.O. Box 12195
Research Triangle Park, NC
27709-9990



Fold and Tape

Please do not staple

Fold and Tape



Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

GC31-8658-03

