IBM Firewall for AIX

User's Guide

Version 3.1.1

GC31-8419-00

IBM Firewall for AIX

User's Guide

Version 3.1.1

GC31-8419-00

Note: Before using this information and the product it supports, be sure to read the general information under Appendix A, “Notices” on page 159.

Second Edition (July 1997)

This edition applies to the IBM Firewall Version 3.1.1 licensed program and obsoletes the first edition, dated May 1997.

Publications are not stocked at the address given below. If you want more IBM publications, ask your IBM representative or write to the IBM branch office serving your locality.

A form for your comments is provided at the back of this document. If the form has been removed, you may address comments to:

IBM Corporation
Department CGM
P.O. Box 12195
Research Triangle Park, North Carolina 27709-2195
U.S.A.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Contents

Figures

Tables

About This Book

This book describes how to configure and administer the IBM Firewall Version 3.1.1 on an AIX/6000 system using the configuration client.

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, please see the user's guide for your TCP/IP client programs.

Use the Installation Instructions attached to the CDROM case to install the product before you use this book.

After you start the configuration client, the online help information will help you fill in the configuration client fields and move from menu to menu.

The SMIT interface is available but is not described in this book.

The chapters in this book basically follow the configuration client navigation tree.

Prerequisite Knowledge

It is important that you have a sound knowledge of TCP/IP 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.

What Is New in This Release

The IBM Firewall offers a rich variety of functions.

Java-based Graphical User Interface

In addition to the command line and SMIT interfaces, the IBM Firewall can be administered through a Java**-based graphical user interface (configuration client). The configuration client allows an administrator to perform remote configuration and administration. To ensure confidentiality and integrity the remote configuration connection can be authenticated using any of several mechanisms and encrypted using Secure Sockets Layer (SSL).

Secure Remote Login

An encrypted secure login is provided for remote IPSec Windows 95** clients and for configuration clients through the version 2 Secure Sockets Layer (SSL) technology. The supported SSL cipher specifications for both clients are:

- RC2,MD5,Export=06 with a session key size of 40 bits
- RC4,MD5,Export=03 with a session key size of 40 bits

Enterprise Firewall Management

Enterprise Firewall Management (EFM) provides the means to manage a group of remote firewalls from a single site. This is accomplished by creating an enterprise firewall server that maintains all the configuration files for all of the firewalls. All data is encrypted as it is sent. Access to the enterprise firewall is through the configuration client. An administrator can clone a firewall to create a new one and replicate configuration files to create or update another firewall.

Network Security Auditor

Network Security Auditor is a tool that checks your network for security holes or configuration errors. You will want to periodically verify that the firewall has not been modified in a way that creates a security vulnerability.

By periodically running the Network Security Auditor, you can make sure nothing has changed, especially after you put the firewall on-line.

Secure Remote Client

The Secure Remote Client is software that is installed on a client PC or an AIX workstation offering secure communication. Data sent between a PC and the firewall is encrypted with the 56-bit Data Encryption Standard (DES) and is authenticated. The Secure Remote Client follows IPSec standards.

The Secure Remote Client does not tie you to a specific Point-to-Point Protocol (PPP) server. The TCP/IP address that is assigned by your PPP server is irrelevant. You can change PPP server and TCP/IP addresses and it does not affect the operation of the Secure Remote Client. Other vendors are sensitive to the specific TCP/IP address and if you change the address, you must reconfigure your client.

Report Utilities

Report Utilities generates files of administrative information that are organized and formatted for easy mapping to relational database tables. These tables help the firewall administrator analyze:

- General usage of the firewall
- Errors in the firewall process
- Attempts at unauthorized access to the secured network

The format of the firewall log record is generally not readable. Using the report utilities, the administrator can create a readable text file of the messages. Additionally, tabulated files can be generated and imported into tables in a relational database system, such as DB2/6000 or DB2/2. The administrator can then use the Structure Query Language (SQL), or other tools like IBM's Visualizer or Query Management Facility to query the data and generate reports.

Logging Enhancements

Real Time Log Monitor notifies the administrator of a detected threshold condition on a real time basis.

Log Viewer is a tool for viewing logs from the configuration client.

Alerts viewer provides a view of the alerts through an easy to read formatted screen.

Mail

The IBM Firewall now supports its own Safemail mail gateway. Sendmail 8.7.X has been dropped from the firewall.

Password Rules

Password rules for the firewall now match AIX password rules. The administrator sets passwords to expired, thus requiring users to change passwords on the first use.

Transparent Proxy

Transparent proxy provides easy access from the secure side of the firewall (your private network) to the nonsecure side of the firewall. You can telnet or FTP transparently through the IBM Firewall. Transparent proxies require no firewall authentication, therefore users of transparent proxies do not have to be defined as firewall proxy users.

Filter Enhancements

The filter rules have been enhanced to allow for time-of-day, day, and date selection. For example, you can specify: `permit ftp to IP address between 8:00am and 6:00pm`, or you can restrict the filter validity to a particular day or set of days.

Filter rules allow IP addresses for interfaces (versus secure/nonsecure) to better support multiple interfaces.

Filter storage allocation has been changed from static to dynamic . This allocates less storage than currently required for 512 rules, while allowing the storage to dynamically grow as filters are added.

Host Address Pricing

Firewall licenses are offered on a tiered basis with the price based upon the number of concurrent sessions.

When you purchased the IBM Firewall, you purchased a license for a certain number of hosts. The IBM Firewall tracks the number of unique hosts (IP addresses) and determines when the purchased number of hosts has been exceeded. When the limit has been exceeded, all overflow addresses are logged in the `local4` facility. The following types of messages are displayed in the log file when the limit has been exceeded:

```
License file has been altered, license limit 50 has been loaded.   License
file has been deleted, license limit 50 has been loaded.  License limit xxx
has been read and loaded.  Host xxx exceeds the allowable number of licenses.
```

Concurrent Sessions

TCP and UDP sessions are tracked. There is a maximum of concurrent active TCP and UDP sessions. Once the threshold for each type of session is reached, no additional sessions are allowed unless a grace period had been configured. Sessions that are not allowed can be optionally logged. This function is only available to EFM firewalls.

SNMP

The Simple Network Management Protocol (SNMP), which is widely used in the TCP/IP environment for network management, can also be be used to monitor IBM Firewall server status and generate traps. There are a significant number of SNMP managers existing in customer environments that can be used to monitor the resources and components without introducing the overhead of a management framework and requiring new application programs. Therefore, using SNMP with the IBM Firewall is a natural extension of management of IBM Firewall servers.

HTTP Proxy

A Hypertext Transfer Protocol proxy efficiently handles browser requests and responses through the firewall. Filter rules permit or deny HTTP transactions.

HACMP

The IBM Firewall continues to provide protection in the event of a hardware failure. Firewall operations are automatically shifted to a backup system. The technology for maintaining business critical applications is called High Availability Cluster Multi Processor (HACMP) for AIX version 4.2. It is the leading high availability technology for UNIX. If a hardware failure occurs, a backup system takes over within seconds to maintain network availability.

For more information on HACMP see:
http://hawww.ak.munich.ibm.com/HACMP/HA-FW/HA-FW.html.

SP Support

Necessary changes are implemented to support the AIX/6000 SP processor. Installation and hardening steps are enhanced for SP configuration requirements.

Default User

A default firewall user, **fwdfuser**, is created during installation. If a user attempting to login is not defined to the firewall, the firewall will authenticate the user with the authentication method defined for **fwdfuser**. This feature supports any user-defined authentication method.

Administration Enhancements

You do not need to be user root to perform administrative functions. Any user designated as a firewall administrator can perform administrative functions. These functions are customizable. You can limit an administrator's authority over specific functions, such as administering proxy users.

Stronger Encryption Support

The IBM Firewall enables an export version of DES. This encryption is available in addition to the currently supported CDMF.

AIX 4.1.5 and 4.2 Support

AIX 4.1.5 and 4.2 are supported, exclusive of the AIX Common Desktop Environment.

IBM Firewall Installable Units

The IBM Firewall separate installable components are:

- EFM
  - IBM Enterprise Management System (a firewall that manages other firewalls)
- FW
  - Base IBM Firewall
  - IBM Firewall Common Libraries and Catalogs
  - IBM Firewall Remote Configuration Client
  - IBM Firewall Report Generation Utilities
- Netscape.NAV
  - Netscape Navigator**
- ipsec
  - IPSec Client
- nsauditor
  - Network Security Auditor
  - Network Security Auditor HTML Interface
- sva
  - System View Agent for AIX
  - SystemView Agent for AIX SNMP Mapper

For directions on how to install the Windows 95 secure remote client, refer to Chapter 14, “Using the Windows 95 Secure Remote Client” on page 103.

To install the PDF version of this manual and the *IBM Firewall Reference* download the following files from the `fwbooks` directory on the IBM Firewall CDROM to your workstation:

- `fwuser.pdf`
- `fwref.pdf`

Use the Adobe Acrobat** Reader to view these books. If you do not have the Adobe Acrobat Reader installed, you can go to the Adobe Web site at: **http://www.adobe.com/prodindex/acrobat/** to learn more about the Adobe Acrobat Reader and to get a copy.

Entering IP Addresses

When you configure your firewall, you will be asked to enter IP addresses. You should enter a complete dotted-decimal IP address, with all 4 octets, in the format:

```
nnn.nnn.nnn.nnn
```

where each `nnn` is a set of three numbers in the range 000–255.

How to Access Online Help

When using the configuration client to configure or administer the IBM Firewall, you can click on the Help button to get online help for the menu you are using.

Where to Find More Information

For additional information about security on the Internet, see the Bibliography.

Additional information about the IBM Firewall can be found on the firewall home page at URL **http://www.ics.raleigh.ibm.com/firewall.**

How to Call 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.

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 Firewall is a network security program for AIX. In essence, a firewall is a blockade between a secure, internal private network and another (nonsecure) network 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 two jobs:

- The firewall lets users in your own network use authorized resources from the outside network without compromising your network's data and other resources.
- The firewall keeps unauthorized users outside of your network.

Firewall Concepts

The any-to-any connectivity of the Internet can give you many security problems. You need to protect your own private data and also protect access to the machines inside your private network against abusive external use. The first step to achieving this is to limit the number of points at which the private network is connected to the Internet. It is best to have a configuration where the private network is connected to the Internet by just one gateway. If you have only one path, you gain a lot of control over which traffic to allow into and out of the Internet. We call this 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.

In the case of the firewall, the classic solution with a screening router is not sufficient today to ensure security. A better strategy is to permit only the applications you have tested and have confidence in. If you follow this strategy, you have to exhaustively define the list of services you must run on your firewall. Each service is characterized by the direction of the connection (from internal to external or external to internal), the list of users authorized, the list of machines where a connection can be issued, and perhaps the range of time of day you authorize this service.

Screening Filters

The first and most commonly used strategy is to separate the private IP network from the Internet by inserting a router between them, as shown in Figure 1.

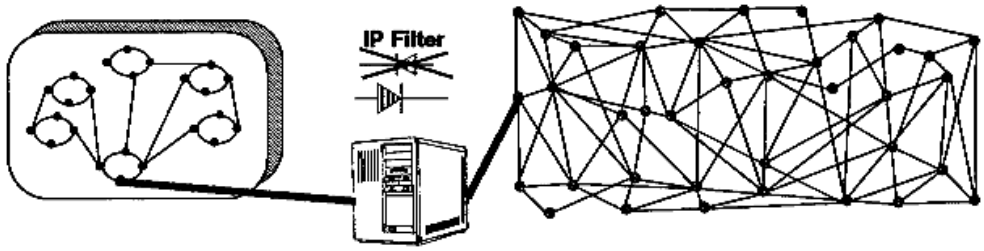

Figure 1. Screening Filter

This router filters all IP packets passing through and is called a screening filter. A screening filter can prevent access to machines or to ports in the private network and also the reverse; the Internet from an internal machine. But if you do this, there is no way to control what is happening at the application layer. That is, you may want to allow one type of traffic across the gateway but not another. You could manage this at the application host itself, but the more machines on which you have to impose controls, the less control you have. Nonetheless a screening filter is a very useful tool to use in conjunction with other tools as a security building block.

Bastion

A bastion is a machine placed between the secure and nonsecure network where the IP forwarding is broken, which means no IP packet can go through this machine. Because the routing is broken, the only place from which you can access both networks is the bastion itself. Therefore, only users who have an account on the bastion, with a double identification (one for the bastion and one for the remote host), can use services on both the networks, as shown in Figure 2.

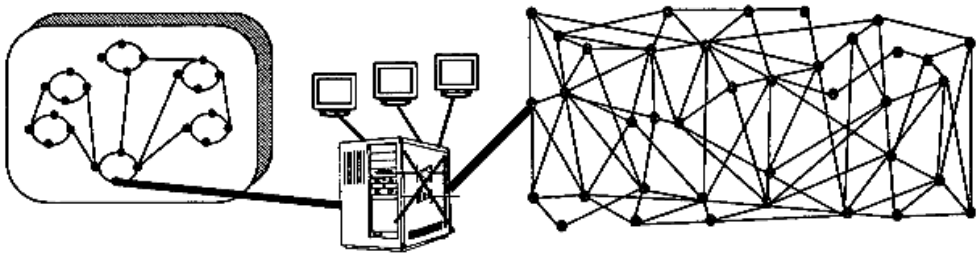

Figure 2. Bastion

It is important to enforce good password control on the bastion because if an intruder manages to break into a user ID, he or she can then impersonate the user and get into the private network. Also, supporting a great number of users will require a big machine. To avoid having many users logged in to this machine and to reduce load on the machine, the bastion concept is combined with the socks server.

Dual-Homed Gateway

One good solution is to combine a screening filter and a bastion, as shown in Figure 3.

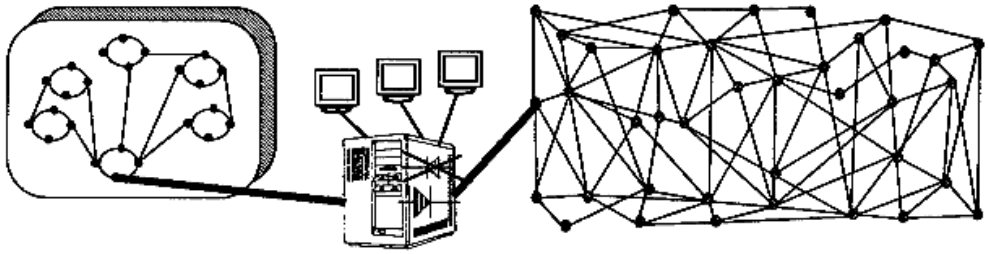

Figure 3. Dual-Homed Gateway

In this case, you can protect the dual-homed gateway from external attacks with filtering. For example, if you forbid external access to the telnet daemon, you reduce the threat of an external attack. If you have some nomadic machines that are hosted outside but need to connect to hosts inside the private network, you can limit the exposure by using a proxy server and perhaps using smart card authentication techniques.

The problem with this configuration is that the firewall machine can become very complex, so if an intruder does break into it, it may take some time to track him or her down. For example, there are a great number of IP ports used by so-called well-known services. Some of these are, in fact, not well-known at all and intruders regularly use them as a back door into a computer. So it is important to block as many such ports as you can, without impacting the services that you do want to work.

Bastion Behind a Screening Filter

A better solution is to use the same solution as above but use two machines as shown in Figure 4.

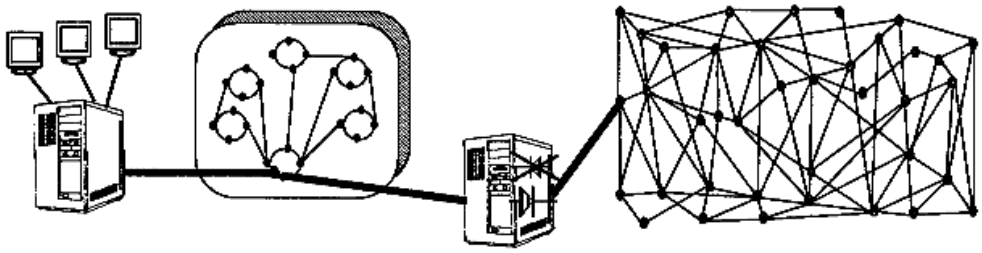

Figure 4. Bastion Behind a Screening Filter

In this configuration, the bastion is protected from external attack by the screening filter. This one is a very simple router without a daemon inside. Because there is no daemon it is very hard to break into this machine.

Screened Subnet

A further development of this is to use the subnetwork between the screening filter and the bastion as a site for application services. This is increasingly common, as organizations want to provide machines that are widely available (such as Web servers) but still have strong protection for their private network. The screening filter provides some protection for the service machines, without unduly limiting access. A possible example of this is shown in Figure 5.

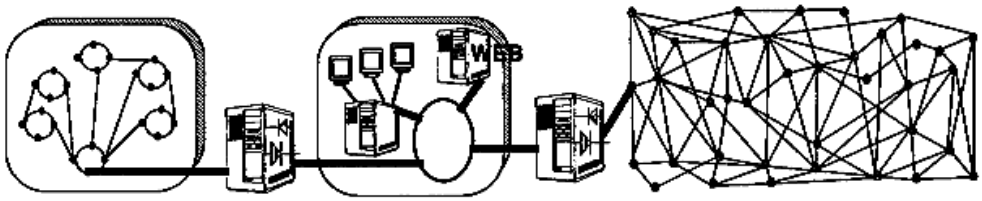

Figure 5. Screened Subnet

This network is composed of two screening filters and one or more bastions. When you start considering this sort of solution, the cost becomes a major factor because, for reasons of integrity, each component in the design should ideally be a dedicated machine. The architecture is simplified when each machine on this subnet performs only a single task(s) depending on the number of bastions you have.

Making Your Firewall Choice

The choice of your firewall architecture depends on your security requirements and also on the size of your organization. A small organization may choose a bastion behind a screening router; a larger organization may choose a screened subnet. Your firewall architecture also depends on the software you choose, and the IBM Firewall is an excellent solution.

As you read further through this book you should be able to decide in more detail which of the IBM Firewall features are appropriate for your environment. Before going into the detail, however, it is worth stating some basic rules:

- Anything that is not explicitly permitted is, by default, denied.

  When you set up your firewall you should be able to state exactly what traffic you want to pass through it. It should not be possible for any other traffic to pass.

- Keep outside users out of your internal network whenever possible.
- Do thorough auditing and logging.

IBM Firewall Tools

The IBM Firewall is like a tool box you use to implement different firewall architectures: both screening filter and bastion. 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 net-

works. 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:

- IP filters
- Proxy servers
- Authentication/Encryption
- Socks servers
- Specific services such as domain name service (DNS) and mail handling
- Virtual Private Networks
- Network Address Translation

IP Filters

IP filters are tools to filter packets at the IP level, based on IP address, direction, and TCP or UDP port. 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 subnetwork. One interface is declared nonsecure and the other(s) declared secure. The filter acts between these two interfaces, as illustrated in Figure 6.

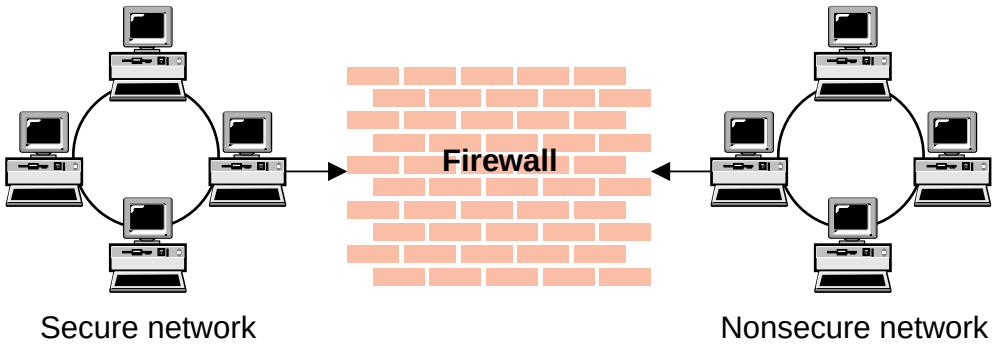

Figure 6. Firewall with IP Filtering

Objectives of the Filters

IP filtering provides the basic protection mechanism for the firewall, allowing you to determine what traffic passes across it 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

This tool allows an internal user, using normal client commands such as telnet, to access the nonsecure network. For example, users use Telnet to access a bastion. They have to be authenticated using a password in the normal way. Having successfully accessed the bastion (and thus authenticated themselves) they now have to again issue the telnet command to reach the desired machine on the nonsecure network. Figure 7 on page 6 illustrates a firewall with a proxy server.

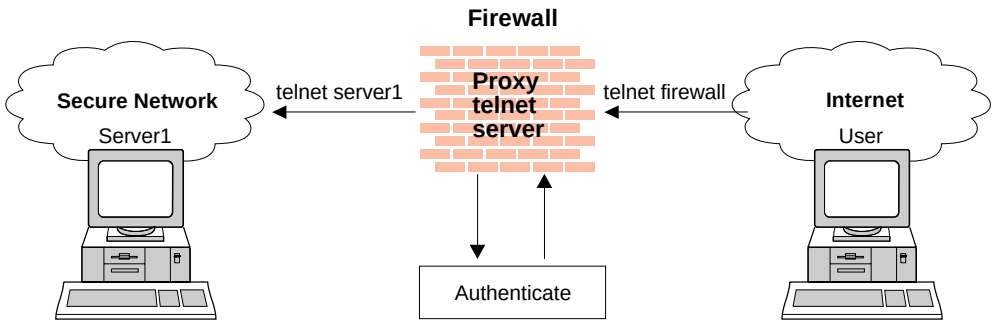

Figure 7. Firewall with a Proxy Server

This method is usable also from the nonsecure network to the secure network but raises other important security problems. If you use the Internet to connect to the bastion, you have to enter your login name and password to be identified. But you cannot know what machines your session may pass through and some intruder may be looking for login names and passwords, by wire tapping or using IP trace commands. If the intruder catches your ID, he or she could impersonate you and thus get into the organization with your identity. (This would only get the intruder to the firewall; they would need to trace the second telnet command to get past it.) In this scenario, with the proxy server, you can use a more sophisticated tool of authentication such as a security identity card. This mechanism generates a unique key which is not reusable for another connection. Two security identification cards are supported by the IBM Firewall: the SecureNet** card from AssureNet Pathways and the SecurID** card from Security Dynamics** Technologies, Inc.

The proxy services available are telnet, FTP, and HTTP. HTTP proxy handles browser requests through the IBM Firewall eliminating the need for a socks server. 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 also have to authenticate themselves, using one of a number of authentication methods.

Once connected, the appearance and the behavior are unchanged. But, the proxy server method needs a double connection: one from the client machine and one from the firewall machine. This double connection has an impact on performance.

A major 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 does not require the overhead of a more conventional proxy server.

The Socks server results in a similar bastion configuration, since the session is broken at the firewall. The difference is that the user does not need to perform the double login manually. Instead of directly starting a session with the `telnetd`

daemon, the session goes to the `sockd`daemon on the IBM Firewall host. `sockd` 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 8 illustrates a firewall with a socks server.

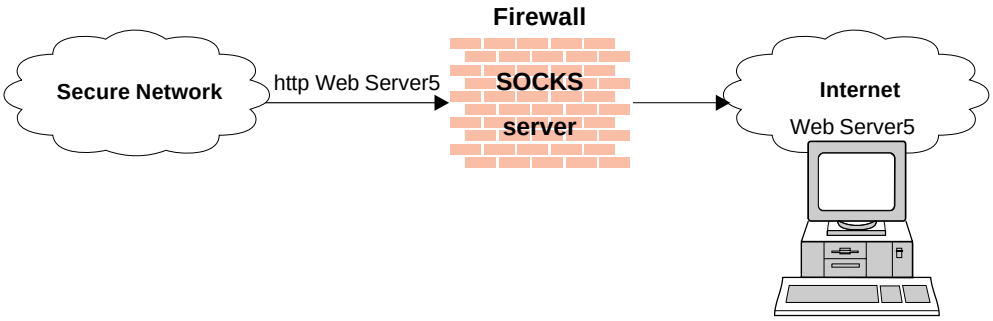

Figure 8. Firewall with a Socks Server

Socks needs to have special versions of the application code (called socksified clients) and a separate set of configuration profiles on the firewall. However, the server machine does not need modification; it is unaware that the session is being relayed by sockd.

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. Socks is not intended to handle inbound sessions, because it does not provide for secure password delivery and the user ID checking could possibly be subverted by an intruder.

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.

Configuration is performed through the configuration client. You have to specify the following:

- Secure domain name
- Secure domain name server
- Nonsecure domain name server

This generates all the necessary configuration files for the firewall domain name server. It is still necessary to configure the name server in the secure network to use the firewall service.

The operation of DNS on the firewall relies on three features:

1. The forwarders function, so that the name server inside the secure network can receive information about hosts outside its domain from the firewall name server, but the reverse cannot happen.
2. The caching capability which allows the firewall name server to get name information from the nonsecure network without pre-definition.
3. The fact that name resolution requests can be directed to any name server, whether or not the host from which the request is coming is a name server itself. This allows the firewall to be able to resolve names inside the secure network without giving those names away to hosts in the nonsecure network.

Figure 9 shows how name resolution requests flow for different combinations of requester and node.

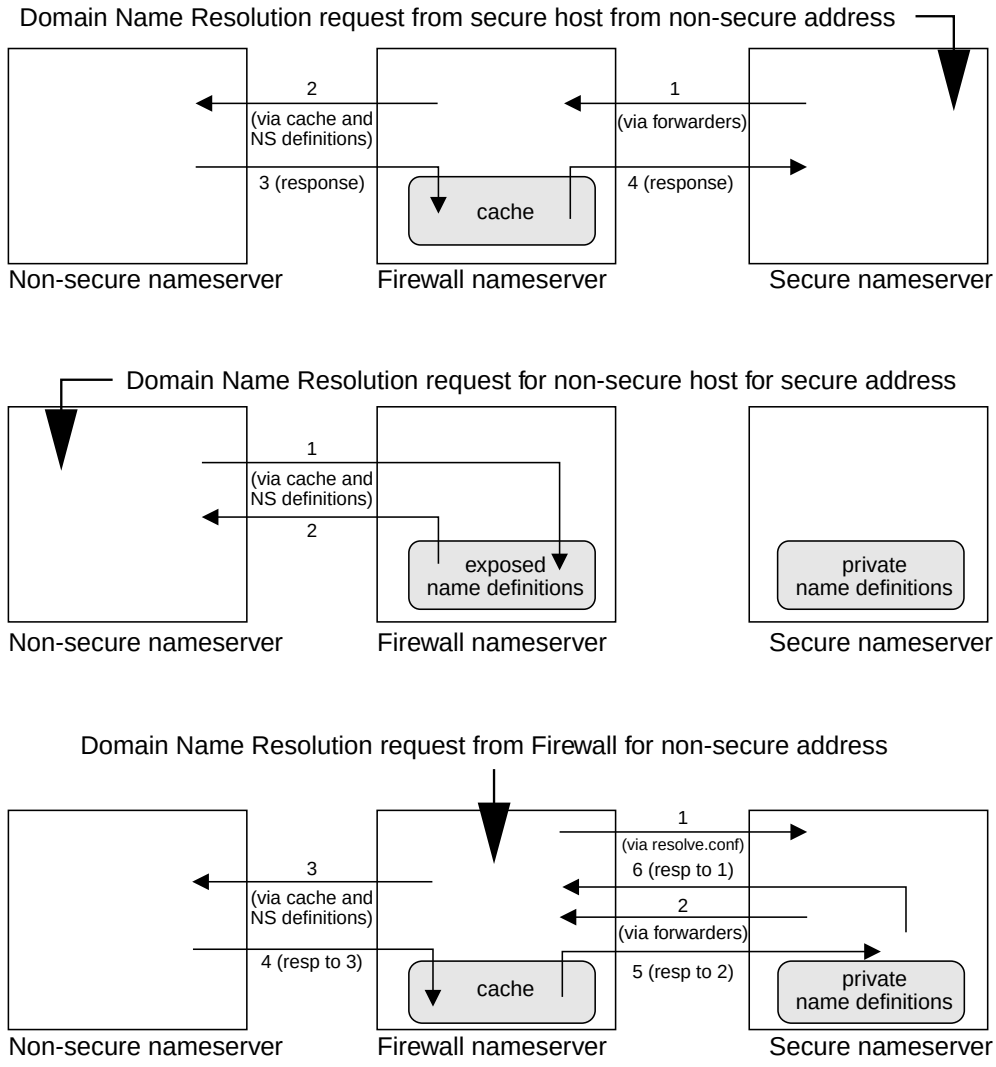

Figure 9. Name Resolution Flows

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.

Mail Handling

Mail is one of the primary reasons why an organization would want to access the Internet. The mail handler on the firewall will forward all incoming mail to centralized mail handlers and route mail from the hosts in the secure network. These centralized mail handlers are referred to as your SMTP gateways or secure mail servers. When you configure the firewall domain name service and mail gateway, you provide the configuration client with the host name, secure domain name, and public domain name for the secure mail servers.

The mail handler on the firewall will also forward all outgoing mail to destinations in the unsecure network.

Configure all users in the secure domain to send all mail to a secure mail server.

Configure the secure mail server to:

- Accept mail destined for `company.com`.
- Deliver incoming mail and all other mail destined for hosts in the secure network.
- Forward all outgoing mail to the firewall host. Mail not destined to your local domains should be relayed to the firewall, which in turn is delivered externally.

Virtual Private Network

A virtual private network (VPN) is two or more networks connected by one or more tunnels. A secure IP tunnel permits a private communication channel between two private networks over a public network such as the Internet. The two private networks are each protected by an IBM Firewall. The two IBM Firewalls establish a connection between them and they encrypt and authenticate (or both) traffic passing between the private networks. Secure IP tunnels can also exist between non-IBM Firewalls. Figure 10 illustrates a secure IP tunnel and a VPN.

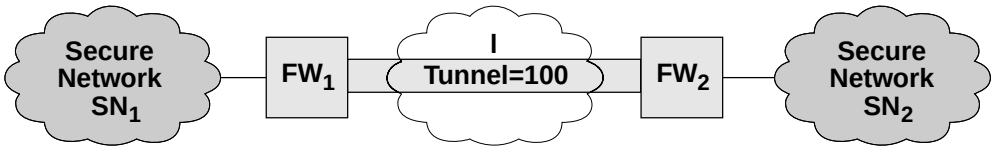

Figure 10. Tunnel, All IP Traffic between Two Secure Networks. FW_1 and FW_2 represent nonsecure interface IP address and mask. SN_1 and SN_2 represent any host in the secure network. The shaded area of the picture represents a VPN.

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.

Combining the Tools

You have now been introduced to the tools provided by the IBM Firewall. In reality, you might not need to use all of the tools available, but it is likely that you will need to use tools in combination. For example, most of the server functions need to be protected by a set of IP filters to prevent them being bypassed.

Chapter 2. Migration and Planning

You need to think about the layout of your network before you configure the IBM Firewall. Read the migration section before you get started and then use the checklist and the planning worksheets to help you with your network configuration.

Migration

If you are migrating from the Secured Network Gateway 2.2 to the IBM Firewall 3.1.1, consider that the filter files built primarily for the Secured Network Gateway 2.2 predefined filter rule sets will map very closely to the IBM Firewall 3.1.1 predefined services.

During installation, if a `sockd.conf` configuration file or a `fwfilters.cfg` file already exists from the previous release, use the `fwxmigrate` utility to generate network objects for the contents of these files.

If you are an existing Secured Network Gateway 2.2 sendmail user, you can set it up on the secure side of your network.

Planning Checklist

1. To enable DNS, install the AIX file set `bos.net.tcp.server`.
2. 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?
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. If socks usage is desired, obtain socksified clients. For information on using socks, see Chapter 9, “Configuring the Socks Server” on page 57.
5. If using Proxy support, what type of authentication is required?
   - If you are going to use the Security Dynamics ACE/Server** to authenticate users, install the ACE/Server client code at the firewall host. We suggest that you install the ACE/Server server code at some other 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 the AssureNet Pathways SecureNet Key card is to be used, purchase cards independently of the IBM Firewall.
   - If you use your own authentication method, see User Supplied Authentication in the *IBM Firewall Reference*
6. 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.

7. Decide on a method for archiving log files. This is an ideal candidate for cron job process. See Chapter 16, “Managing Log and Archive Files” on page 115.

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 Domain Name _____

Public Domain Name _____

Secure Domain Name _____
Registered IP addresses for NAT _____
_____

IP address of the configuration client _____

IP address of the remote client(s) _____

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 for the IBM Firewall.

Setting Up the Configuration Server

The configuration server is the configuration client's interface to the firewall. When you initially install the IBM Firewall, configuration can only occur from the local machine.

The configuration file for the configuration server is `/etc/security/rcsfile.cfg`. This file cannot be updated through the configuration client, SMIT, or the command line. It can only be updated by editing it directly. The defaults in this file are local=yes and encr=none. The options that can be specified in it are:

local= Indicates whether or not the configuration can occur on any machine or just from the local machine.

- **local=yes** The configuration can occur only on the local machine; this is the default.
- **local=no** The configuration can occur from any machine.

encr= Indicates what type of encryption should occur for data being sent to and from the remote configuration server.

- **encr=none** No encryption will occur; this is the default.
- **encr=ssl** SSL encryption will occur.

sslfile= Indicates the name of the keyfile to be used with SSL encryption; the default is `/etc/security/fwkey.kyr`. For information on how to create the keyfile, see the *IBM Firewall Reference.*

If the configuration server detects an error in `rcsfile.cfg`, it uses the default values for all parameters.

The configuration server listens on port 1014, which is the default, but it can be changed. To change the port number, modify the entry for `ibmfwrcs` in the `/etc/services` file and refresh the `inetd` daemon.

Setting Up the Configuration Client

When you install the IBM Firewall, the configuration client is automatically installed. The configuration client can also be separately installed on any AIX machine without the firewall.

To set the logon timeout value for faster/slower machines, change the parameter `name=TIMEOUT value=20`, where 20 equals the number of seconds to wait for a con-

nection to occur. Faster machines can be set to 10 and slower machines should accept the default value.

To increase the level of debug information in the JAVA console, change the parameter `name=DEBUG value=false`, where false equals no console logging and true equals console logging enabled.

To enable the Enterprise mode login panel, change the parameter `name=ENTERPRISE value=false`, where false equals normal login panel and true equals enable Enterprise login panel.

When the configuration client is started, you must first log on to the configuration server using a username defined on the machine where the configuration server is running.

Only user root and any usernames designated as firewall administrators that have the appropriate administration authentication can use the configuration client to log on to the configuration server.

After the firewall is installed, no usernames are designated as firewall administrators. So use the configuration client to log on to the configuration server using the root username and define the firewall administrator usernames. See Chapter 10, “Administering Users at the Firewall” on page 61 for information on how to define firewall administrators using the configuration client.

Log On to the Configuration Client

To log on to the configuration client (on the local or remote machine) for the firewall administrator:

- The user must be a firewall administrator
- The firewall administrator must have an administration authentication defined
- The user must have the authority to perform specific configuration functions

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:

- Is enabled for remote login.
- 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). This applies to logging on to the configuration server 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 of the user attributes can be set through the Modify User menu in the configuration client, SMIT, or the command line. User root will have all of the above fields set appropriately after installation of the Firewall. Refer to Chapter 10, “Administering Users at the Firewall” on page 61 for more information.

Chapter 4. Using the Configuration Client

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

When you first install the IBM Firewall, you can only run the configuration client from the firewall machine.

To launch the Netscape browser and start the configuration client applet, type **fwconfig** at the AIX command prompt.

A mouse is required to use the configuration client.

How to Log On to the Configuration Client

Log on to the configuration client is accomplished through a logon panel, as shown in Figure 11.

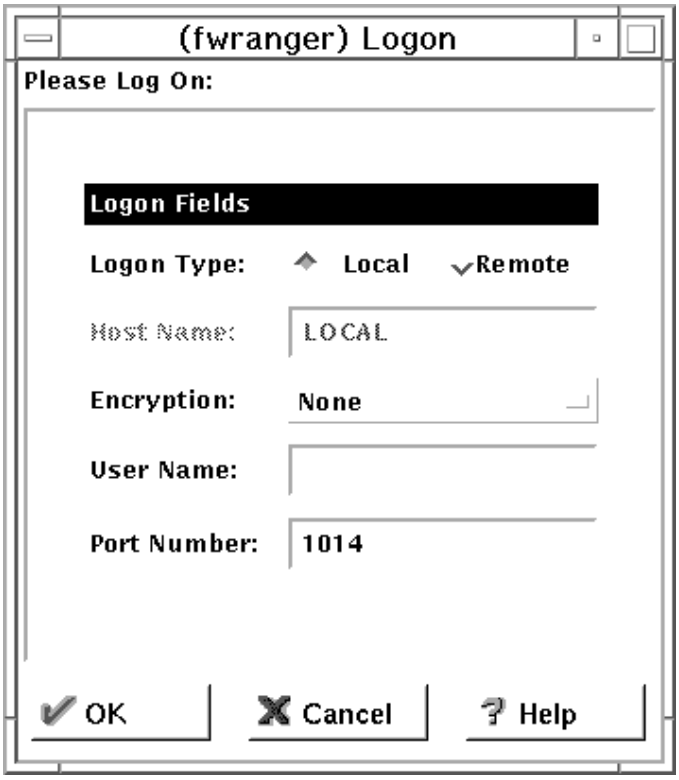

Figure 11. Configuration Client Logon Panel

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 configuration server. For the Client, the default for Local is None and the default for Remote is SSL.

4. Enter your username.
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 the firewall machine that you are logging on to. With host administration, the administrator can locally or remotely update one firewall at a time. Configuration files are updated directly on the firewall machine. Select Enterprise if you want to configure another firewall machine. With Enterprise Firewall Management (EFM) administration, the administrator is able to modify managed configuration information from the configuration client. With the exception of proxy files, configuration files for each firewall are stored on the central EFM administration server. These files can be transmitted to the managed firewall during subsequent download processing. For more information see Chapter 11, “Enterprise Firewall Management” on page 71.
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 username. 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 may be none. In this case you will immediately get the IBM Firewall configuration client panel.
8. After you have successfully been authenticated, you will see the main configuration panel that the administrator has authority to see.

The Navigation Tree

The configuration client has a collapsible tree-style navigation aid along the left side, as shown in Figure 12 on page 19.

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 12. Configuration Client Navigation Tree

General Features on the Main Panel

Above the Alerts Display you will see the following three buttons as shown in Figure 12:

Help — A **Help** button is located near the top of the configuration client main panel. Click **Help** to see what to do to get your IBM Firewall up and running.

User's Guide — A **Users Guide** button is located near the top of the configuration client main panel. Click **User's Guide** to see this softcopy publication.

Reference — A **Reference** button is located near the top of the configuration client main panel. Click **Reference** to see this softcopy publication.

Buttons that you will encounter on the main panel are:

Latest — A **Latest** button is located at the bottom of the configuration client main panel. Click **Latest** to see the most recent alerts.

Logoff/LogOn — A **Logoff/LogOn** button is located in the upper right-hand 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 close Netscape.

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

Previous A **Previous** button is located at the bottom of the configuration client main panel. Click **Previous** to see earlier alerts.

Other Features

Other features on the configuration client panels are listed here.

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

Buttons that you will encounter on many of the configuration client menus follow and some of them appear in Figure 13.

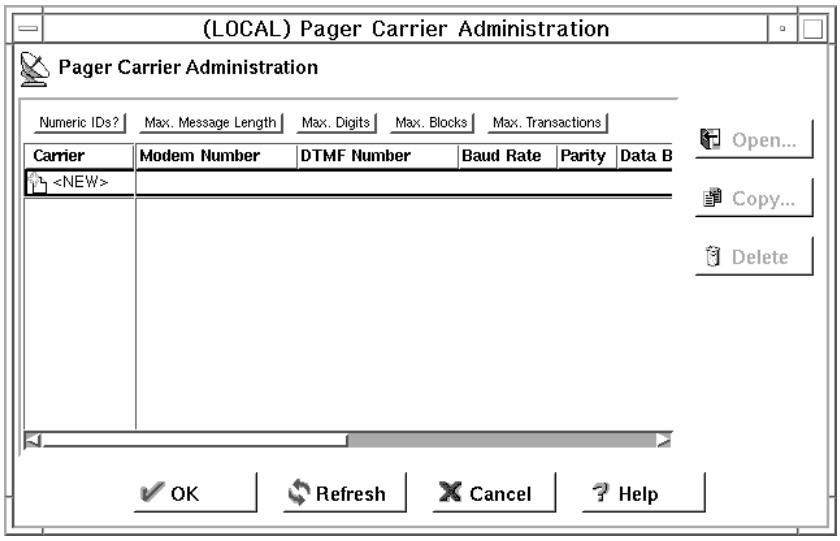

Figure 13. Other Features

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. Clicking **Copy** to create an item that is similar to the selected item will open a new item that 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

Depending upon what selections you have made earlier, some fields are not selectable.

Common fields that you will encounter on many of the configuration client menus 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 because these are used as SMIT and file delimiters. 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.

Anomalies

There are several unique features of the configuration client you need to be aware of.

If you hold down the left mouse button to proceed through a spin control and accidentally drag the mouse away without releasing the mouse button, the spin control continues. To stop it, click one of the spin control directional arrows with the left mouse button.

If you are selecting an item from a pull-down menu or list box, and you accidentally hit two mouse keys at once, the list box remains open. To alleviate the problem, click Cancel to exit the current panel and start over again.

In AIX, if you click above or below the elevator control in the scroll bar, you can only scroll one line at a time instead of page by page, when viewing lists for Users, Connections, and so forth.

The number keypad works even when Num Lock is off.

On the Log Facilities panel, if a log file is not highlighted and you scroll with the scroll bar, an outline appears every time you press the arrow. When you release the arrow, the outline disappears.

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 14.

Figure 14. The Alerts Display

The alert records displayed are obtained from the file identified by the first `local1` facility defined in the `/etc/syslog.conf` file. If no `local1` facility is defined you will see a blank display. See “Add Log Facilities” on page 115 for help in defining a `local1` 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.

The Log Viewer

Clicking Log Viewer brings up a log viewer window, as shown in Figure 15. The log viewer allows you to view Firewall log records. You can specify a log file and a record count (default is 25).

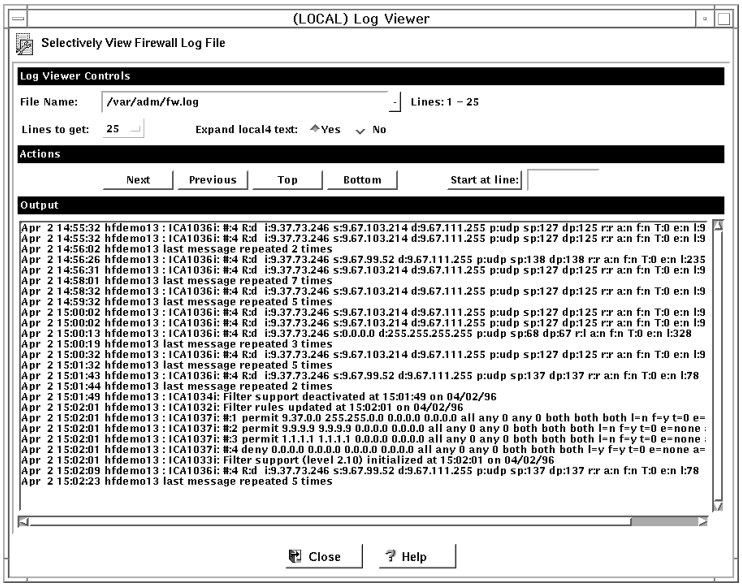

Figure 15. Log Viewer

The default log is the file identified by the first `local4` facility defined in `/etc/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.

See “Log File Creation and Archiving Using the Configuration Client” on page 115 and Chapter 17, “Monitoring the Firewall Logging” on page 121 for more information about log files, facilities, monitoring and alerts.

Chapter 5. Getting Started on the IBM Firewall

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

Basic Configuration Steps

For a basic IBM Firewall setup, perform the following steps:

1. Plan for your IBM Firewall setup. It is important to decide which functions of the firewall you want to use and how you want to use them as much in advance as possible. The following sections are helpful:

   - Chapter 1, “Introducing the IBM Firewall” on page 1
   - Chapter 2, “Migration and Planning” on page 11
   - “Planning Considerations” on page 43

2. Tell the firewall which of its interfaces are connected to secure networks. You must have a secure interface and a nonsecure interface in order to have your firewall work properly. From the configuration client navigation tree, open the System Administration folder and click Interfaces. You should 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 26 for more information.

3. Set up your general security policy. An easy way to do this is to access the Security Policy dialog, available inside of 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 27 for more information.

4. Set up your domain name service and mail service. Access these functions from inside the System Administration folder on the configuration client navigation tree. First read “Handling Domain Name and Mail Services” on page 30.

5. Define key elements of your network(s) to the firewall. Do this by using the Network Objects function in the configuration client navigation tree. Network Objects are especially important for controlling traffic through the firewall. You should define the following key elements as network objects:

   - Secure Interface of the firewall
   - Nonsecure Interface of the firewall
   - Secure Network

   See “Network Objects” on page 28 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. It is important to remember that 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 7, “Examples of Services” on page 43 for information on how to set up connections that support certain services.

7. Set up firewall users. If you are going to allow users to use services such as Proxy Telnet or Proxy FTP, you need to define these users to the firewall. See Chapter 10, “Administering Users at the Firewall” on page 61 for more information.

Following these steps should help you to get a basic firewall configuraton up and running. The IBM Firewall provides other functions, such as system logs to help you ensure the security of your network. See Chapter 16, “Managing Log and Archive Files” on page 115 for more information.

Designating Your Network Interface

Throughout this book we make the distinction between the secure and nonsecure interfaces, networks, and hosts. Secure interfaces connect the IBM Firewall host to the network of hosts in your own internal network, the network that you want to protect. **You must have a secure interface for your firewall to work.** Nonsecure interfaces connect the IBM Firewall to the outside network or networks, or to the Internet itself.

The IBM Firewall must have a minimum of two network interfaces. At least one interface connects to the internal, secure network, and at least one interface connects out to the Internet or to some untrusted, nonsecure network.

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

Select System Administration from the configuration client navigation tree. Double-click the file folder icon to expand the view. Select Interfaces. The addresses of the network interfaces are displayed. All known interfaces (adapters) will be shown and identified as secure or nonsecure. At this point you can add or delete a secure network interface.

To identify a network interface as a secure network interface:

1. Select an interface and click Change.

   If you have selected the only nonsecure interface and want to change it to secure, you will get a Change Warning menu. Click Yes to make the change or No if you do not want to make the change.

   There should be at least two interface addresses, one that connects to a nonsecure network and one that connects to your internal, secure network.

2. Repeat as necessary.
3. Click Close.

To identify a network interface as a nonsecure interface:

1. Select the network interface that you want to specify as no longer secure and click Change.
2. Select Close.

Using the Configuration Client to Define a Security Policy

One of the first things to consider when configuring the IBM Firewall is to determine the general security policy for your installation. The IBM Firewall provides a dialog to assist you in setting up your security policy, as shown in Figure 16.

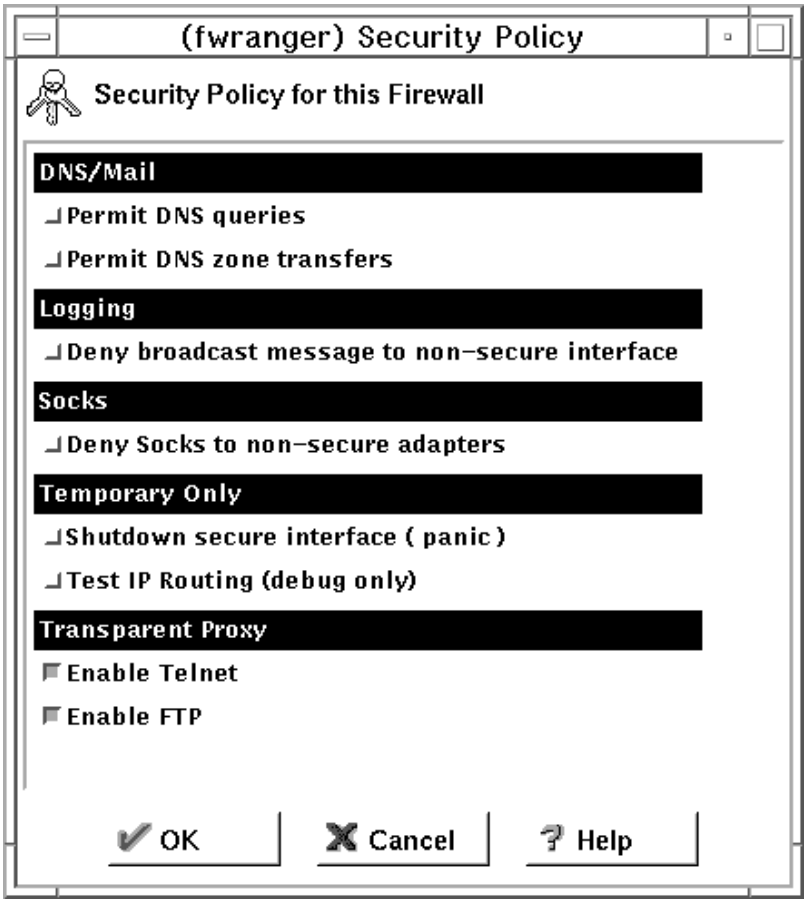

Figure 16. Security Policy

The Security Policy provides a quick and easy way for administrators to set blanket policies for the firewall . Most of the checkboxes 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.

The items on this menu enable services and makes them globally available.

Note that anytime you select a checkbox 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.

You are presented with the following list of checkboxes 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.
- Select Permit DNS zone transfers to allow Domain Name Service data files to be transferred from nameserver to nameserver.
- Select Deny broadcast message to nonsecure interfaces to deny 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 adapters to disallow socks traffic to enter the firewall from the nonsecure network.
- 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 checkbox, you must save it by clicking OK and Activate it through the Connection Activation window. **Warning: Use of this Service can cause security exposures for your firewall. Use it with extreme caution.**
- Select Enable Telnet to allow Transparent Proxy Telnets. If you change the value of this checkbox, you do not need to activate through the Connection Activation window. Just click OK for the changes to take place.
- Select Enable FTP to allow Transparent Proxy FTPs. If you change the value of this checkbox, you do not need to activate through the Connection Activation window. Just click OK for the changes to take place.

Network Objects

Network objects are representations of objects that already exist in your network such as hosts, networks, routers, virtual private networks, or users. You use them to designate source and destination addresses of 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, which can be found in “Network Configuration Planning Worksheet” on page 12, you are ready to build objects.

During installation, if a `sockd.conf` configuration file or a `fwfilters.cfg` file already exists from the previous release, use the `fwxmigrate` utility to generate network objects for the contents of these files.

You can create single or group objects. All network objects, except type User, require the attributes of IP address and subnet mask. If you select User as the object type, then you must select a valid user name.

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 menu appears. Double-click NEW. The Add a Network Object menu appears, as shown in Figure 17.

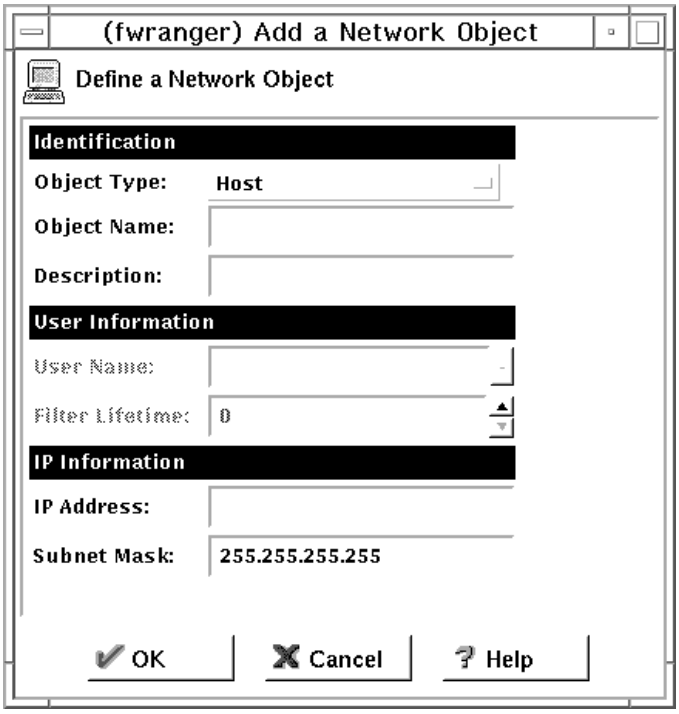

Figure 17. Add a Network Object

1. Select the object type from the pull-down menu. Object types are:
   - Host - a particular node on your network. It generally will have a mask of 255.255.255.255.
   - Network - a collective range of network addresses that is characterized by an address and subnet mask.
   - Firewall - a single machine with a firewall installed on it. Only an object of type firewall, can be the target of an IBM or a manual tunnel.
   - Router - a host which routes traffic between two or more networks.
   - Interface - a network adapter on a machine. It does not have to be an adapter on the Firewall.
   - VPN - Virtual Private Network (VPN) is the network on the other side of a tunnel.
   - User - This is a secure remote client user. See Chapter 14, “Using the Windows 95 Secure Remote Client” on page 103 for more information.
2. Fill in the object name.
3. Fill in the description. This field is optional.
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.

6. For object type User, fill in the user name and timeout fields. The user name replaces the IP address and subnet mask fields. The timeout field is used to remove unused filter rules after the secure remote client has disconnected ungracefully. See Chapter 14, “Using the Windows 95 Secure Remote Client” on page 103 for more information.
7. Click OK.

Network Object Groups

A group represents a collection of network objects. They are used as a convenience in setting up connections and can eliminate repetitive work. An example would be to group some users, individually represented by network objects, together 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 menu appears. Double-click NEW GROUP. The Add a Network Object menu 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
  George 191.1.10.1
  Susan  191.1.10.3
  Helen  191.1.10.5
  Peter  191.1.10.7
  John   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.

Handling Domain Name and Mail Services

This section explains how to configure Domain Name Service (DNS) in relation to the IBM Firewall. It also explains what you need to consider when configuring your mail service using the configuration client.

DNS in Relation to the IBM Firewall

The goal for DNS is to provide full domain name service to hosts inside the secure network while providing minimal information to hosts outside the secure network. Three domain name servers are required to accomplish this:

- One at the IBM Firewall
- One inside the secure network
- One outside the secure network

Refer to Figure 18 on page 31 to see how DNS works with the IBM Firewall.

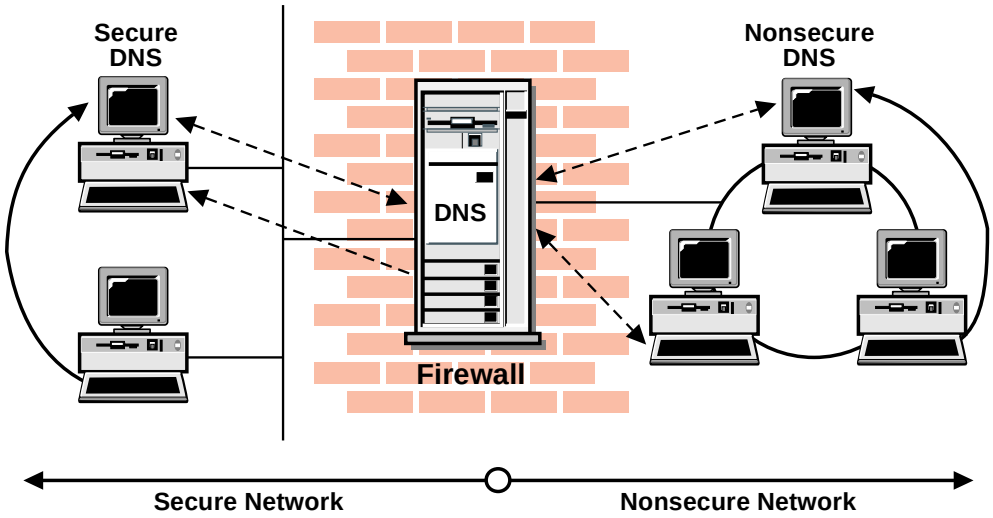

Figure 18. DNS

When Domain Name Services (DNS) is configured on the firewall, the following files are created:

- /etc/fwnamed.boot - boot file
- /etc/fwnamed.loc - reversion resolution data file for localhost
- /etc/fwnamed.ca - cache file

As a result, the firewall becomes a caching-only name server. If DNS is not provided by your Internet Service Provider (ISP) for your external domain, you might have to create additional data files on the firewall.

Configuring DNS Using the Configuration Client

To configure domain name services, 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 domain name services configuration.

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

1. The **Secure Domain Name** field is the domain name of your internal, secure network. Enter the name of the secure network domain.
2. The **Secure Domain Name Server** field contains the IP addresses of domain name servers that are inside, and provide name resolution only for, your internal network. Enter the IP address of the secure network domain name server.
3. The **Nonsecure Domain Name Server** field contains the IP addresses of the domain name servers outside your secured network that are provided by your Internet service provider. These are the domain name servers to which your internal domain name servers go for authoritative name resolution outside your own network. Enter the IP address of the parent zone domain name server.

DNS on the IBM Firewall

When configuring domain name service on th firewall, consider the secure domain name, which is the name of the internal domain that is protected from the Internet by the IBM Firewall.

When configuring domain name servers on the firewall, consider the following two servers:

1. The name server that resolves names and IP addresses for the hosts protected from the Internet by the IBM Firewall. We will refer to this as the secure name server.
2. The name server that resolves names and IP addresses outside of your organization or company. The address of this name server is usually supplied by your Internet service provider.

Note: Hostname on the firewall must not be fully qualified.

Secure Name Server

If you have a secure name server, you must have the following information in the data file for the secure domain.

```
firewall        99999999  IN   A   sss.sss.sss.sss
mail            99999999  IN   A   mmm.mmm.mmm.mmm
```

`sss.sss.sss.sss` represents the IP address of the secure interface on your firewall. `mmm.mmm.mmm.mmm` represents the IP address of the internal central mail server.

To forward requests that cannot be resolved locally to the firewall and to the Internet, add the following to the boot file on your secure name server:

```
forwarders   sss.sss.sss.sss
cache        .               /etc/named.cache
```

Also /etc/named.cache should be:

```
. 99999999 IN NS firewall.secure.company.com.
firewall.secure.company.com.  99999999 IN A sss.sss.sss.sss
```

No Secure Name Server

If you do not have a secure name server, you must modify the files created by the IBM Firewall.

1. Point to the firewall as the secure name server.

   For an AIX client this is accomplished by adding the following information to the /etc/resolv.conf file:

   ```
   domain secure.company.com
   nameserver sss.sss.sss.sss
   ```

2. Create a reverse name resolution file on the firewall to associate the IP address of secure interface with a name.

   Reverse name resolution takes an IP address and returns a host name. For example, if your internal network is the class C network 199.2.3 and your firewall is named firewall.secure.company.com at 199.2.3.27, then you would make the following entry in your /etc/fwnamed.boot file:

   ```
   primary 3.2.199.in-addr.arpa       /etc/fwnamed.rev
   ```

Create the /etc/fwnamed.rev file by copying the /etc/fwnamed.loc file and modfying it to look like the following except DO NOT split the first two lines as shown in this example:

```
3.2.199 in-addr.arpa IN SOA firewall.secure.company.com.
root.firewall.secure.company.com. (
                        9        ; Serial
                        86400    ; Refresh after 1 day
                        300      ; Retry after 5 minutes
                        654000   ; Expire after 1 week
                        3600  )  ; Minimum TTL of 1 day
```

Tips If You Don't Have DNS

If your ISP does not provide DNS and you want to receive mail through your firewall or you have hosts, such as web servers, on the nonsecure side of your firewall, you must create data files for name and address resolution. To create these data files:

1. Add entries for the data files in /etc/fwnamed.boot as follows:

   ```
   primary 2.1.204.in-addr.arpa /etc/fwnamed.rev
   primary company.com /etc.named.data
   ```

   company.com represents your external domain name. 2.1.204 represents the reverse of an address such as 204.1.2.1. of the secure interface on your firewall.

2. Create the domain data file fwnamed.data as follows:

   ```
            @ IN SOA firewall.company.com. root.firewall.com. (
                        9        ; Serial
                        86400    ; Refresh after 1 day
                        300      ; Retry after 5 minutes
                        654000   ; Expire after 1 week
                        3600  )  ; Minimum TTL of 1 day
   company.com.IN NS firewall.company.com.
      firewall IN A 204.1.2.3
   ```

3. Create the /etc/fwnamed.rev file by copying the /etc/fwnamed.loc file and modfying it to look like the following:

   ```
   @ IN SOA firewall.company.com. root.firewall.company.com. (
                           9        ; Serial
                           86400    ; Refresh after 1 day
                           300      ; Retry after 5 minutes
                           654000   ; Expire after 1 week
                           3600  )  ; Minimum TTL of 1 day
     IN NS firewall.company.com.
   1 IN PTR firewall.company.com.
   2 IN PTR www.company.com.
   ```

4. Refresh the name daemon using the refresh -s named command.

5. Go to your parent name server for the two domain names that you created and add name server records for your domain.

If you decide to create a data file on the IBM firewall, the name server will respond to DNS queries from the external network with information only from this data file. If you create this file it should have:

- An address record that associates a name with the nonsecure interface's IP address.
- A mail exchanger record that identifies the nonsecure interface as the place to connect when sending mail to the organization.

DNS requests from the internal network are sent to the name server inside the secure network. If the query cannot be answered by the internal name server, the query is forwarded to the firewall. If the query cannot be answered by the firewall name server, the query is forwarded to a name server on the external network. Once the query is resolved, the answer is returned to the originating name server which sends the answer to the originator of the query.

Mail Service

The mail handler on the firewall will forward all incoming mail to centralized mail handlers. These centralized mail handlers are referred to as your SMTP gateways or secure mail servers. When you configure the secure mail server (as described in "Configuring Secure Mail Servers Using the Configuration Client"), you provide the configuration client with the host name for the secure mail servers.

The mail handler on the firewall will also forward all outgoing mail to destinations in the outside network.

Configure all users in the internal network to send all mail to a secure mail server.

Configure the secure mail server to:

- Accept mail destined for hosts in your internal network.
- Deliver incoming mail and all other mail destined for a host in the internal network.
- Forward all outgoing mail to the firewall host. Mail not destined to a host in your internal network should be relayed to the firewall, which in turn delivers it externally.

The SMTP commands, EXPN and VERIFY, are not available.

Configuring Secure Mail Servers Using the Configuration Client

To configure Secure Mail Servers, select System Administration from the configuration client navigation tree. Double-click the file folder icon to expand the view. Select Secure Mail Server. The IBM Firewall displays the Secure Mail Server menu.

1. To add a mail server, select New and click Open. The Add Mail Server menu appears.
2. The **Secure Domain Name** field is the domain name of your internal, secure network. Enter the name of the secure network domain.
3. The **Secure Mail Server Name** field contains the host name of a mail server inside your secure network that handles all e-mail between the external network and the users inside your secure network. Enter the host name or IP address of the secure network mail server (optional).

   Note: If you do not configure a secure mail server, you see this text:

   ```
   Network Mail Server has not been configured
   ```

4. The **Public Domain Name** field is the domain name that represents the next highest level (the parent zone) of host naming, to which your internal network belongs. For example, if your network is `michu.jonesvl.edu`, your parent zone is `jonesvl.edu`. Enter the name of the parent zone domain.
5. Click OK.

Change a Mail Configuration Entry

Follow this procedure to change the mail configuration:

1. Select System Administration from the configuration client navigation tree. Double-click the file folder icon to expand the view. Select Secure Mail Servers. The IBM Firewall displays the Secure Mail Servers menu.
2. Select an entry in the list and click Open. The Change Mail Server Configuration menu appears.
3. The Secure Domain Name field is disabled, but you can change the other fields, as described in “Configuring Secure Mail Servers Using the Configuration Client” on page 34.

   Notes:

   a. If you previously configured a secure mail server and you specify a secure mail server here, this mail server replaces the one you configured earlier.

   b. If you have NOT previously configured a secure mail server and you specify a secure mail server here, this mail server is added to the configuration.
4. Click OK.

Delete a Mail Configuration Entry

1. Select System Administration from the configuration client navigation tree. Double-click the file folder icon to expand the view. Select Secure Mail Servers. The IBM Firewall displays the Secure Mail Servers menu.
2. Select an entry in the list and click Delete. You will get a delete warning.
3. Click OK to delete or Cancel if you change your mind.

Hints for Configuring Mail

Select Help for Mail Server Configuration. Read and follow the procedure for helpful information on how to configure Mail.

Chapter 6. Controlling Traffic Through the Firewall

This chapter tells you how to use the configuration client to configure filters. IP filters are tools to filter packets at the session level, based on IP address, direction, and TCP or UDP port. 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 subnetwork. One interface is declared nonsecure and the other declared secure. The filter acts between these two interfaces.

IP filtering provides the basic protection mechanism for the firewall, allowing you to determine what traffic passes across it based on IP session end details, thereby protecting the secure network from outsiders that use unsophisticated techniques (such as scanning for secure servers) or even more sophisticated techniques (such as IP address spoofing). You should think of the filtering facility as the base on which the other tools are constructed. It provides the infrastructure in which they operate and denies access to all but the determined intruder.

For the Secured Network Gateway V2R2, you built a filter rules base yourself and you edited the filters configuration file. The configuration of filters no longer works that way. You do not build a filter rules base with an editor. You use the configuration client to create network objects, rules templates, services and connections.

Using the Configuration Client to Build Connections

You use the components of the configuration client illustrated in Figure 19 on page 38 to build connections.

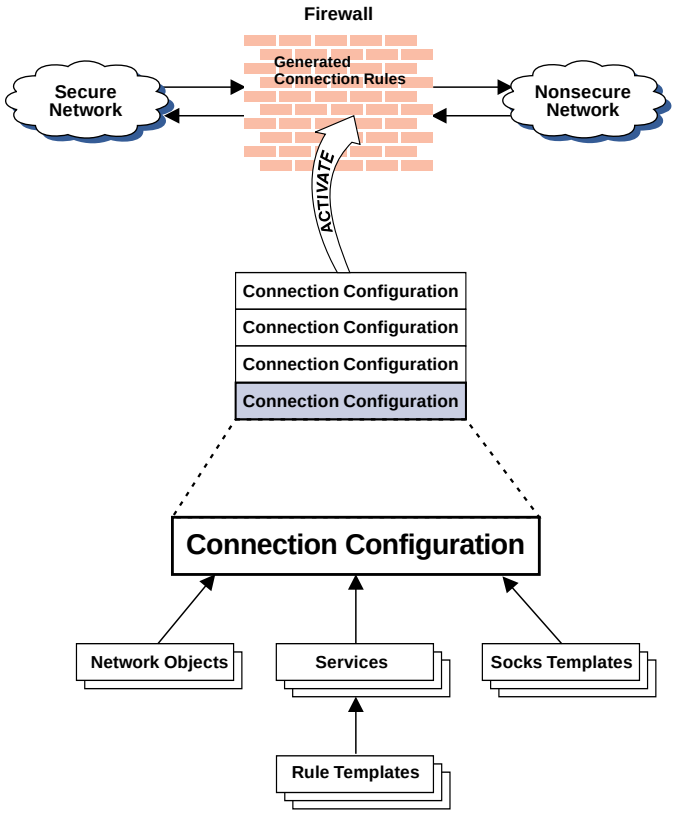

Figure 19. Building Connections

Network Objects Represent the various hosts, users, and subnets which interact with the firewall. Network objects can be grouped.

Network Object Groups
Represent a collection of 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.

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 or deny IP packets based upon their various attributes.

Services Are groups of rule templates which together instruct the firewall to permit or deny some meaningful type of communication. For example, one of the FTP services is comprised of eight rule templates. The IBM Firewall comes preloaded with a default set of 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 53 for more information.

Connections Associate network objects with services and/or socks templates to define the types of communications allowed between endpoints.

Building Connections Using Predefined Services

The Connection function allows you to control the type of network traffic that can take place between two network entities that are connected through the firewall. They permit or deny specified types of communications between two named network objects that serve as endpoints.

After you have defined your network objects, you create connections. In building connections, you will 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 menu appears. Select NEW and click Open. The Add a Connection menu appears, as shown in Figure 20.

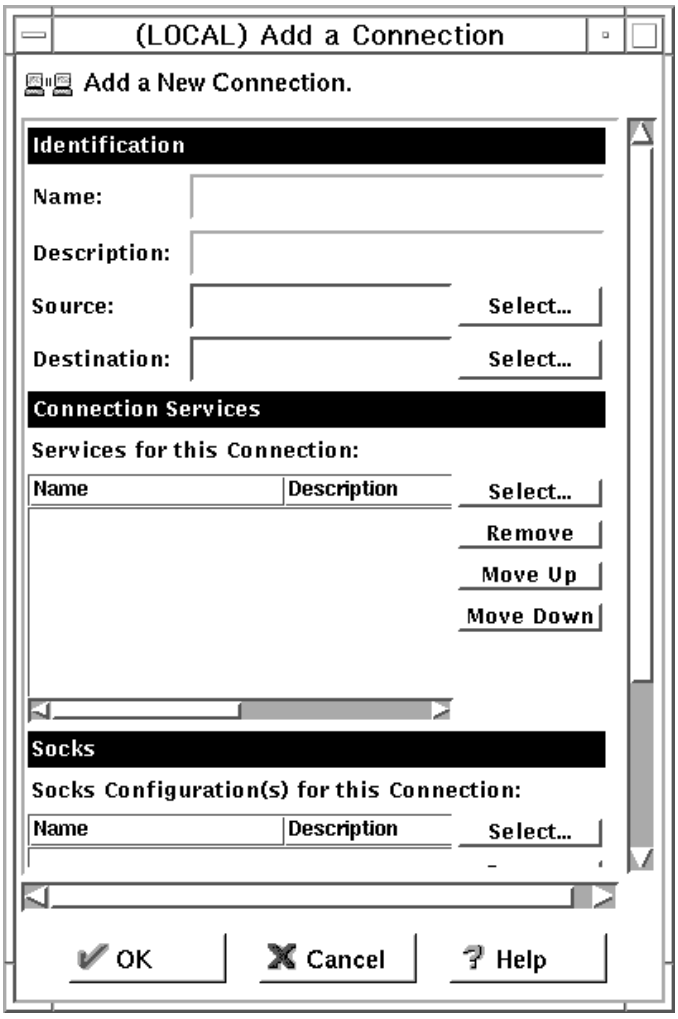

Figure 20. Add a Connection

1. Fill in the name of 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 menu list.
4. For the destination field, click Select and choose a network object from the Network Object menu list.
5. From the services list, click Select to 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 highlighting a service and clicking Move Up or Move Down.
8. You can remove a service by highlighting it and clicking Remove.
9. Use Socks Configuration for this Connection area of the panel, and follow the same procedure to make connections for Socks.
10. After you have everything defined, click OK.
11. Activate all of your connections.

Ordering Connections

You can order connections. When a datagram is received at a network interface, whether going into or out of the firewall host, these rules are applied starting at the top of the generated connection rules. When the information from the datagram exactly matches the information in the rule, the action (permit or deny) is enacted. If the entire file is searched without a match, the request is automatically denied.

Tip: Connection order is significant. Place more specific rules closer to the top and less specific rules 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 A server` before the `Permit telnet for Dept A` connection. If you reverse the order of the connections, the deny connection will never get encountered.

Connection Activation

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

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

Regenerate and activate connection rules
Deactivate connection rules
List Current Connection Rules
Validate Rule Generation
Enable Connection Rules Logging
Disable Connection Rules Logging

The Connection Activation panel appears, as shown in Figure 21 on page 41.

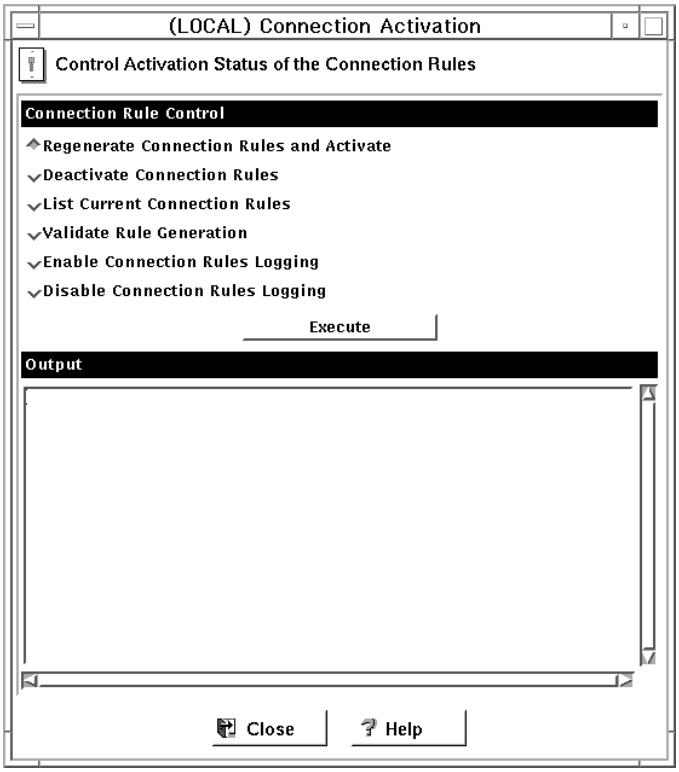

Figure 21. Connection Activation

After you make a selection, click Execute.

If you activate rules, the IBM Firewall builds the generated connection rules from the connection configuration and activates that rule set.

If you deactivate rules, the firewall is now protected by the default rules.

If you list the current rules active in the firewall, you see the most recent generated connection rules set. If you did a deactivate, these rules are not being used.

If you enable logging, the IBM Firewall logs selected traffic to the `local4` log.

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. Default filters are in place to protect your network from intrusion when you first install the firewall.

   Firewall Access:

   - The default filter configuration means that only local access is permitted from the secure and nonsecure interfaces.

- All other access is denied. That means that all routing is denied between the secure and nonsecure interfaces.

2. The configuration is active but has errors.

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

   Firewall Access:

   - All local access permitted only from secure interface to the firewall.
   - All other access is denied. That means that all routing is denied between the secure and nonsecure interfaces.

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 datagram 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 datagram, there is an implied "deny all" rule that denies access.

Chapter 7. Examples of Services

This chapter describes the steps needed to use the IBM Firewall configuration client 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, see Figure 19 on page 38. These connections define the types of communication which are allowed or prohibited between pairs of endpoints. Therefore, it is critical to plan your connection in terms of these endpoints.

As described in Chapter 6, “Controlling Traffic Through the Firewall” on page 37, 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, “Migration and Planning” on page 11 and create the network objects necessary to represent your network.

The examples in this chapter will 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 which are accessible through the firewall's secure interface. This could easily be a network object group that could contain several distinct domains, each of which is represented by its own network object.

The World This predefined network object will serve as our nonsecure network for the purposes of these examples.

Remote Firewall This object will represent a firewall which defends a network for with which we will be establishing a VPN tunnel.

Remote Host This object will represent a host inside a network defended by for the Remote Firewall. This host will be the target of our communication within the VPN.

Each desired type of communication must be envisioned in terms of the endpoint-to-endpoint communication involved. In this stage, consider whether your firewall will be proxying these communications 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; unless NAT is used, the secure host's IP address will be exposed to the network.

If you will use the firewall as a proxy, then the endpoints of your communication will include the firewall, as shown in, Figure 22 on page 44.

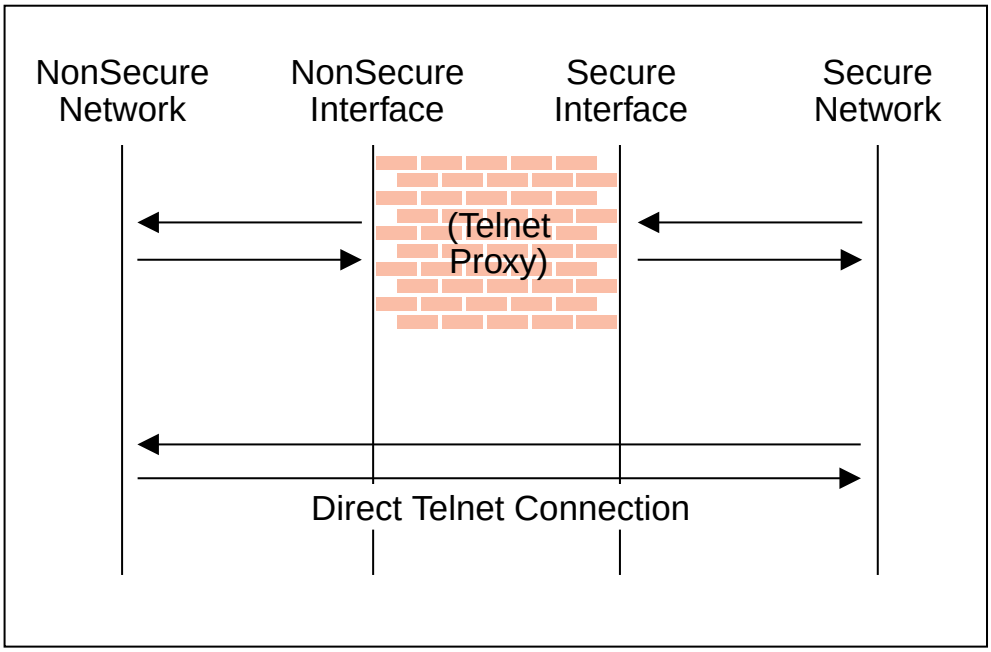

Figure 22. Telnet Proxy

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 22, 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
NonSecure Interface	The World	Telnet Proxy out 2/2

Example of Routed Telnet

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

Table 2. Routed Telnet

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

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

Example of Proxy HTTP

HTTP is one of the hottest protocols going these days. 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 routed 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 simple 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 23.

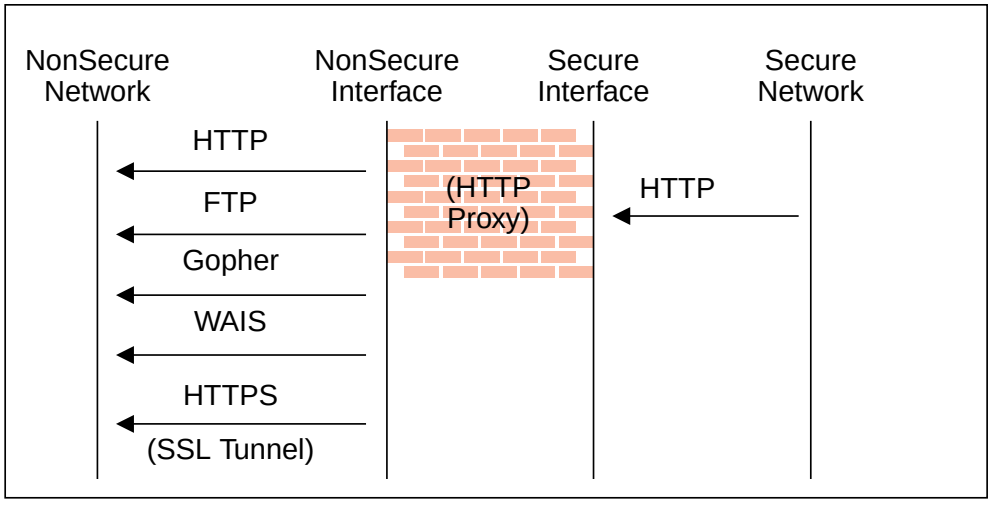

Figure 23. Proxy HTTP

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... • 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, “Using Proxy Servers” on page 87.

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 daemon and the client. Socks is more flexible than the HTTP proxy in that it can accommodate any TCP-oriented protocol and that the daemon may 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... • HTTP proxy out 2/2 • FTP proxy out 2/2 • Telent 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... • permit socksified HTTP • permit socksified FTP • permit sockisfied 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 9, “Configuring the Socks Server” on page 57.

Example of Virtual Private Networks

To establish a tunnel connecting Virtual Private Networks requires an even more 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 encrypted and 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 directed into its tunnel 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, the firewall now sends the encapsulated packet to the remote firewall directly, where the packet will be un-encapsulated and sent to its destination.

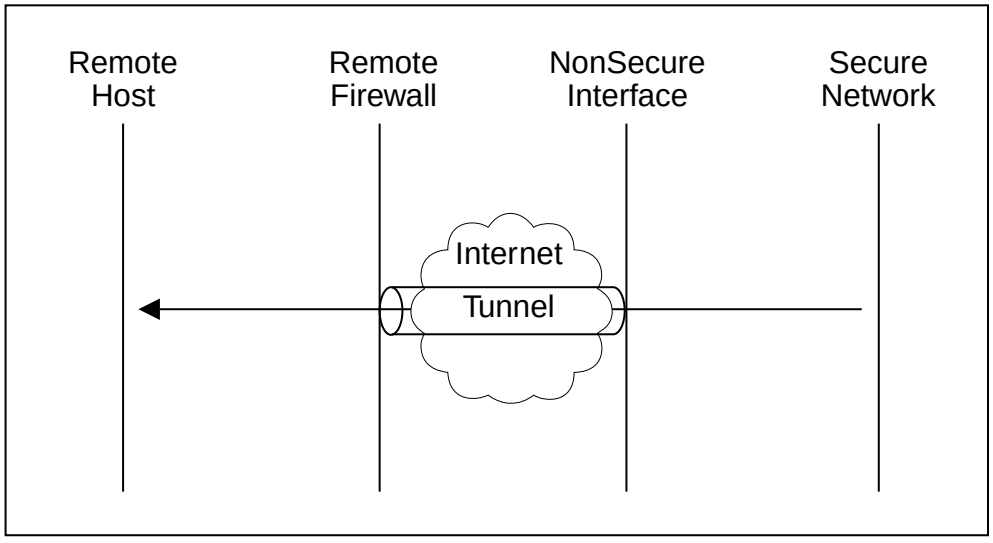

Figure 24. Virtual Private Networks

Such a configuration requires the following connections:

Table 5. Virtual Private Networks

Source Object	Destination Object	Services Required
Secure Network	Remote Host	• VPN traffic 1/2 • VPN traffic 2/2
NonSecure Interface	Remote Firewall	• VPN encapsulation • VPN key exchange (IBM tunnel only)

For more information on VPNs, see Chapter 13, “Creating a Virtual Private Network” on page 95.

Hints for DNS

Very little communication will take place efficiently if you do not provide DNS resolution. See “Handling Domain Name and Mail Services” on page 30 for details on configuring DNS, and do not forget to enable "Permit DNS Queries" in your Security Policy.

Hints for NonSecure Socks Clients

The Security Policy panel contains a checkbox for "Deny Socks to nonsecure interface". This service will reject any packets addressed to your socks daemon from any nonsecure interface, and will make your firewall much more secure.

Chapter 8. Customizing Traffic Control

This chapter helps you to define filter rules and services. You can tailor any predefined services to your particular needs or create new 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. You can also delete services. Socks services apply to socksified connections.

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 menu, double-click NEW.

   The IBM Firewall displays an Add IP Rule menu, as shown in Figure 25 so that you can define a rule.

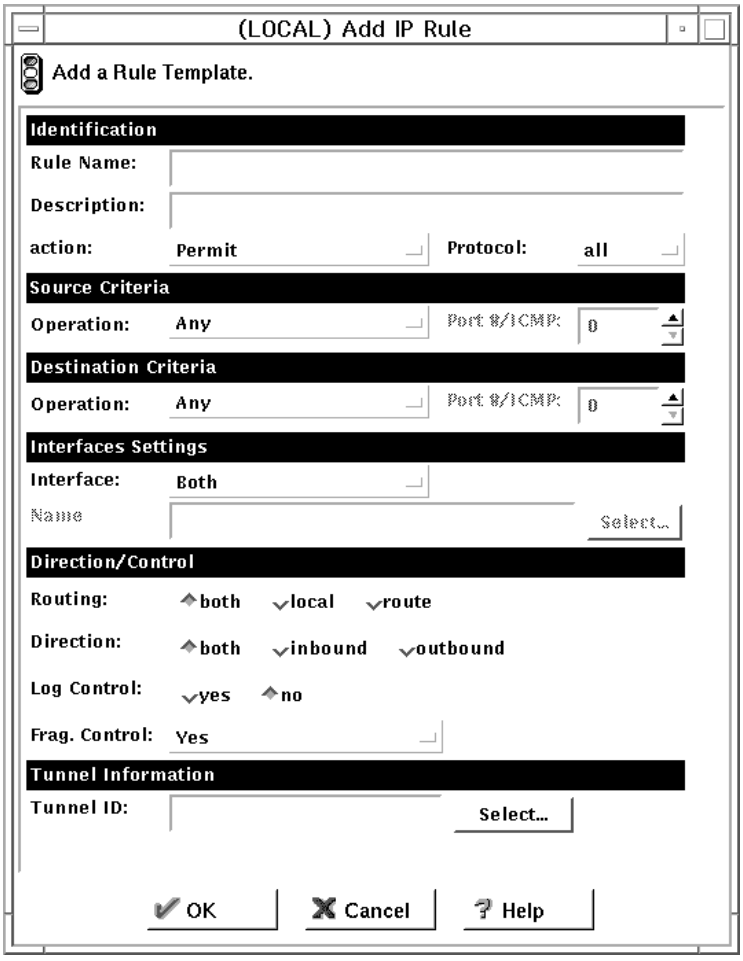

Figure 25. Add IP Rule

3. Enter the Rule Name.
4. Enter the Rule Description. This field is optional.
5. Open the pull-down menu and choose to either permit or deny access to the firewall.
6. To choose a protocol, open the pull-down menu and select from the following list:

   all Any protocol will match this rule.

   tcp The datagram protocol must be transmission control protocol (TCP) to match this rule.

   tcp/ack The datagram protocol must be TCP with acknowledgement to match this rule.

   udp The datagram protocol must be user datagram protocol (UDP) to match this rule.

   icmp The datagram protocol must be internet control message protocol (ICMP) to match this rule.

   ospf The datagram 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 datagram 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 datagram protocol must be encapsulating security protocol used by the tunnel for sending encapsulated IP packets to match this rule.

   ah Authentication header protocol is the datagram protocol used by the tunnel 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.
8. The logical operation and port number operands are used together. The source opcode and destination opcode operands are logical operations that state a relationship between the port number (destination or origin) for the datagram and the source port# and destination port# operands. For example, if the datagram destination port is port 20, and the destination opcode and destination port# operands are “ge 15,” the datagram matches because port 20 is greater than or equal to port 15.

   If you use a source opcode or destination opcode with the value **any**, the filter does not look at the port number; **any** means that 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 oper-

ator specified is applied to the type or code and, as for ports, an operator of any, means that any type and/or code value will match the rule. The port number cannot be changed in this case.

The values for source opcode, destination opcode, and icmp opcode 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:

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

In the rule values, note these constraints:

- The values for the port numbers must be in the range 0 through 65535.
- Your port number will never be less than 1 or greater than 65535.

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
5	1 - Redirect for host

9. Open the pull-down menu to select the type of interface (adapter) you want.

both	For datagrams coming or going on either the secure or the nonsecure interface
secure	For datagrams coming or going on the secure interface
nonsecure	For datagrams coming or going on the nonsecure interface
specific	Use with the interface name field when selecting an 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 datagrams going out from or into the selected interface

    inbound For datagrams coming into the selected interface from the network

    outbound For datagrams 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 gets recorded in the `local4` log with priority level `Error`. If this parameter is not specified, for permit, the default is No and for deny the default is yes.
14. Open the pull-down menu 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 you have a network object with firewall as the type, you can choose a tunnel ID. Click Select and choose a tunnel ID from the Select a Tunnel screen. Click Apply.

In order 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.

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

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 menu. The Modify IP Rule Configuration menu appears.
2. Modify the appropriate fields as described in Chapter 8, “Customizing Traffic Control” on page 49 and click OK to apply the changes.

Delete Rule Configuration Entry

To delete a rule template:

1. Select a rule from the Rules List and click Delete.
2. The configuration client asks you whether you really want to delete this entry.
3. Click OK to delete the rule entry. The configuration client deletes that entry from the filter configuration file.

Defining 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.

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 menu appears. Double click NEW to get the Add Service menu, as shown in Figure 26 on page 54.

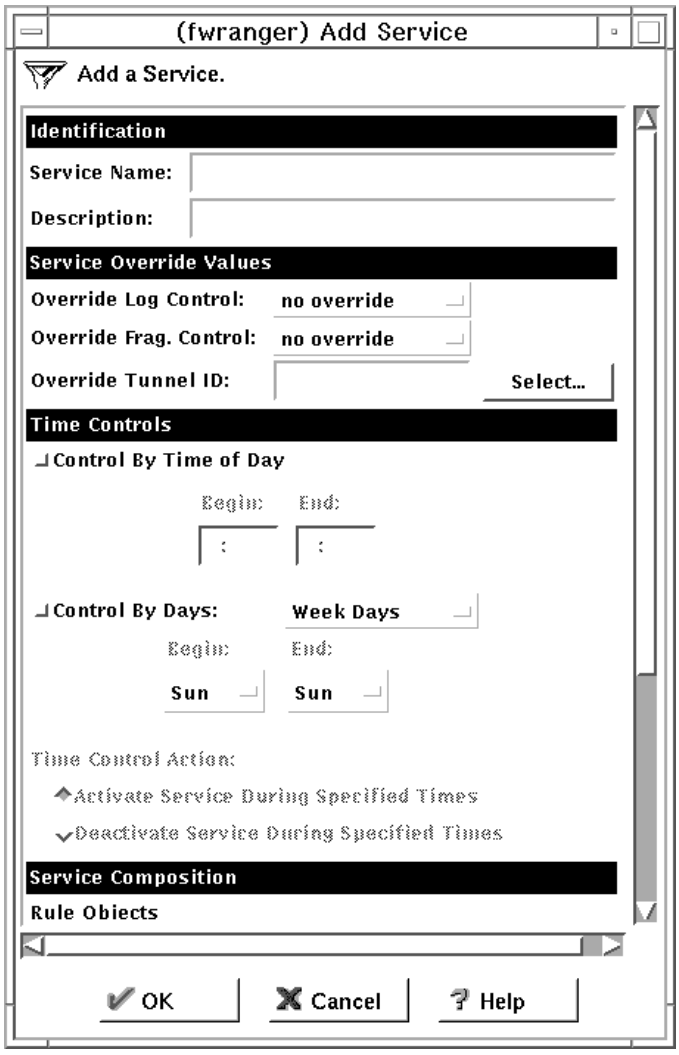

Figure 26. Add a Service

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.

   Note: If you are using the tunnel override field to override the tunnel for a predefined service only, then you must use the configuration client to eliminate the tunnel ID, if you later wish to delete this tunnel.
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 Click if you want this service to be activated or deactivated according to begin and end times during the day. Use a 24-hour military format. If this field is not enabled, the Time of Day fields will be in effect everyday.

   Control by Days Click if you want this service to be activated or deactivated according to a schedule based upon either week days 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 datagram is received, the IBM Firewall compares the information in the datagram to the rules in the rules configuration file starting at the top of the file (Rule 1). It stops comparing when the first exact match is found, and performs the action contained in the rule. Thus it is important to have the rules in order.

   Once you have added a series of rules to the service, you can change their order by moving selected rules after others. 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 9. Configuring the Socks Server

This chapter tells you how to use the configuration client to configure the Socks Server.

The socks server (`sockd`) is enabled, but there are no rules in the socks configuration file. For socks clients to use the Socks server, you must configure socks using the configuration client.

See “Example of Socks” on page 46 for an example of how to set up a socks service.

Configuring the Socks Server Using the Configuration Client

Socks templates are rules that control security through the socks server. The Socks Templates function allows you to customize, add to, copy, or delete existing Socks Templates. These socks templates, in turn, can be used in the definitions of Connections on the firewall similarly to the way Rules Templates are used.

Add a New Socks Rule

This section describes how 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 on the file folder icon to expand the view. Select Socks. You get the Socks menu.

1. Double-click NEW to add a new socks template.

   You get the Add a Socks Rule menu, as shown in Figure 27.

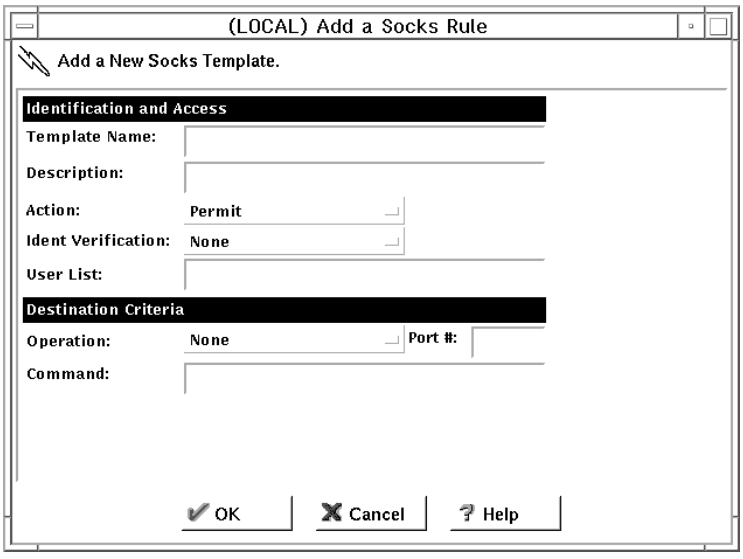

Figure 27. Add a Socks Rule

2. 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 SMIT and file delimiters. Use of these characters will result in unreliable data.

3. Fill in a description.
4. 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. Open the pull-down menu to choose either permit or deny access from a source to a destination.
5. Specify whether `identd` verification should be used for this rule. The `identd` server provides identification of users in your network. If you are using an `identd` server, use this field to indicate how you want the results of the identification to be used. The User List can also be used by adding IDs that will be matched.

   None Use the identification option selected in the `sockd` entry in `/etc/rc.tcpip`, if any are specified.

   ?=I The `identd` must be used to verify the user's identity. Access is denied if connection to client's `identd` fails or if the result does not match the user ID reported by the client program.

   ?=i `?=i` also specifies the use of `identd`, but denies access only if client's `identd` reports a user ID different from what the client program claims.

   ?=n Do not use the `identd` program. This overrides the setting on the socks entry in `rc.tcpip`.
6. In the User List field, you can enter a user ID, a list of user IDs, a file name, or a list of file names. If you enter a list, separate the entries with commas. File names must be fully qualified (including the leading "/"). 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 the following characters pipe symbol (|) double quote character(").
   - If file names are used, they must be fully qualified (with the leading "/" to prevent their being interpreted as user IDs). Each file can contain a list of user IDs, with one or more per line, separated by commas, and optionally including a comment that is delimited with the # character. Full comment lines - those that begin with the # character are also supported. Each line in the file can be up to 1023 characters long and must be terminated by a "newline" character.

Note that when SMIT constructs a rule consisting of user list data obtained from this field, it will accept an arbitrary number of blank characters or a comma as entry delimiters and will build a userlist entry consisting of a contiguous string of entries, separated by commas. This is done at rule creation time, not rule evaluation time. Do not rely on this behavior if you manually edit the configuration file and change the contents of a userlist. A rule created or changed manually to include imbedded spaces (or tabs) will cause that rule to be rejected as invalid.

7. In the Operation field, enter a logical operator code that represents 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 operator 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.

8. In the Port # field, enter the number of a port. The Port Number is used with the Operation field 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 this pair is omitted, the line applies to all destination port numbers.

9. In the Command field, enter a command string to be executed when the conditions in this rule are satisfied. The following substitutions occur before the string is presented to the Borne shell for execution:

   ```
   %A      replaced by the client host's domain name if
           known, by its IP address otherwise
   %a      replaced by the client host's IP address
   %c      replaced by connect or bind, the command
           sockd is asked to execute
   %p      replaced by the process id of sockd
   %S      replaced by the service name (for example, ftp)
           if known, by the destination port number other-
           wise
   %s      replaced by the destination port number
   %U      replaced by the user-id reported by identd
   %u      replaced by the user-id reported by the client
           program
   %Z      replaced by the destination host's domain name
           if known, by its IP address otherwise
   %z      replaced by the destination host's IP address
   %%      replaced by a single %
   ```

   You can string together several shell commands in a line with a | or ; symbol.

Each screen of information builds a PERMIT or DENY rule in the `/etc/sockd.conf` file that controls access through the Socks server. Use this Add a Socks Rule menu 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 menu.

   The Modify a Socks Rule menu appears.

2. Change the appropriate fields as described in “Add a New Socks Rule” on page 57, and click OK.

Delete a Socks Rule

Select an entry from the Socks menu 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 `local4` log file.

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 specific client that socks implements, refer to that client documentation. For additional information refer to:

http://www.raleigh.ibm.com/sng/sng-socks.html
http://www.socks.nec.com

Chapter 10. Administering Users at the Firewall

This chapter describes how to do the daily administrative tasks with the IBM Firewall. These tasks involve:

- 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
- Setting up the idle proxy environment.

The configuration client provides menus and online help to guide you through administering the IBM Firewall. 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 menus or command line.

Adding a User to the IBM Firewall

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 menu appears.
2. Select New from the User Administration menu and click Open. The Add User menu appears, as shown in Figure 28 on page 62.

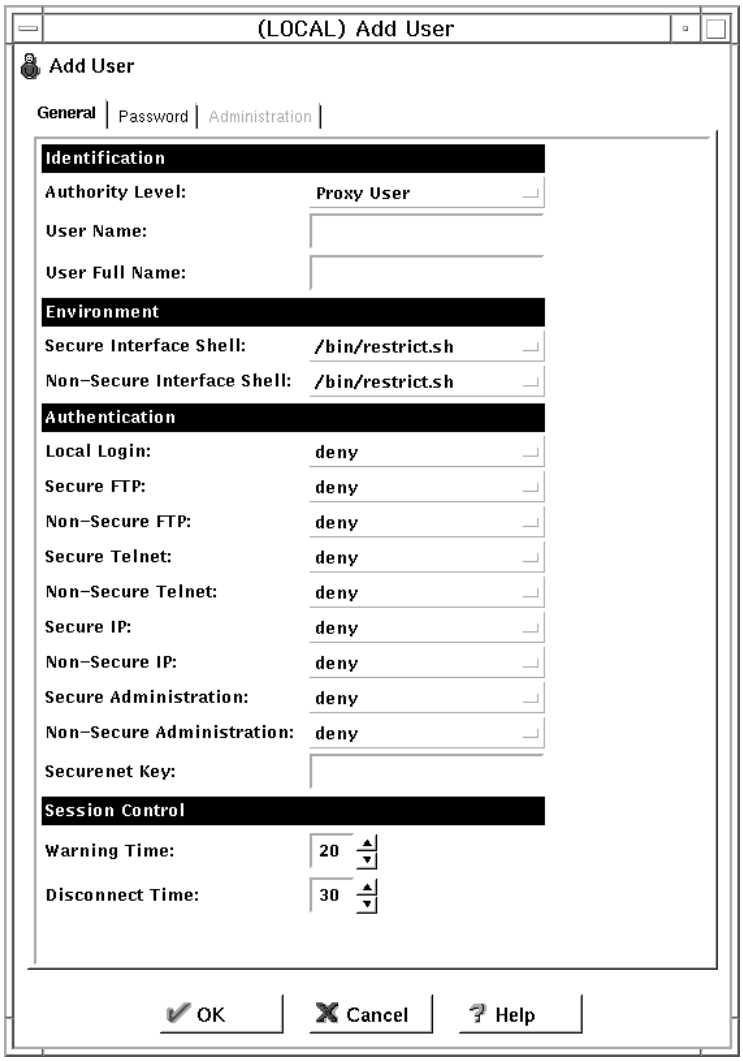

Figure 28. Add User

3. Provide this information:

Authority Level Specifies the authority level for this user. Users other than root can perform firewall administrative functions. Open the pull-down menu to select user type.

Proxy User Has no administration authority. It is the default.

Firewall Administrator

Has authority to administer the firewall. Note that only user root can create users with firewall administrator authority.

All firewall administrator actions are logged to the `local0` log facility. Only root has access to the `local0` facility through SMIT. Logged data includes administrator username, command executed,

arguments passed, and return code.

For more information, see “Administrator Authority Level by Function” on page 69.

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 username can consist of from 1 to 8 characters, including:

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

The firewall comes with two preinstalled users:

a. Default User Authentication, which is a user that is authenticated by whatever method has been specified for the default username `fwdfuser`. You can implement default user authentication for usernames that have not been authorized as proxy users. Any user authentication method can be called to validate these usernames, for example, the username can be authenticated by a remote server that has access to a centralized user ID database.

   At installation, when the `fwdfuser` is created, all authentication methods are set to deny. 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, `fwdfuser` cannot be used as a username for telnet or FTP requests. Password and none are not valid authentication types for `fwdfuser`. For more information, refer to the *IBM Firewall Reference.*

b. `fwdpuser` shows the default values of the various attributes for the Add User panel. Because of `fwdpuser`, the administrator can choose to have uniform attribute values for all users. The administrator does not have to retype all of the attribute values each time they add a new user. If the administrator changes the values of `fwdpuser`, any subsequently added users would display the changes reflected. `fwdpuser` cannot be deleted.

User Full Name Specifies a description of the user.

Secure Interface Shell
Specifies the shell program that will run when this user logs in from the network connected to the secure interface.

Open the pull-down to see alternative shell names. The choices are:

/bin/restrict.sh	The firewall restricted shell. This is the default.
/bin/csh	The C shell
/bin/ksh	The Korn shell
/bin/bsh	The Bourne shell
/bin/oneact.sh	A firewall shell that performs a single action and only allows telnet or ftp through the firewall.

Non-secure Interface Shell
Specifies the shell program that will run when this user logs in from the network connected to the nonsecure interface. Open the pull-down menu to see the alternate choices:

/bin/restrict.sh	The firewall restricted shell. This is the default.
/bin/csh	The C shell
/bin/ksh	The Korn shell
/bin/bsh	The Bourne shell
/bin/oneact.sh	A firewall shell that performs a single action and only allows telnet or ftp through the firewall.

The following fields refer to authentication methods:

Local Login Authorizes login from the console.

Secure FTP Specifies the level of authentication this user needs to use FTP to access the firewall from the secure network. Open the pull down to select from the list of choice. They are explained in “User Authentication Methods” on page 66.

Nonsecure FTP Specifies the level of authentication this user needs to use FTP to access the firewall from the nonsecure network. Open the pull-down menu to select from the list of choices. They are explained in “User Authentication Methods” on page 66.

Secure Telnet Indicates whether this user's identity, when logging in from the secure network, must be authenticated by some means. Open the pull down to select from a list of choices. They are explained in “User Authentication Methods” on page 66.

Nonsecure Telnet Indicates whether this user's identity, when logging in from the nonsecure network, must be authenticated by some means. Open the pull-down menu to select from a list of choices. They are explained in “User Authentication Methods” on page 66.

Secure IP Specifies the authentication method to be used for the secure remote client when logging in from the secure side. Open the pull-down menu to select from a list of choices. They are explained in “User Authentication Methods” on page 66.

Nonsecure IP Specifies the authentication method to be used for the secure remote client when logging in from the nonsecure side. Open the pull-down menu to select from a list of choices. They are explained in “User Authentication Methods” on page 66.

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. Open the pull-down menu to select from a list of choices. They are explained in “User Authentication Methods” on page 66.

Nonsecure Administration
Specifies the authentication method used to log on from the configuration client through a nonsecure interface. Open the pull-down menu to select from a list of choices. They are explained in “User Authentication Methods” on page 66.

SecureNet Key Specifies the character sequence to be entered by a remote user who has a AssureNet Pathways SecureNet Key card. Enter the key code with which you will also prime the key card. See your SecureNet Key information for instructions on selecting and installing a key code.

Notes:

a. This field is not used for the SecurID card.

b. You must create a unique random key for each user.

c. When you install the key in the SecureNet key card, use the AssureNet Pathways installation procedure and select **Mode 5**.

Warning Time The warning time is the maximum time in minutes that the user has remained idle before a warning message is issued to disconnect the user. See “Setting Up and Administering the Idle Proxy Environment” on page 70 for more information.

Disconnect Time The disconnect time is the maximum time in minutes that the user has remained idle before they are disconnected. The disconnect time must be greater than the warn time.

See “Setting Up and Administering the Idle Proxy Environment” on page 70 for more information.

User Authentication Methods

The choices for user authentication are:

Deny The user is denied access.

None No authentication is needed.

Password The user must be prompted for, and enter, a valid password. When this panel is complete, the IBM Firewall prompts you to specify a password for this new user. This password is used for both telnet and FTP requests.

Notes:

1. Passwords should be no longer than eight characters in length. Characters in excess of eight will be truncated.
2. 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.
3. AIX allows placing limits on passwords when changed by users. Password rules are:

   Login retries
   Number of days to warn the user before the password expires
   Number of passwords before reuse
   Weeks before password expiration
   Weeks before password lockout
   Maximum age of the password
   Minimum length of the password
   Minimum alphabetic characters
   Minimum other characters
   Maximum number of repeated characters
   Minimum number of different characters

Click the password tab to customize these values for each user, as shown in Figure 29 on page 67.

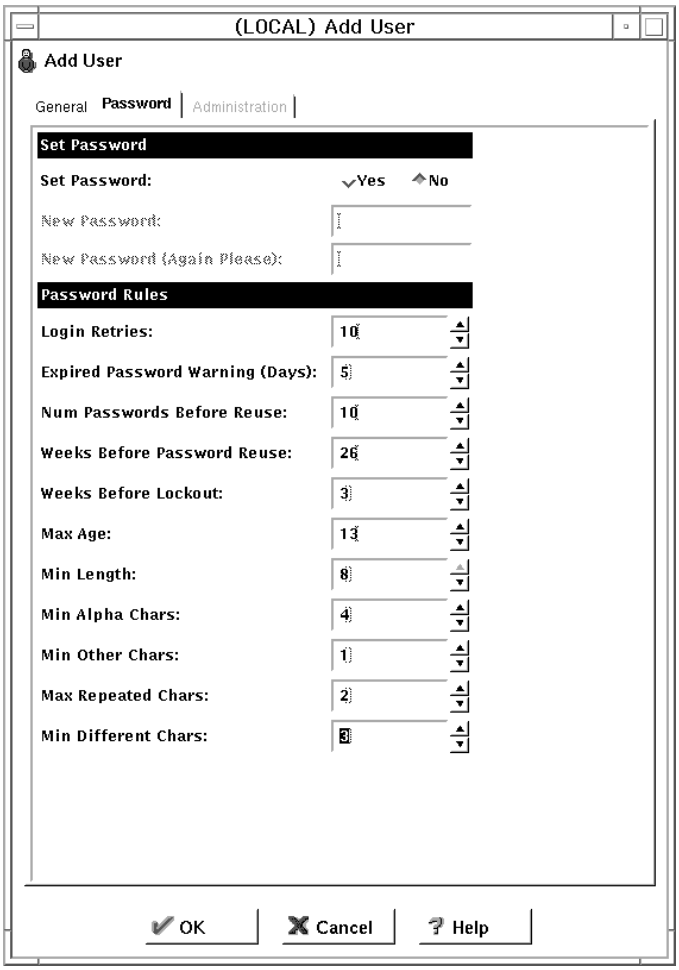

Figure 29. Password Tab

SecureNet Key

Authentication is done using an AssureNet Pathways SecureNet Key.

In the SecureNet Key field, enter the key code with which you also prime the SecureNet Key card.

Notes:

1. You must create a unique random key for each user.
2. The random key must be in the range 1–377 for each 8 octal values
3. When you install the key in the SecureNet key card, use the AssureNet Pathways installation procedure and select **Mode 5**.

SecurID Card

Authentication is done using a Security Dynamics SecurID security card or pinpad card. DO NOT use the SecureNet Key field. 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.

User-Supplied Authentication

Authentication is supplied by the user. You can only have one user-supplied authentication method on the firewall at any given time. For information on how to create and compile a subroutine for user supplied authentication, refer to the *IBM Firewall Reference*.

Changing a User's Access

After you add a user to the Firewall, you can change that user's security attributes.

1. Select the user you want to change from the Users menu and click Open.
2. When the Modify User menu appears, change the appropriate fields. See "Adding a User to the IBM Firewall" on page 61 for a list of user attributes that you can change.
3. When you have made the changes, click OK.

For security, be sure the **root** user's Nonsecure Telnet authentication level is set to deny. This will prevent anyone from logging in as root from the nonsecure interface and is the default.

Changing the User's Password

If you checked set Password on the Modify Username menu, follow this procedure.

1. Enter the password for the user.
2. Enter the new password again to confirm it.

   Note: 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.

Deleting a User from the IBM Firewall

Note: Do not Delete the user root, fwdfuser, or fwdpuser.

An IBM Firewall user is simply an AIX user with additional configuration definitions. Deleting a user from the firewall, deletes all of the additional configuration definitions relating to the firewall, and it also removes the user definition from the underlying AIX system.

The root user must remain as a firewall user as long as the IBM Firewall is installed on the system.

1. Select a user from the User Administration menu and click Delete. A confirmation box appears and you will be asked if you wish to delete the user.
2. Click OK to delete the user.

Administrator Authority Level by Function

If the configuration client is logged in as **root**, **root** and only **root** can create and modify administrators and determine which firewall functions they will have authority over. For example, you can limit a particular administartor to just having the authority to perform the Users and Log Monitor fucntions.

If an administrator copies user **root**, to create a new administrator, the new admistrator maintains most of root's attributes except that remote logins are enabled. The new administrator will not have root authority over the AIX system in general.

On the Add User menu, select Firewall Administrator for the Authority Level field. See “Adding a User to the IBM Firewall” on page 61 for more details on completing the Add User menu.

Then, select the Administrator tab at the top of the Add User menu. You get a screen similar to the one shown in Figure 30.

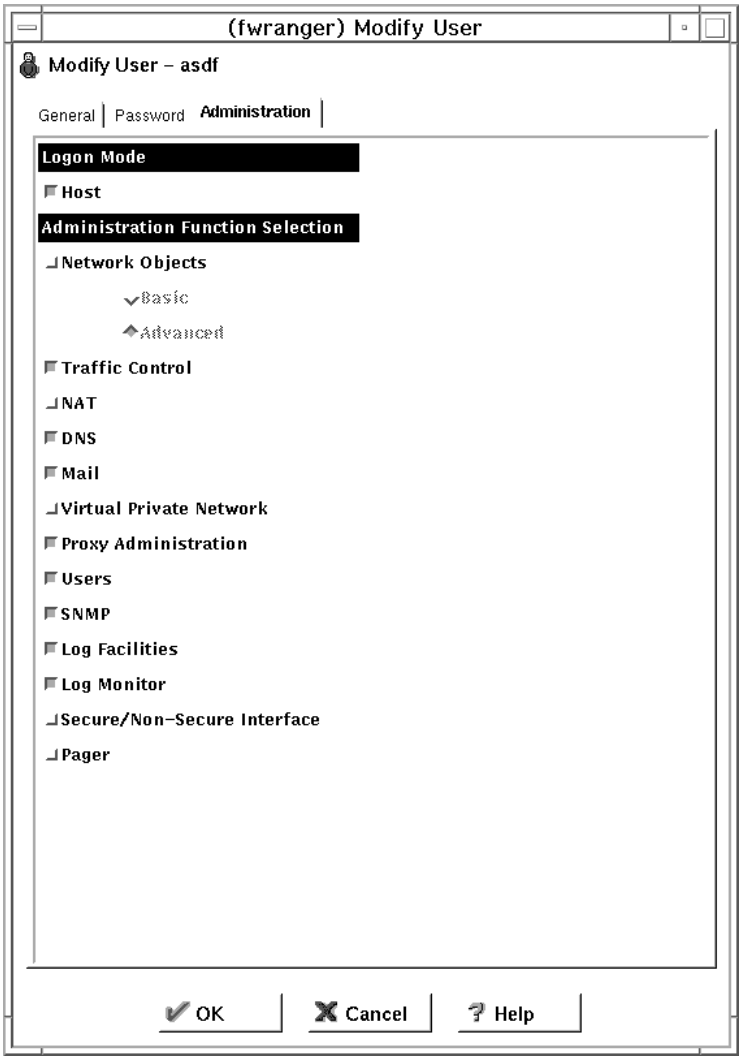

Figure 30. Administrator Tab

Select which functions the particular administrator will have authority over.

Setting Up and Administering the Idle Proxy Environment

An administrator can disconnect proxy connections to the firewall that have been idle for a specific period of time. Users are first warned and if their connection continues to remain inactive for an additional specified period of time, they are disconnected.

Safeguards for the Proper Working of Idle Proxy

To ensure smooth and correct functioning for idle proxy, follow these safeguards:

- Because this process disconnects other processes, it is essential that this process be run by root only. No other firewall users should be allowed to run this process.
- The idle proxy process disconnects all non-interactive sessions that exceed the disconnect time.

  Note: If a batch job produces output to the terminal, the job is terminated if the disconnect times are met. So, if you are running applications that use the terminal as a standard output device, consult your firewall or system administrator to modify the disconnect times or user IDs accordingly.
- The most efficient or convenient means of running this process would be to set it up as a cron job. This periodically executes the idle proxy process at a predetermined frequency.

  Root must set up the crontab file to determine the frequency of execution of Idle proxy.

  The idle proxy process can be run either from the command line by issuing the `fwidleout` command or by setting up the process as a cron job. The ideal situation would be to set up the process as a cron job, periodically checking for inactive users and disconnecting their processes. For example, if you want to set up the idle proxy to run every 10 minutes for every day of the year, type `crontab -e` and add the following line:

  ```
  0,10,20,30,40,50 * * * * /usr/bin/fwidleout
  ```

  Or, if you want to set up the idle proxy to run every 30 minutes on every alternate day for all days of the year, the system crontab entry could look like this. (Use the crontab -e command to edit the root's cron table).

  ```
  0,30 * 1,3,5,7,9,11,13,15,17,19,21,23,25,27,29,31 * * /usr/bin/fwidleout
  ```
- This process writes a log record using the standard firewall syslog facility. It is logged to `local4`.

Chapter 11. Enterprise Firewall Management

This chapter describes the Enterprise Firewall Management (EFM) function, which allows an administrator to control and update firewalls from one central location.

How EFM Works

An administrator logs on to the EFM Firewall (a central server) in EFM mode. The administrator selects the firewall for which he or she wants to perform configuration tasks. The administrator can then configure functions for that managed firewall. A copy of the managed firewall's configuration files are kept on the EFM Firewall. During configuration, these local files are the ones that are updated. When configuration tasks are complete, the EFM administrator distributes the changes he or she made to the managed firewall. The configuration files that are sent to the managed firewall are not activated until an EFM administrator activates them.

EFM allows an administrator to clone a new firewall's configuration definitions from a firewall that is already managed at the EFM.

Before configuring functions for a managed firewall, an EFM administrator must first create a firewall object for that managed firewall. The EFM administrator then assigns a security agreement to the managed firewall object. The security agreement indicates which functions can be configured by the EFM Firewall and which functions can be configured by the managed firewall itself. Each function can only be configured in one location.

In order to configure a function for a managed firewall, an administrator must:

- Have the authority to log on to the EFM Firewall in EFM mode
- Have the authority to configure that specific function
- Be configuring a function that can be configured by the EFM according to the security agreement for that firewall

Note that the user root always has the authority to logon in both EFM and host mode, and can always perform all configuration tasks.

EFM uses IPSec tunnel transport and security features to communicate and transmit data to the managed firewalls. Communication between the EFM and remote firewalls can be encrypted and/or authenticated. DES (US and Canada) or CDMF encryption schemes can be used for VPN tunnel sessions. Frequency for automated key exchange can be set as desired at the EFM. The EFM owns the connection.

Installation

The EFM fileset is installed as a separate component of the IBM Firewall. You must install it on the EFM Firewall (the firewall that will manage other firewalls). Do not install it on the managed firewall.

Note: Configuration clients that were installed from IBM Firewall V3R1 must be reinstalled from IBM Firewall V3R1.1.

Setup

To set up your EFM Firewall to manage a remote firewall, do the following:

1. Create a tunnel connection between the managed firewall and the EFM Firewall and activate it. See Chapter 13, “Creating a Virtual Private Network” on page 95 for information on how to create a tunnel connection.
2. Log on to the EFM Firewall in EFM mode and create a managed firewall object for the managed firewall.
3. You must get the tunnel definitions you just created from the managed firewall to the EFM Firewall. You can do this by either:
   a. Logging on in EFM mode and duplicating the tunnel definitions created on the managed firewall.
   b. Copying the tunnel configuration files `fwmctx`, `fwmctx.manual`, and `fwpolicy` from the `etc/security` directory on the managed firewall into the `/etc/security/efm/firewallname` directory on the EFM Firewall, where firewallname is the name of the managed firewall object.
4. On the managed firewall, make the following changes so that the EFM firewall can communicate with the managed firewall.
   a. Add the following line to `/etc/services`:

      ```
      efmd    1024/tcp
      ```

      Note that you can use a different port number than 1024, but whatever is used must match the port number specified when creating the managed firewall object.
   b. Add the following line to `/etc/inetd.conf`:

      ```
      efmd stream tcp nowait root /usr/sbin/efmd efmd
      ```

   c. Issue the following command or reboot:

      ```
      refresh -s inetd
      ```

   d. On the EFM Firewall, create a connection between the EFM Firewall and the managed firewall using the EFM predefined service.
   e. On the managed firewall, create a connection between the managed firewall and the EFM firewall using the managed firewall predefined service. Note that if traffic control is going to be an EFM controlled function, you will need to also create this connection in the managed firewall's definitions on the EFM Firewall.

Logon and Managed Firewall Object

In order to configure managed firewalls using EFM, you must specify Enterprise mode when logging on with the configuration client. After you have logged on in EFM mode, the configuration client displays the name of the EFM Firewall and the managed firewall, when appropriate, so that you always know which configuration files are being modified.

Figure 31. Firewall Configuration Client Panel

Administrator Logon

You log on to the EFM Firewall with mode set to Enterprise to perform EFM administration.

If you initialize the configuration client with mode set to Enterprise, the names of the EFM and managed firewall are displayed on the firewall configuration client panel. From this panel you can activate the Select function to display the list of firewalls that are administrated by the EFM. This information is presented on the Select Firewall to Configure panel, see Figure 32 on page 74. You update configuration files for the displayed Firewall name.

Authorized Functions

Only authorized functions that can be performed for the firewall and by the EFM administrator are displayed in the left window of this panel. Authorized functions are defined by the security agreement for the managed firewall and the administrator's authority. See “Administrator Authority Level by Function” on page 69 for more information.

Alerts and Log Viewer

The Alerts/Log Viewer Window on the main configuration client panel can display alert and log information for firewalls managed from the EFM. Alerts and log information for the EFM Firewall can also be displayed in this window. You can redirect alert and log information to the EFM Firewall by following current redirection procedures for the `syslog.conf` file. This file must be updated on the managed firewall to redirect alert and/or log information to the EFM Firewall's `local1` and `local4` log facilities. Alert information is directed to `local1` while log viewer information is directed to `local4`.

This information can also be directed to another firewall or a stand-alone server used for report generation. Specifically, the syslog daemon on each remote firewall can write log records to the `local4` facility on a stand-alone server that is used for report generation.

You might want to direct alert information to the `local1` on the managed firewall and EFM Firewall. Due to the high volume of log records, you might want to record this information on the managed firewall and a stand-alone server that is used for report utility processing.

Managed Firewall Objects

Click Select on the main firewall configuration client Panel to get the following panel, which displays the list of managed firewalls supported by the EFM.

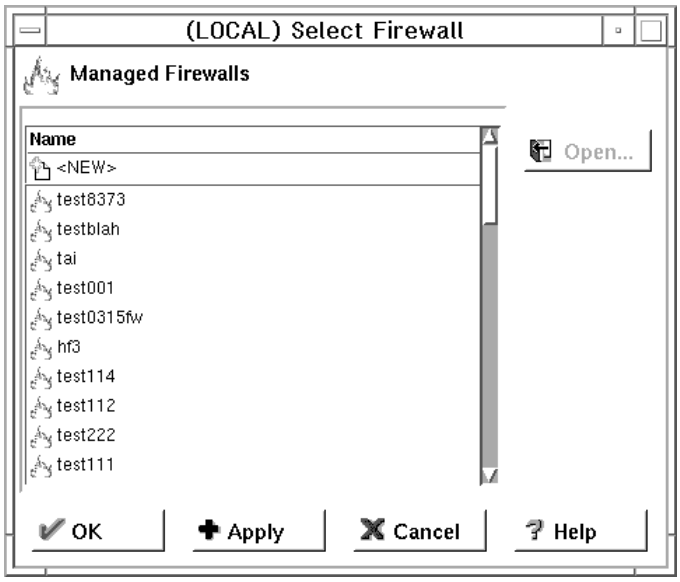

Figure 32. Select Firewall to Configure

Select the firewall to configure and click OK. The selected firewall name is displayed as the Managed Firewall on the firewall configuration client Panel.

You can create a new firewall object by selecting NEW and clicking Open. The Managed Firewall panel Figure 34 on page 76 is displayed. An EFM administrator must be authorized to perform the Managed Firewall Objects function to create a new managed firewall object. Any EFM administrator can display the Select firewall to Configure panel from the main configuration client panel.

The following panel displays the list of managed firewalls supported by the EFM, when you activate the managed firewall list item on the main configuration client panel.

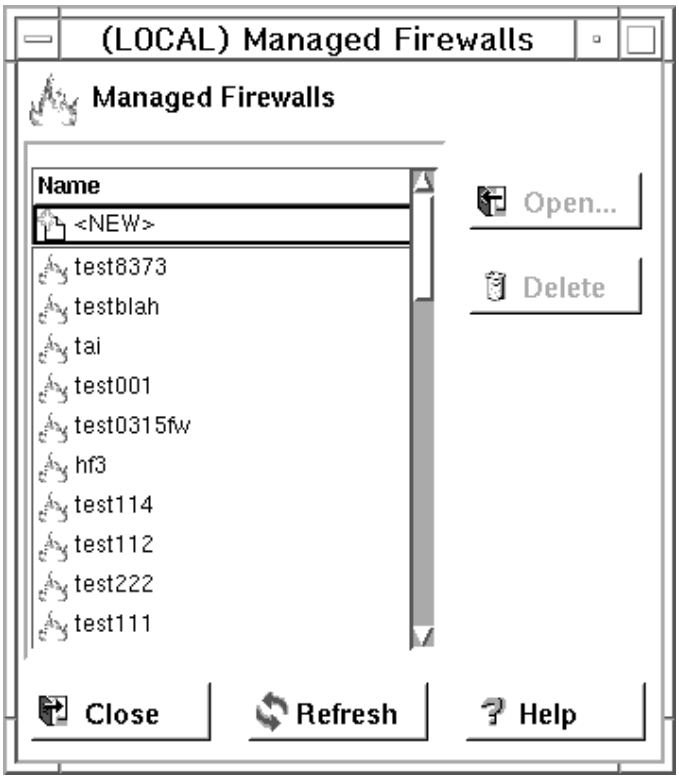

Figure 33. Managed Firewalls

You can view information for a managed firewall by highlighting the firewall name and clicking Open. In this case, the Managed Firewall panel Figure 34 on page 76 is displayed with detail information for the firewall. Any EFM administrator is able to view this information regardless of their administrator authority level.

You can delete a managed firewall object by highlighting the desired firewall and clicking Delete. You must be authorized to perform the Managed Firewall Objects function to delete an existing firewall object. When an authorized EFM administrator attempts to delete an existing managed firewall object, a message box appears requesting the administrator to confirm the deletion. If the deletion is confirmed, all configuration files for the managed firewall will be deleted at the EFM.

The managed firewall object is only accessible when you are logged on in Enterprise mode. It will not be listed as a network object when you are logged on the EFM Firewall in Host mode.

A managed firewall object can be created with the following panel.

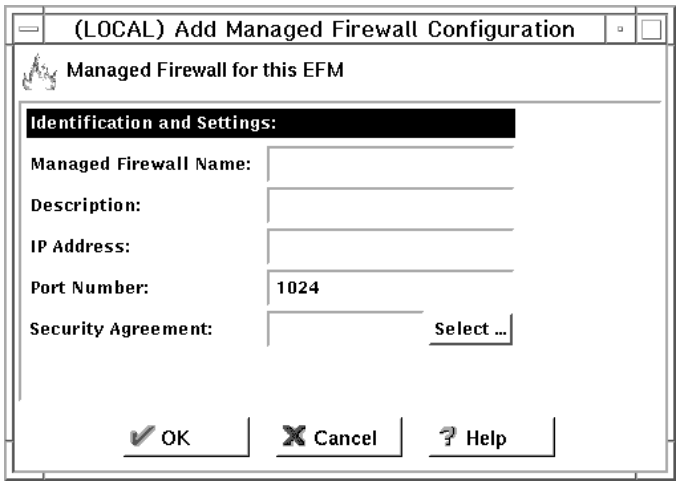

Figure 34. Managed Firewall

Managed firewall name, IP address and port number are required fields. The managed firewall name is the object name that is displayed on all EFM panels. It is recommended that you use the fully qualified hostname for the object name. However, you can use an alias or IP address for the object name. The IP address is a valid address that is used for the tunnel connection between the EFM and the firewall. Default port number 1024 is displayed when the panel is initialized. You can change the port number but it must match the port number you specified when setting up the managed firewall. See 4a on page 72.

Security Agreement

When creating a managed firewall, you must specify a security agreement in order to define which functions are managed by the EFM and which functions are managed locally. The default security agreement specifies that all functions are managed locally.

You must be authorized to perform the Managed Firewall Objects function to create, change, or delete a Security Agreement.

If desired, you can assign another security agreement by activating the select button to display the Select Security Agreement panel Figure 35 on page 77. If a different security agreement is specified through the select process, this name will be displayed as the assigned security agreement when you return to the Managed Firewall panel.

The following panel is displayed when you activate the Select button on the Managed Firewall panel.

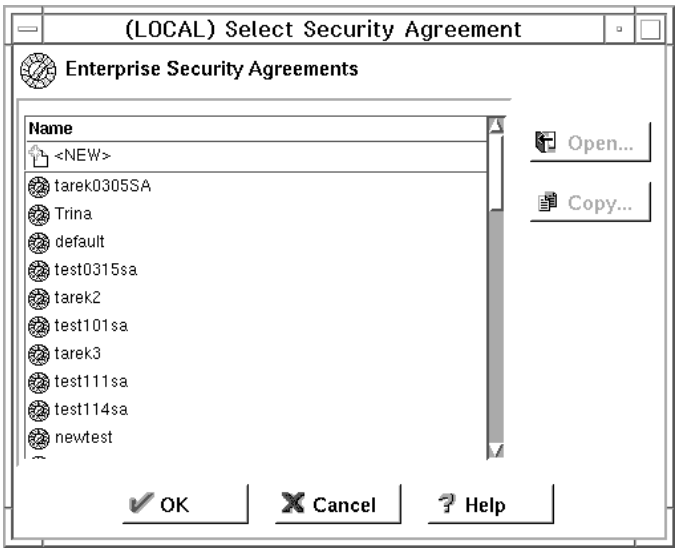

Figure 35. Select Security Agreement

You can assign a different security agreement to the new firewall by selecting a listed security agreement and clicking OK. You can view detailed information for a security agreement by highlighting an agreement entry and clicking Open. In this case, the Open Enterprise Security Agreement panel, which is shown in Figure 37 on page 79, is displayed with applicable security agreement information.

If needed, you can also add a new agreement, if you are authorized, by selecting NEW and clicking Open. The Open Enterprise Security Agreement panel, shown in Figure 37 on page 79, is displayed with blank information. Enter the desired information to create a new agreement record.

You can also copy a security agreement by highlighting an agreement name and clicking Copy. The Open Enterprise Security Agreement panel, shown in Figure 37 on page 79, is displayed with information for the selected agreement. However, the security agreement name is blank.

The security agreement file entry defines which resource (for example, the EFM or remote firewall) controls a particular function. Administrators at the EFM, who may have authority to perform a function, are not permitted to perform the function if it violates the security agreement. For example, administrator 1 at the EFM may be authorized to perform proxy user updates. However, the security agreement for a particular firewall specifies that all proxy user updates are to be performed by the local administrators at the remote firewall. In this case, administrator 1 is not able to modify proxy user information for the managed firewall.

The following function categories are defined in the security agreement record:

Address Translation
DNS
Log Facility (2)
Log Monitor
Mail
Network Objects
Pager
Users
Proxy Administration (3)

Secure/Non-Secure Interfaces
SNMP
Traffic Control (1)
VPN

1. Includes configuration updates for Security Policy excluding transparent proxy.
2. Includes Report Utilities activation.
3. Includes RealAudio**, HTTP proxy and transparent proxy configuration updates.

Figure 36 is displayed when the Security Agreements list item is activated from the main configuration client panel.

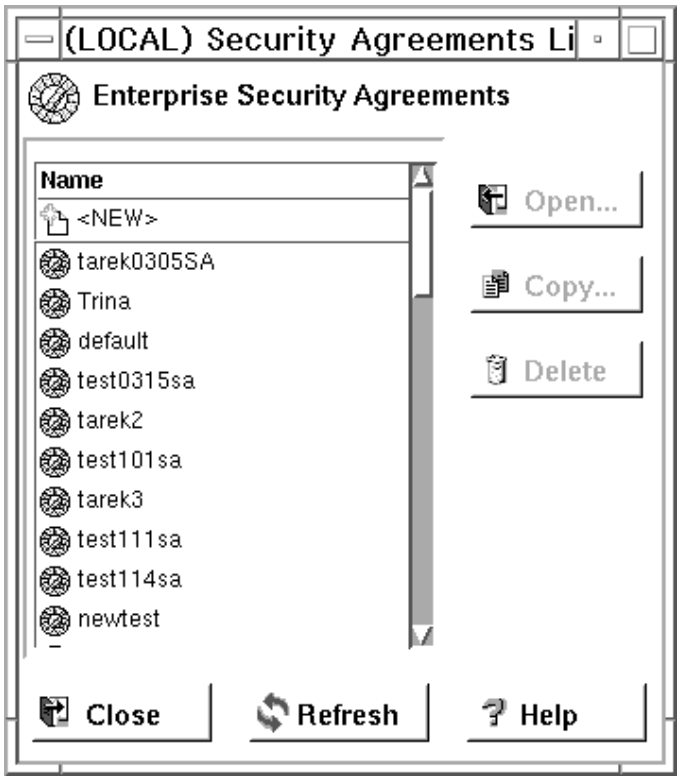

Figure 36. Security Agreement Selection List

It is used to add, copy or delete security agreements at the EFM. If you click Open, the Open Enterprise Security Agreement panel is displayed, as shown in Figure 37 on page 79. If you try to delete a security agreement, all firewall definitions are checked to verify that the agreement is not assigned to a firewall.

Use the following panel to define the security agreement for firewall management. The agreement record is used to define configuration functions that are controlled by administrators at the EFM or at the managed firewall.

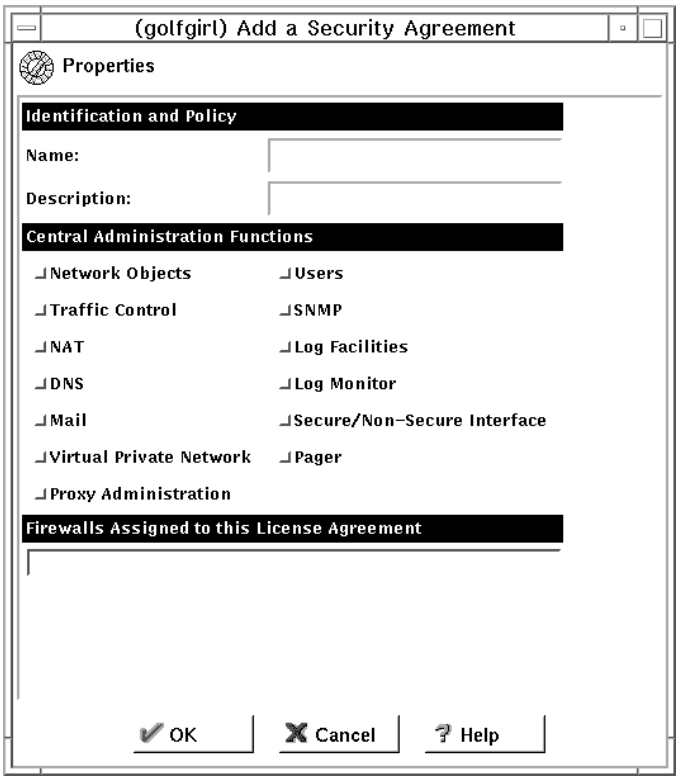

Figure 37. Open Enterprise Security Agreement

The names of the firewalls assigned to the security agreement are displayed in the lower window of this panel.

The security agreement record when created or changed must be transferred and activated to the remote firewall by an authorized administrator at the EFM.

Configuring a Managed Firewall

After you have created a managed firewall object and selected that object as the one to be managed, you can now make configuration changes to that firewall. You will see the list of functions that can be configured on the left side of the panel in Figure 31 on page 73. Configuration changes are kept at the EFM machine until you transfer and activate them at the managed firewall.

Some configuration tasks allowed in host mode are not allowed in EFM mode due to the nature of those tasks. Those items do not show up on the configuration client panel. Some of these are: NAT activation and deactivation, tunnel connections by tunnel ID, and disablement of NAT logging and filter explosion.

Session Monitor

Before you specify the maximum number of TCP and UDP sessions, it is important to evaluate the total number of TCP sessions because TCP sessions are used to distribute files and activate from the EFM Firewall.

An administrator with session limit authority is permitted to control the number of concurrent sessions on a managed firewall. Figure 38 on page 80 identifies the session monitor information an administrator can enter.

License requirements for host address pricing levels are part of the managed firewall. The number of IP host addresses for secure to nonsecure connections are monitored on the managed firewall.

Figure 38 is displayed when the Session Monitor list item is selected from the System Administration folder on the configuration client navigation tree.

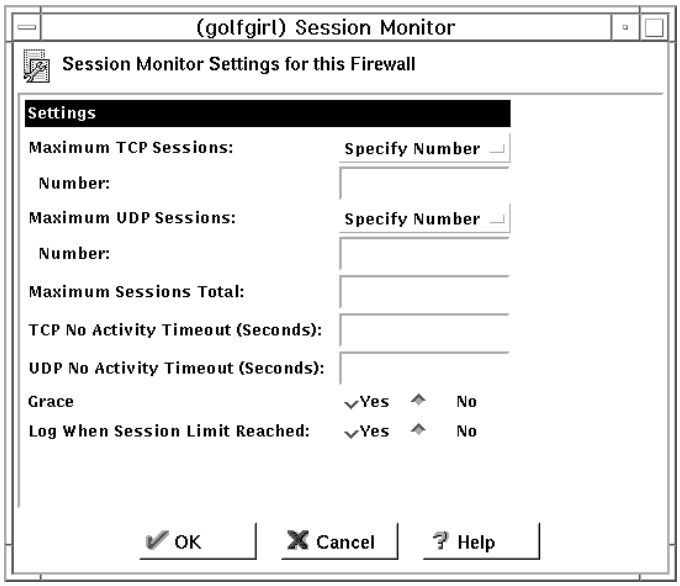

Figure 38. Session Monitor

This panel is only displayed when you have logged on in Enterprise mode.

You specify the maximum number of TCP and UDP sessions for the managed firewall. If these values are unlimited, you can select Unlimited from the maximum TCP or UDP sessions pull-down menu. If you need to define a limit, you select Specify Number from the list and enter a value in the number field. The maximum sessions total is calculated and displayed for you.

The minimum number of TCP sessions is 10 and the minimum number of UDP sessions is 10. The combined minimum number of TCP and UDP sessions is 50. The maximum number is 1,000,000.

You define the TCP and UDP no activity timeout values. The range is -1 for no timeout up to 9999999 maximum timeout.

You can elect to implement a hard stop if the limit is reached for the session type by setting the grace button to No. If set to yes, any session type requests over the maximum session type value will be allowed. If the grace button is set to Yes, any session type requests over the maximum session type value will be allowed.

You can specify if logging should occur when the TCP or UDP limit is exceeded. Because logging for excessive sessions could significantly impact firewall performance, you can determine whether logging should always occur for this event. If logging is set to No, an error message will be written to the log if the grace period is set to No. Messages will not be written to the log if logging is set to No regardless of the grace period setting.

Firewall Clone

You can use the firewall clone feature to quickly create initial configuration files for a new firewall from an existing firewall's configuration files. Once the cloning function has been completed, you can change other configuration processes to modify the initially created definitions.

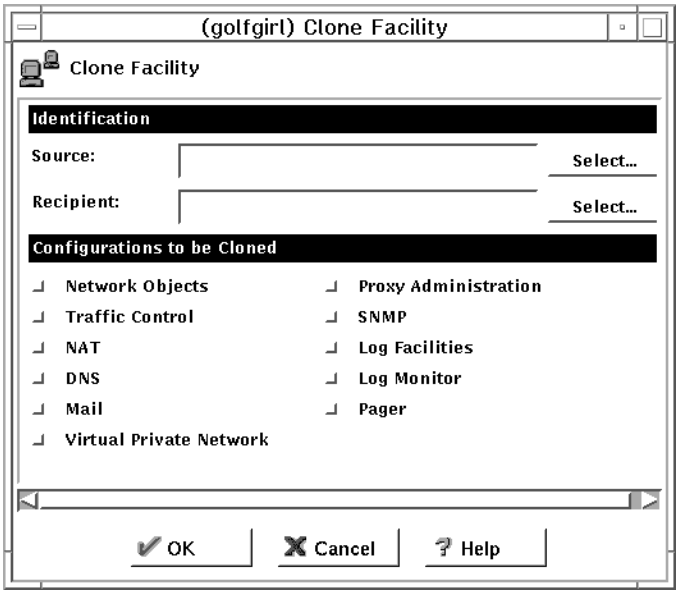

Figure 39. Firewall Clone

Required configuration files for each function are copied from the directory of the source firewall to the directory of the recipient firewall. Configuration categories are displayed for functions that are supported at the EFM for the source firewall. The source firewall's security agreement record will be checked to identify these functions.

You must first create a firewall object for the recipient firewall before it can be cloned. The EFM administrator must have managed firewall objects administrator authority to perform the clone function.

File Integrity Checker

File Integrity Checker is not used on the EFM for configuration files that are maintained for managed firewalls. However, it can be performed at the managed firewall when new configuration files are activated. The checksums for managed files must be updated on the managed firewall.

When logged on in Host mode on a managed or non-managed firewall, you must have root authority to perform file system integrity checking.

Users

If users are managed from the EFM Firewall, user updates are sent directly to the managed firewall. The user request is immediately processed by the managed firewall and appropriate files are updated. Instead of updating a user file that is located on the EFM firewall, the EFM configuration client (at the usual file update point) will issue a request to ship the update transaction to the managed firewall.

Note that if the machine that is managing Users changes from the EFM to the local machine or from the local machine to the EFM machine, you must immediately

transfer and activate the Security Agreement. Otherwise the two machines will be out of synch and both machines will be able to make User changes.

Security Policy and Transparent Proxy

Configuration values for transparent proxy are set in the Security Policy configuration client panel. Security policy administration authority is grouped with traffic control. Administration authority for transparent proxy is controlled by proxy administration. Depending on an administrator's authority and approvals in the security agreement record for the firewall, select fields are enabled or disabled when the security policy panel is displayed. Security policy and transparent proxy fields are enabled if you are authorized to perform traffic control and proxy administration and the firewall's security agreement record also authorizes these updates. If your administration record and security agreement record do not authorize these functions, security policy or transparent proxy fields for the unauthorized functions, will not be enabled for input.

Distribution and Activation

Configuration changes at the EFM for a managed firewall do not take effect until they are distributed and activated. Distribution sends the configuration updates to the managed firewall. Activation puts those changes into use at the managed firewall.

Configuration File Transmission Processing

If authorized to perform configuration file transmission transactions, you can ship files for one or multiple firewalls based on the following selection criteria:

- Elect to transmit files for functions whose configuration definitions have changed since the last transaction
- Force the transmission of files for select functions

You are asked to identify or select functions that should be updated. The detail files to be transmitted are not presented on the panel. However, the names of the actual files that were transmitted will be listed in syslog.

On the EFM's Firewall, a message is written to the syslog file to record the transmitted event. A corresponding message is written to record the successful or unsuccessful load of a file on the remote firewall.

Figure 40 on page 83 is used to transmit or distribute configuration files from the EFM to the remote firewall.

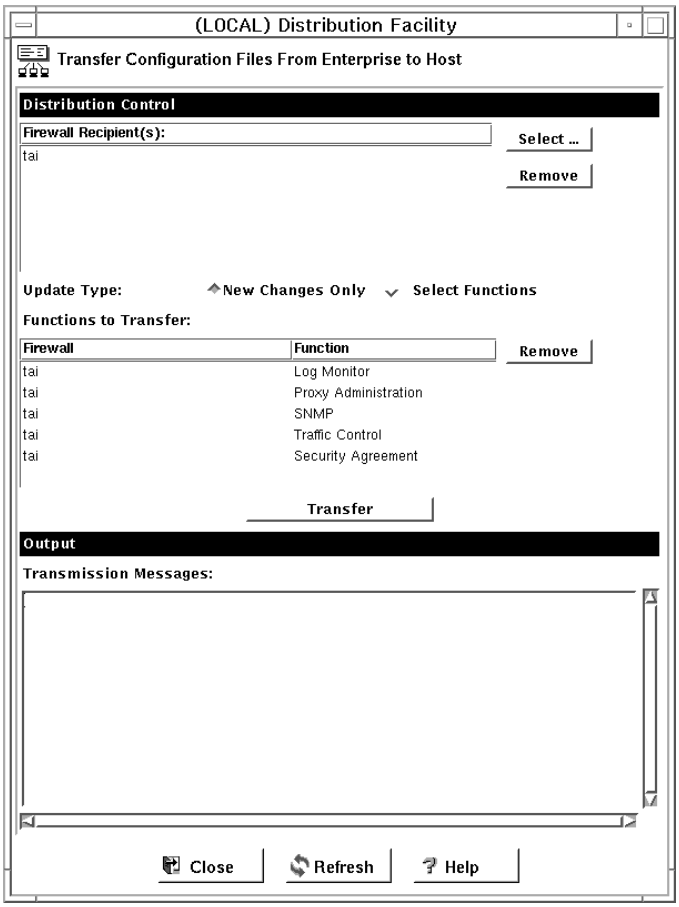

Figure 40. Distribution Facility

You can elect to transmit configuration files that have changed since the last transmission (net changes only). You can also force the transmission of configuration files for select functions. The names of functions whose configuration files are available for ship per managed firewall are displayed in the Functions to Transmit window. If net changes are to be transmitted, only functions with changed definitions are displayed. If select functions are to be transmitted, all functions supported by the EFM (per the security agreement record) are displayed. You can elect not to transmit configuration files for displayed functions by clicking Remove to remove the function name from this window. Files applicable to functions displayed in the Functions to Transmit window, are sent to the remote firewall when the Transmit button is activated.

Traffic control is dependent on the most current network object information. Any time traffic control files are transmitted, the network objects file should also be transmitted if it is controlled at the EFM and if it has been changed.

The following functions can be listed in the function window for a firewall based on authorizations in the assigned security agreement: DNS, Mail, Log Facilities, Log Monitor, Proxy Administration, VPN, Pager, Network Objects, Traffic Control, Address Translation, SNMP, Interfaces, Session Monitor, and Security Agreement. Session Monitor and Security Agreement are not defined in the security agreement file. By default, these are always controlled by the EFM.

Message responses indicating successful or failed update on the remote firewall are displayed in the Transmission Messages window.

Activation Processing

After files have been transmitted and stored in the holding directory at the remote firewall, an EFM administrator must activate the changes. During activation, files are copied to required directory paths and commands are processed or daemons refreshed to activate configuration definitions.

The following functions can be listed in the function window for a firewall based on previously transmitted file information and information in the security agreement: DNS, Mail, Log Facilities, Log Monitor, Proxy Administration, VPN, Pager, Network Objects, Traffic Control, Address Translation, SNMP, Interfaces, Session Monitor, and Security Agreement. Session Monitor and Security Agreement are not defined in the security agreement file. By default, these are always controlled by the EFM.

Managed Firewall Activation: The panel shown in Figure 41 is used to activate previously transmitted configuration file definitions on the remote firewall. An EFM administrator can also use this function to activate managed firewall functions even if a modified configuration file was not transmitted.

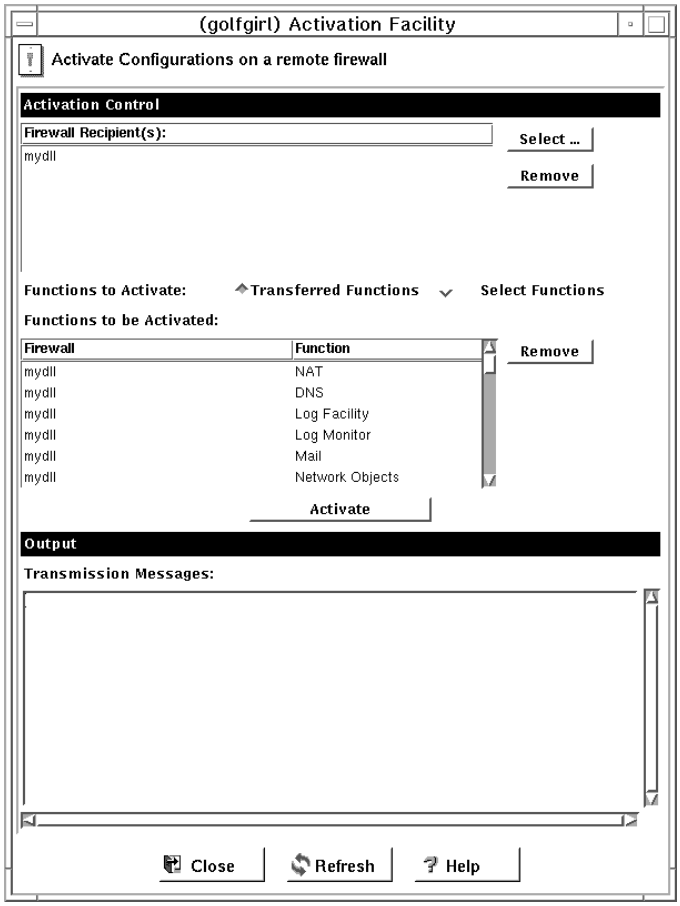

Figure 41. Activation Facility

Messages denoting successful or failed activation are sent from the managed firewall machine to the EFM. These messages are displayed in the Output window. The messages are also written to the syslog of the EFM's firewall.

VPN Connectivity to Remote Firewalls: A secure IP tunnel is implemented between the EFM and each remote firewall. The VPN connection is used to pass configuration file information when transmit requests are initiated by the EFM.

Log Facilities Definitions: When configuration files related to the definition of log facilities are activated, the definitions received from the EFM Firewall overwrite any existing definitions that are currently on the managed firewall.

VPN Definitions: When VPN definitions are received from the EFM Firewall and activated, any VPN definitions that already exist on the managed firewall are deleted.

When configuring the VPN definitions on the EFM Firewall, the administrator indicates which tunnels should be activated and which tunnels should be deactivated. When the managed firewall receives a command to activate VPN definitions, the managed firewall will activate and deactivate the tunnels as the administrator indicated.

Reconnecting to a Managed Firewall

If the managed firewall's connections or VPN definitions are misconfigured, it is possible that the EFM Firewall will be unable to communicate with the managed firewall. If this occurs, follow these instructions for reestablishing communications between the EFM Firewall and the managed firewall:

- Log on locally to the managed firewall with the root password.
- Change to the `/etc/security/` directory.
- Copy `fwconns.cfg.BAK` to `fwconns.cfg`. This will put a working copy of the filter connection file in place to be activated. If there are problems preventing communication with the managed firewall other than a bad connection, you might have to copy all of the fw*.cfg.BAK to the corresponding cfg file.
- Edit `secag.cfg` and change the following two lines:
  1. Traffic:efm to Traffic:host
  2. VPN:efm to VPN:host
- Start the configuration client and log in to the managed firewall as root in Host mode.
- Open the activation window under Traffic Control. Regenerate the Connection Rules from this panel. This will recreate a working set of filters and activate them.
- Open the Virtual Private Network window under Traffic Control. Chose the VPN going to the EFM Manager and activate this VPN. This should allow the manager to regain a connection to the managed firewall.
- On the EFM manager fix the problem that caused the connection to be lost. Force the security agreement to be distributed and activated with the corrected filter rules. Reactivate from the manager. The manager and managed firewall should be back to the original state before the problem occurred.

Chapter 12. Using Proxy Servers

This chapter contains general information about how to use the proxy servers from workstations both inside and outside your secure network. For specific host names and procedures, be sure to ask your network administrator or the person who administers your firewall.

The proxy servers are started automatically as required.

HTTP Proxy

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.

The administrator has to configure HTTP proxy and start it. The user needs to change the proxy pointer on the configuration page of their browser to point to the IBM Firewall and the proper port (the proxy port number field).

If you change any parameters, you must stop your proxy and then start it again for the change to take effect.

You must allow DNS queries before HTTP Proxy can work properly. An easy way to do this is to click Security Policy from inside the System Administration folder on the configuration client navigation tree and click Permit DNS Queries.

A predefined service allows an HTTP client to reach the proxy server on port 8080, but does not allow the client to pass through to the nonsecure side of the firewall.

Click HTTP in the navigation tree and ensure that the HTTP Proxy configuration settings are suitable for your purposes. Note that port 8080 is the default here. If you change this, the port number must also be changed in the Services that are set up for this configuration. If you change any of these settings you must restart the phttpd process.

Ensure that the phttpd process is running. If it is not, you can type `phttpd` at the AIX command line to start it. To have it start automatically each time the firewall is rebooted, uncomment the phttpd line in the `/etc/rc.tcpip` file so that it looks like the following:

```
## If you want the HTTP proxy daemon to always             #FW#
## start at boot time, uncomment the following line.       #FW#
/usr/sbin/phttpd
```

See “Example of Proxy HTTP” on page 45 for an example of how to set up a connection on the nonsecure side of your network.

Configuring HTTP Proxy Using the Configuration Client

To configure HTTP Proxy, select HTTP from the configuration client navigation tree. The IBM Firewall displays the HTTP Proxy menu, as shown in Figure 42.

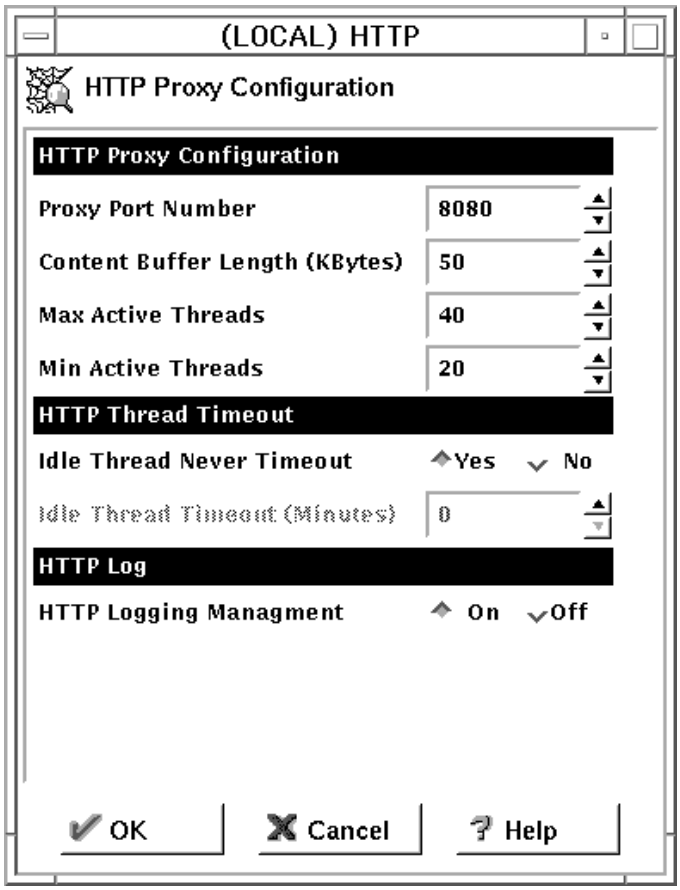

Figure 42. HTTP

The administrator configures the following parameters on the HTTP Proxy menu.

Proxy Port Number

Use this parameter to specify the port number the proxy should listen to for requests. If you change the port number, you must configure your filters to allow or disallow flow through the ports. Port numbers less than 1024 are reserved for TCP/IP applications. Common ports used for proxy Web servers are 8080 and 8008.

The default filter rules are set to disallow inbound, nonsecure traffic on port 8080, but allow secure traffic on that same port. Adjust these rules accordingly. The default is `8080`.

Content Buffer Length

Use this parameter to set the size of the buffer for dynamic data generated by the server. Dynamic data is output from CGI programs, server-side includes, and API programs. It is data that does not come from a proxy.

The value can be specified in bytes (B), kilobytes (K), megabytes (M), or gigabytes (G). It does not matter if you have a space between the number and the value (B, K, M, G). The default is `50 K`.

Max Active Threads

Use this parameter to set the maximum number of threads that you want to have active at one time. If the maximum is reached, the proxy holds new requests until another request finishes and threads become available. Generally, the more power a machine has, the higher the value you should use for this parameter. If a machine starts to spend too much time on overhead tasks, such as swapping memory, try reducing this value. The default is `40`.

Min Active Threads

Use this parameter to set the minimum number of threads that you want the proxy to have available for use. The server will not close threads below this minimum even if the threads are idle. Generally, the more power a machine has, the higher the value you should use for this parameter. If a machine starts to spend too much time on overhead tasks, such as swapping memory, try reducing this value. The default is `20`.

Idle Thread Never Timeout

Use this parameter to specify how long the proxy should keep an idle thread available. A thread becomes idle after the last request to use it completes. If the number of threads already available or active is greater than the value on MinActiveThreads and the proxy does not use the thread again within the specified time, it closes the idle thread.

Specify the time value in minutes. If you use the default value of `forever`, the server does not close any idle threads.

HTTP Logging Management

This parameter tells the proxy to log startup/shutdown and all proxy requests to the AIX Syslog. It uses the LOG_NOTICE level of logging. Set this to on if you wish to monitor HTTP request activity. Events are logged in `local4`.

Starting the HTTP Proxy (Command Line Only)

During installation, a start command for the executable (phttpd) was added as a comment to the `/etc/rc.tcpip` file. You can modify this file to uncomment the start command and also add additional parameters (such as -p for a different port), or you can start the proxy from the command line. It is not advisable to register it as a subsystem and start with the STARTSRC command, although technically it will work, since it will leave a parent process under user name **root** until you stop the subsystem. Normally phttpd runs under user name **nobody** authority.

Stopping the HTTP Proxy (Command Line Only)

You can stop the HTTP proxy by logging on to the firewall as root and using the **kill** command. To give the proxy time to write the log entry showing shutdown has occurred, use the **kill cat /etc/httpd-pidfile** command.

User Actions

In the client browser configuration, point to the HTTP proxy on the IBM Firewall. Use the port that the administrator configured in the configuration client as the port that the proxy will be listening on.

SSL Connections

SSL tunneling for HTTP Secure Connection to other servers is supported. The IBM Firewall acts as a gateway in this case. The tunnel goes from the client through the firewall to the server. Use the standard port 443 for HTTP Secure Connection as shown in the following example:

```
http://www.ibm.com:443
```

Also, use the predefined service `HTTS proxy out 2/2`.

Secure News is supported and uses port 563. These ports can be selected in the URL of the destination. See the following example.

```
http://www.ibm.com:563
```

Methods Supported

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

- HTTP
- FTP
- Gopher
- WAIS

FTP with the Proxy Server

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

   ```
   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: me_user
   ```

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 61). See “Authenticating Users at the Proxy Server” on page 92 for information about how users are authenticated by proxy servers.

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. Again, remember that this is most likely a different user name and password from those you used to FTP to the Firewall.

FTP with Transparent Proxy

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 our host name.)

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

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

   ```
   username:
   ```

3. Enter your user name at the nonsecure network:

   ```
   username: username@remote_site_host_name
   ```

4. You are then prompted for your password for the target host.

   ```
   password:
   ```

5. Enter your password.

Telnet with the Proxy Server

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 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 as authorized to use the firewall:

   ```
   login: me_user
   ```

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 61). See “Authenticating Users at the Proxy Server” on page 92 for information about how users are authenticated by proxy servers.

After you are authenticated, the proxy server displays a command prompt. You can now use telnet to access the foreign host:

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

And this time, the foreign host asks for your user name and password, as you are known on that host. Remember, these might be different from the user name and password that you used on the firewall proxy server.

Telnet with Transparent Proxy

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:

   ```
   username:
   ```

3. Enter your user name at the nonsecure network:

   ```
   username: username@remote_site_host_name
   ```

4. You are then prompted for your password for the target host.

   ```
   password:
   ```

5. Enter your password.

Authenticating Users at the Proxy Server

When you use telnet or FTP to access the proxy server, the server tries to authenticate your identity using the method that was specified when your user name was added to the firewall (see “Adding a User to the IBM Firewall” on page 61). The possible choices are:

- Deny
- None
- Password
- SecurID card
- SecureNet Key card
- User Supplied Authentication

Deny

The IBM Firewall prohibits access to the server.

None

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

Password Authentication

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

```
Password:
```

Enter your password. This is the same password with which your user name was added to the Firewall.

SecurID Card Authentication

Use this method if you have a SecurID card and your network uses the Security Dynamics ACE/Server.

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, 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.

SecureNet Key Authentication

Use this method if you have a Assurenet Pathways SecureNet Key card.

The proxy server will ask for a response provided by your SecureNet Key card, before letting you proceed.

```
Use SNK for challenge
######## for user user_id
Ed:
```

The challenge `########` is an 8-digit number that you enter into the SecureNet Key card.

1. When you receive this prompt, activate your SecureNet Key card and enter your PIN code. The PIN code was given to you along with the card.
2. Enter the challenge as provided by the server.

   For example: you log into the server; the server prompts:

   ```
   Use SNK for challenge
      78987648 for user NEWUSER
   Ed:
   ```

   Enter the value 78987648 into the SecureNet Key card. The card then displays the response, which you provide to the proxy server.
3. Enter this response to the server.

   If the SecureNet Key card displayed `8AE222A9` in response to your challenge, then you enter `8AE222A9` to the server:

```
logon: NEWUSER
   Use SNK for challenge 78987648 for user NEWUSER
Ed:8AE222A9
```

User Supplied Authentication

You can use the User Supplied Authentication method for FTP and telnet. See the *IBM Firewall Reference* for more information.

Chapter 13. Creating a Virtual Private Network

A **Virtual Private Network** is comprised of one or more **secure IP tunnels** and two or more networks. A secure IP tunnel describes the process of encapsulating a complete IP packet, including its header information, in a new IP packet seen by only the source and destination firewall hosts. The original IP packet is protected during the transmission between the two firewall hosts. You can configure a secure IP tunnel between any two firewall host addresses. The tunnel defined policy specifies that the data (original IP packets) be either:

- Encrypted
- Authenticated
- Encrypted and then authenticated
- Authenticated and then encrypted
- Neither encrypted nor 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 and/or authenticated data is integrated with the IP filtering rules. This provides the level of granularity, for defining which IP data packets will be encrypted and/or authenticated, at the IP address and port level. Any particular tunnel policy must have the same specification at both ends or the transmitted packets will be discarded.

Encryption and message authentication support requires that each of the two parties (tunnel end points) have a shared secret key. For IP tunnel support, manual administration of the shared master keys (and some additional data) is required for setup.

The actual keys used for encryption and message authentication are derived algorithmically from the base master key value and can be automatically refreshed on a timed basis. This significantly reduces the need to change the master key.

Two versions of the IBM Firewall are available. You can use either the Data Encryption Standard (DES) or the Commercial Data Masking Facility (CDMF) level of encryption (user selection) for the IP tunnel.

Because an operational IP tunnel requires administration definitions at both of the tunnel end points, those end points and the administration tasks associated with IP tunnel operations, are defined in terms of **tunnel owner** and **tunnel partner**.

Tunnel Types

The IBM Firewall allows you to create three kinds of tunnels: an IBM tunnel, a non-IBM or manual tunnel, and a dynamic tunnel.

- IBM tunnels are used between two IBM Firewall host addresses and feature an automatic key refresh mechicanism. For other firewall products you can define a non-IBM or manual tunnel.
- A manual tunnel uses the IPSec standard and can be used between an IBM Firewall and a non-IBM Firewall.

- A dynamic tunnel is used between a secure remote client and the firewall. It is configured but not activated until the remote client starts the tunnel. For more information see Chapter 14, “Using the Windows 95 Secure Remote Client” on page 103.

Figure 43 is an illustration of a tunnel and a VPN.

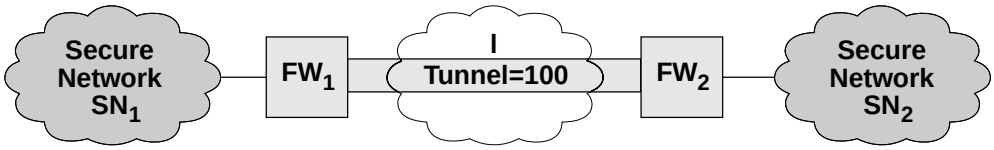

Figure 43. Tunnel, All IP Traffic between Two Secure Networks. FW_1 and FW_2 represent nonsecure interface IP address and mask. SN_1 and SN_2 represent any host in the secure network. The shaded area of the picture represents a VPN.

IP Tunnel Configuration and Activation (IBM and Manual Tunnels)

To configure and activate tunnel(s) you have to:

1. Add connections to allow the required firewall-to-firewall communication.
2. Create your tunnel.
3. Export a tunnel partner's policy and context appendages from the tunnel owner to the tunnel partner
4. Import (load) the tunnel partner policy and context appendages
5. Associate filter rule(s) at both IP tunnel end points, with a tunnel context
6. Activate tunnel policy for both IP tunnel end points

Note: Do NOT try to edit the tunnel files yourself. Do all IBM Firewall administration using the IBM Firewall configuration client menus or commands.

Steps 1, 5, and 6 must be taken for the firewall at each end of the tunnel. Steps 2 and 3 are the tunnel owner, step 4 is the tunnel partner. The IBM Firewall commands and a configuration client provide the ability to perform these steps. IP tunnel context descriptions can be incrementally added by any IBM Firewall. For example, a particular firewall may be a tunnel owner for one set of tunnels context definitions and a tunnel partner for other tunnel context specifications. Both tunnel policy and IP tunnel context definitions can be changed dynamically. Either or both of these definitions can be changed with a complete replacement of the *active* definition. With a complete replacement, the appropriate updates, deletions and additions are performed.

Configuring Tunnels Using the Configuration Client

This section describes how to use the configuration client to configure your tunnel(s) 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 menu, you can add, delete, import, export, activate, deactivate, and shutdown a tunnel.

Add a Tunnel

1. Select NEW from the Tunnels menu and click Open.

   A dialog asks you to specify the values required for a tunnel context ID specification, as shown in Figure 44.

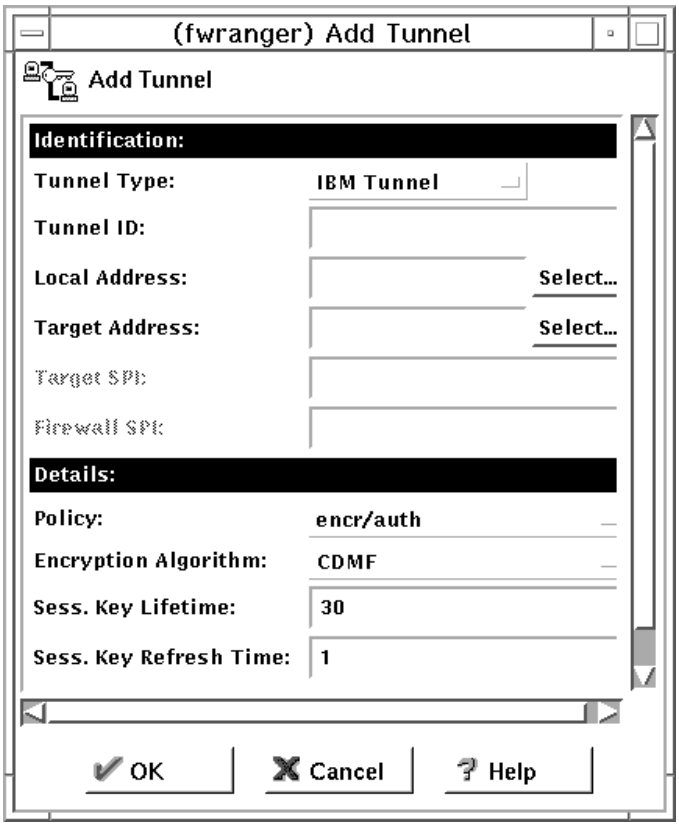

Figure 44. Add a Tunnel

2. Enter the following values:

Value	Description
Tunnel Type	Open the pull-down menu to select a tunnel type. An IBM tunnel is used between two IBM Firewall host addresses. A manual tunnel is used between an IBM Firewall host address and a non-IBM Firewall host address. A dynamic tunnel is only used with the Windows 95 secure remote client. It is based upon the firewall IP address and a user's ID and password rather than a source and destination address. **Note:** In order to define a dynamic tunnel, you must also have authority to define user network objects and the security agreement must allow you to define network objects. See “Administrator Authority Level by

Function” on page 69 and “Security Agreement” on page 76 for more information.

Tunnel ID

The identification number for the tunnel. Put a number in the entry field. This number must be unique at both ends of the tunnel. It can be 1 to 6 characters.

Local Address

IP address of the local firewall nonsecure interface to be used by the tunnel. Click Select to get the Interface list. Select an interface and click Apply. The local address will be added to the Tunnels screen.

Target Address

For an IBM tunnel, click Select. You get a list of firewall network objects. Select a network object or create a new one. Click OK. The address of that network object is entered in the target address field.

For a manual tunnel, click Select. You get a list of all of the network objects. Select a network object or create a new one. Click OK. The address of that network object is entered in the target address field.

Target User

If you selected dynamic tunnel, this option replaces Target Address. Click select and choose a network object from the User Network Object menu or create a new one. Click OK.

Target SPI

For a manual tunnel, specifies the security parameter index (SPI) value the tunnel partner will use. It is usually decided by the tunnel partner. All SPIs are 32 bit random numbers. They can be entered in either decimal or hex format.

Note that the SPI value 0 is reserved to indicate that no security association exists. The set of SPI values in the range of 1 through 255 are reserved to the Internet Assigned Numbers Authority (IANA) for future use.

Firewall SPI

Firewall Security Parameter Index is assigned by the firewall when you add a manual tunnel or use dynamic tunneling for the secure remote client. You cannot set or change this value.

Encryption Algorithm

For an IBM tunnel, specifies the algorithm used for IP packet encryption. If used, must specify either DES_CBC or Commercial Data Masking Facility (CDMF). Open the pull-down menu to choose from either CDMF, DES_CBC_4, or DES_CBC_8. DES_CBC_8 uses a 64 bit initialization vector and DES_CBC_4 uses a 32 bit initialization vector. Click OK and the encryption algorithm you chose is added to the Tunnels screen.

ESP Algorithm
For a manual tunnel, specifies the algorithm used for IP packet encryption. If used, must specify either DES_CBC or Commercial Data Masking Facility (CDMF). Open the pull-down menu to choose from either CDMF, DES_CBC_4, or DES_CBC_8. DES_CBC_8 uses a 64 bit initialization vector and DES_CBC_4 uses a 32 bit initialization vector. Click OK and the ESP algorithm you chose is added to the Tunnels screen.

Policy
Allows you to enter a combination of encryption and authentication values. Open the pull-down menu to select a particular policy. Click OK and your selection is added to the Tunnels Definition screen.

Session Key Lifetime
Specifies the time in minutes that an IBM tunnel will be operational. The current session key may be used. The value specified will affect performance (smaller value, bigger performance hit). Generally, this value should be smaller when CDMF is used as the encryption algorithm. Put a value in the entry field. The default is 30 and the maximum time allowed is 1440.

Tunnel Life Time
Specifies the time in minutes that a manual tunnel will be operational. The current session key may be used. The value specified will affect performance (smaller value, bigger performance hit). Generally, this value should be smaller when CDMF is used as the encryption algorithm. Put a value in the entry field. The default is 480 (8 hours) and the maximum time allowed is 99999. Note that for a tunnel connection with Secured Network Gateway 2.2, the maximum tunnel life time value that you can set is 44640.

Session Key Refresh Time
For an IBM tunnel, specfies the time in minutes between a new key start and an old key expiration. Put a value in the entry field. The default is 1 and the maximum time allowed is 720. This value is half or less than half of the session key lifetime.

Initiator
Identifies which partner starts the session negotiations. If both partners are identified as the "initiator," the tunnel logic will resolve the deadlock. At least one of the partners must be set as the initiator. Select either Yes or No.

3. Click OK and your entries are added to the Tunnels screen.

Modify a Tunnel

You cannot modify an active tunnel.

1. Select a tunnel from the Tunnels menu and click Open.
2. Modify the desired fields on the Modify Tunnel menu and click OK.

Delete a Tunnel

If the tunnel that you wish to delete is used in a rule or service, you must first eliminate the reference to the tunnel on the Services panel. To do this, double-click Traffic Control from the configuration client navigation tree. Double-click Connection Templates. Select Services. Double-click on the service you wish to modify on the Services List panel. The Modify Service panel appears. Click Select and choose a tunnel that is not used in any service, from the Select a Tunnel panel. Click Apply to place the tunnel ID in the Override Tunnel ID field. This will eliminate the original tunnel reference and you are now able to safely delete the tunnel.

Do not use SMIT to delete a tunnel that has been overriden by this configuration client process.

1. Select the tunnel you want to delete from the tunnels menu 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.

Note: If you delete a tunnel and then add it again, you have to reexport it.

Export Tunnel Definition Files

As a tunnel owner, after you have defined a set of tunnel context definitions, you will export one or more of these definitions to a tunnel partner.

1. Select a tunnel from the tunnels menu and click Export.
2. Enter a directory name in the field.

   The directory name is an arbitrary name that you create; it does not have to match anything. The directory is where the tunnel information is temporarily placed for export to a partner.
3. Move your cursor to the tunnel ID field (there is no input indicator for tunnel ID because it is not a writeable field).
4. Open the pull-down menu to get a list of tunnel IDs.
5. Select the desired item.

   One or more items can be selected.
6. Click OK after making all selections.

   The tunnel ID(s) that you have selected are added to the Export Tunnel Definition Files screen.

When you have completed this operation, the directory contains the names of the files that need to be moved or exported, (for example by using ftp or by creating a diskette) to your tunnel partner's machine. Because a directory name of `tmp` was used in this example, the format for the files to be exported would be:

```
/tmp/fwexpmctx  (for IBM tunnel)
/tmp/fwexpmctx.manual  (for Manual tunnel)
/tmp/fwexppolicy  (for migration of an IBM or Manual tunnel)
/tmp/fwexpmctx.3.1.1 (for a new installation of an IBM and Manual tunnel)
```

Import Tunnel Definition Files

1. Select a tunnel from the Tunnels menu and click Import.

   A dialog screen appears.

2. Enter the name of the directory where you have restored the files you have imported from the firewall.
3. Click Enter.

   All is the default if you do not want to choose from the list of tunnel IDs.

4. Open the pull-down menu to get a list of tunnel IDs.
5. Select the desired item. (One or more items can be selected).
6. Click Enter after making all selections.

Tunnel Activation Status

Use this procedure to activate or deactivate a tunnel(s). When you activate a tunnel, the IBM Firewall enables the use of that tunnel. Any filter rules, that reference the activated tunnel will be operable.

Activate a Tunnel

1. Select a tunnel from the Tunnels menu and click Activate.

   The Activate a Tunnel menu appears.

2. Open the pull-down menu to get a list of tunnel IDs.
3. Select the desired item.

   One or more items can be selected. The default is all, which gives you all of the tunnels in the list.

4. Click OK after making all selections.

Deactivate a Tunnel

Use this procedure to stop communication at an IBM or manual tunnel. The session key engine will continue to run.

1. Select a tunnel from the Tunnels menu and click Deactivate.
2. Open the pull-down menu to get a list of tunnel IDs.
3. Select the desired item.

   One or more items can be selected.

4. Click OK after making all selections.

   The tunnel ID to be deactivated is displayed.

Shutdown the Session Key Engine

Use shutdown only if you want to stop all tunnel activity for an extended period of time. Or use shutdown for security reasons if you need to stop all tunnel activity immediately.

If you have multiple tunnel partners, each partner machine must restart because of the shutdown.

1. Select a tunnel from the Tunnels menu and click Shutdown.

   A dialog screen appears.

2. The configuration client asks you to confirm the request.
3. Click Yes to shutdown.

Activating an IP Tunnel

You can also use the command interface to configure and activate an IP Tunnel. For more information on the tunnels commands, see the *IBM Firewall Reference.*

Chapter 14. Using the Windows 95 Secure Remote Client

This chapter describes how to install, configure, and use the Windows 95 secure remote client.

Windows 95 Secure Remote Client

With the Windows 95 secure remote client, a mobile user or a home user can access their network through a secure tunnel. The tunnel authentication is based upon the policy set by the administrator. The user first dials into their Point-to-Point Protocol (PPP) server.

Then the user logs on to the firewall. The user enters the firewall IP address, a user ID, and a password to establish a tunnel with the firewall. After the tunnel is established, all IP traffic goes into the tunnel. Once users make a connection, they have full TCP/IP access to whatever servers are behind the firewall and can use FTP, telnet, HTTP, and mail applications.

After the remote client user has connected to the PPP server, he or she can connect a tunnel, disconnect a tunnel, When the user selects Connect Tunnel, the configuration client starts a secure socket layer (SSL) control session with the firewall. The SSL Server application authenticates the remote client based on the ID and password, sends the remote client the tunnel policy, and then activates the dynamic tunnel, dynamic filters, and dynamic policy for the remote user on the firewall. Then the SSL control session is terminated.

The dynamic tunnel, dynamic filter(s), and dynamic policy remain active until the remote client disconnects the tunnel or the tunnel times out.

When the user selects Disconnect Tunnel, the configuration client starts another SSL control session with the firewall. The SSL Server application authenticates the remote client based on ID and password and deactivates the dynamic tunnel, dynamic filters, and dynamic policy for the remote user. Then the SSL control session is terminated.

When using the Windows 95 secure remote client, set the default route to the firewall. This enables the Windows 95 secure remote client's PPP Internet IP address to be routed back to the firewall so that is can return to the Windows 95 secure remote client in a tunnel.

If you have more than one firewall attached to the Internet, specify one of the firewalls as the designated firewall for the Windows 95 secure remote client. You must set the default route to this designated firewall.

Configuring the Firewall

The Windows 95 secure remote client supports a maximum MTU size of 500 bytes. You can configure the MTU size per interface in AIX by using SMIT. Select:

```
Communications Applications and Services
  TCP/IP
    Further Configuration
       Network Interfaces
            Network Interface Drivers
```

Select the nonsecure interface, and set the Maximum IP PACKET SIZE for THIS DEVICE to 500.

If you do not want to limit the MTU size of your primary nonsecure firewall interface to 500 bytes, then install another nonsecure interface and use this interface for the Windows 95 Secure Remote Client traffic. Set the MTU size of this interface to 500 bytes.

Before you can configure the Windows 95 secure remote client, perform the following steps to configure the IBM Firewall.

1. Create a user. Select Users from the configuration client navigation tree. Fill in the Add User menu fields in the following way:
   - Authority Level: Proxy User
   - User Name: New User ID
   - Nonsecure IP: password
   - Set the Nonsecure FTP and Nonsecure Telnet if the user needs access to FTP or telnet to the firewall.
   - Set the password by clicking the Password tab
   - Type in the new password
   - For all other fields, use the defaults
   - Save the user name and the password for the secure remote client configuration.
2. Create a single network object. Select Network Objects from the configuration client navigation tree. Double-click NEW.

   Fill in the following Add a Network Object menu fields for the first network object and refer to “Using the Configuration Client to Define Network Objects” on page 29 for explanations of these fields.
   - Object Type: User
   - User Name: Select the User you previously created
   - Description: Anything you would like
   - Filter Lifetime: 60 minutes
3. Create a tunnel definition. Select Traffic Control from the configuration client navigation tree. Double-click the file folder icon to expand the view. Select Virtual Private Network. Double-click NEW.

   Fill in the Add Tunnel menu fields and refer to Chapter 13, “Creating a Virtual Private Network” on page 95 for explanations of the following fields.
   - Tunnel Type: Dynamic Tunnel
   - Tunnel ID: Choose a value
   - Local Address: Nonsecure Interface
   - Target User: Select the user you previously created
   - Target SPI: Select a number over 256
   - Policy: Choose a policy from the pull-down menu
   - Encryption Algorithm: Choose a value from the pull-down menu
   - Session Key Lifetime: Select a value
4. Create a Connection. Select Traffic Control from the configuration client navigation tree. Select Connection Setup. Double-click NEW.

   Fill in the Add a Connection menu fields in the following way:
   - Name: SSL Connection
   - Source: The World

- Destination: Nonsecure Interface

In the Connection Services section of the Add a Connection screen, click Select. Apply the service "SSL Server - Permit SSL Server traffic to remote SSL agents".

Activate this connection. Select Regenerate Connection Rules and Activate. Click Execute.

5. Create the keyfile. Refer to the *IBM Firewall Reference* for information on how to do this.

Installing the Windows 95 Secure Remote Client

Install the Windows 95 secure remote client on a Windows 95 system. This system must have the Microsoft Dialup Adapter installed. See your Window's documentation for information. Microsoft's ISDN Accelerator Pack 1.1 is a prerequisite and must be installed first. If you do not have the Microsoft ISDN Accelerator Pack 1.1, go to the following Web page for information:
`www.microsoft.com/windows/software/isdn.htm`.

To install the Windows 95 secure remote client:

1. Start Windows 95.
2. Choose Start - Run.
3. Type **x:\win\ipsec\setup**, assuming the IBM Firewall CDROM is in the **x** drive, and where **x** is your CDROM drive.
4. Click OK.
5. Follow the instructions on the menu-driven installation program.
6. When setup is complete, click OK. You then get detailed instructions on how to install the device driver, as described in “Installing the Secure Remote Client Device Driver.”

Installing the Secure Remote Client Device Driver

The next step of the secure remote client installation is the installation of the required device driver from a Windows 95 workstation.

1. Opening the Network Icon
   a. Click Windows 95 Start.
   b. Click Settings.
   c. Click Control Panel.
   d. Double-click Network icon.
2. Adding an Adapter
   a. From the Configuration Page, click Add.

      A Select Network Component Type dialog is displayed.

   b. Select the Adapter entry in the list box.
   c. Click Add.

      A Select Network Adapters dialog is displayed.
3. Installing the device driver

a. Click Have Disk.

b. Enter **x:\win\ipsec\driver** in the Install from Disk dialog entry field, where **x:** is your CDROM drive.

c. Click OK.

d. Select the IBM Ibmisdn software network adapter, then click OK.

e. You get the Configuration Page again. Click OK.

f. You get the Resources Page. Click OK.

g. You get the ISDN Configuration panel. Follow the Wizard instructions as you continue the installation of the driver.

h. You do not need to configure the ISDN adapter or enter information like phone numbers during this driver installation. All dialogs have default choices. Select the defaults.

i. You will be asked to keep the newer version of WAN.TSP. Deny this by clicking NO and install WAN.TSP from the driver disk.

j. You will be asked to insert your Windows 95 installation CDROM. Insert the CDROM and click OK. Then copy files from **x:\win95**, where **x:** is your CDROM drive. Click OK.

k. You might get the "Version Conflict" dialog for the NDIS.VXD file. You have to keep the new NDIS.VXD file by clicking YES, as recommended by this dialog.

4. When prompted to restart your computer, click YES.

Configuring the Windows 95 Secure Remote Client

To configure the Windows 95 secure remote client (with a mouse):

1. If you do not have a modem installed, you need to install a modem using the standard procedure from a Windows 95 control panel, before the configuration is complete.
2. Click the PPPSEC icon in your IBM Firewall folder.
3. Click Line. Select Edit IPSEC Entry to configure this connection. Use the General tab of the edit dialog to enter the default PPP server access phone number in the Phone Number field.
4. If you have a fixed (static) IP address and/or fixed DNS IP address, click the Server Types tab and then click TCP/IP Settings. (We do not support WINS). Enter the alternate DNS server address and/or specify the fixed IP address for this client. Click OK.
5. Do not use any other configuration options on this panel.

Using the Windows 95 Secure Remote Client

To establish the point-to-point protocol (PPP) IP secure connection, you have to perform the following steps from the Windows 95 client.

1. Dial into the PPP server by clicking on the telephone icon or by clicking on Line and choosing Dial a Phone Line from the pull-down menu.
2. You get a Dial dialog. Fill in the Phone Number field. The other fields have default values, but you can change them. Click Dial. When a connection occurs you will see a small modem icon in the taskbar.

3. You then get a User Logon dialog. Fill in your User Name and the required authentication information for your PPP server. Click OK. When a PPP connection is established, a small PPP icon appears on the taskbar, and the Logon button is enabled.
4. Click Firewall. Select Logon. You then get a Firewall Logon menu. Enter your firewall IP address. Then enter your firewall user ID and password. Click OK.
5. Click Firewall. Select Start tunnel.

   After the tunnel is started, the IBM IPSec icon appears on the Taskbar and the Stop tunnel button is enabled.
6. You can start the tunnel and stop the tunnel again anytime as long as the PPP connection exists.
7. To hang up the connection, stop the tunnel first, then click Hangup or select the appropriate menu item.

   After hangup, the client will exit automatically.

Chapter 15. Using the AIX IPSec Client

The AIX IPSec client provides the capability for a client AIX machine to securely and cost effectively communicate with a firewall or another AIX IPSec client through a secure tunnel.

The AIX IPSec client provides an IPSec function for hosts. IPSec tunnel(s) can be created and used between clients or between a client and a firewall.

Manual key distribution is used for IPSec tunnels on the AIX secure client. Tunnels that are defined on an AIX secure client will communicate with manual tunnels on an IBM Firewall 3.1.1 and a Secure Network Gateway Firewall 2.2.

Installing the AIX IPSec Client

There is a separate installp image `ipsec.obj` for the AIX secure client. You can install it on a RISC/6000 machine with AIX 4.1.5 or 4.2.

The AIX IPSec client should not be installed on the same machine as the IBM Firewall because the IBM Firewall already has IPSec function. However, you can install the configuration client, configuration libraries, and report utilities on the same AIX machine as the AIX IPSec client. This allows you to perform secure remote configuration and logging. You need to reboot the AIX operating system after you successfully install the AIX secure client.

If the installation failed, run `installp -C ipsec` before reinstalling the AIX IPSec client.

Configuring and Managing the AIX IPSec Client

After you install the AIX secure client, you use SMIT to do configuration and administration of the tunnels. From the SMIT System Management menu, select IPSec Tunnel Management.

Adding a Tunnel

Tunnels can be created between an AIX client machine and another AIX client or an AIX firewall machine. To add a tunnel, select Add a Tunnel from the IPSec Tunnel Management menu and enter the following values:

Value	Description
Tunnel ID	Specify an integer from 1–99999. It has to be different for each tunnel defined or imported on this AIX client and different than tunnels already existing on the firewall if it is to be exported to the firewall or the other AIX client.
Local IP Address	Specifies an IP address in dotted decimal format. You can choose a value from the selection list.
Target IP Address	Specifies the tunnel partner's IP address in dotted decimal format.

Target SPI

Specifies the security parameter index (SPI) value the tunnel partner will use. It is usually decided by the tunnel partner. All SPIs are 32 bit random numbers. They can be entered in either decimal or hex format.

Note that the SPI value 0 is reserved to indicate that no security association exists. The set of SPI values in the range of 1 through 255 are reserved to the Internet Assigned Numbers Authority (IANA) for future use.

ESP Algorithm

Specifies the encryption algorithm used by the tunnel. Currently DES_CBC_4 (with a 32-bit initialization vector), DES_CBC_8 (with a 64-bit initialization vector), and CDMF (Commercial Data Masking Facility) are available.

AH Algorithm

Specifies the authentication algorithm used by the tunnel. Currently keyed MD5 is the only available algorithm.

Tunnel Life Time

Specifies the maximum time in minutes the tunnel can function. Select a value from 1–99999. A tunnel will stop working after its tunnel life time expires. The tunnel will have to be refreshed again when the time expires.

Protocol

Specifies the type of TCP/IP protocols to go through the tunnel.

Source Port (From)/ICMP Type

If the protocol is TCP or UDP, this field and the Source Port (To) field specify the source port range of the packets. These values must be non-negative integers. If the protocol is ICMP, then this field specifies ICMP type. The value must be a non-negative integer. If the protocol is ALL, this field and the Source Port (To) field specify the source port range of the TCP and UDP packets. All other types of packets will go through the tunnel.

Source Port (To)

If the protocol is TCP or UDP, this field and the Source Port (From) field specify the source port range of the packets. These values must be non-negative integers. If the protocol is ICMP, then this field is not applicable. If the protocol is ALL, this field and the Source Port (From) field specify the source port range of the TCP and UDP packets. All other types of packets will go through the tunnel.

Destination Port (From)

If the protocol is TCP or UDP, then this field and the Destination Port (To) field specify the destination port range. These values must be non-negative integers. If the protocol is ICMP, this field is not applicable. If the protocol is ALL, this field and the Destination Port (To) field specify the destination port range of the TCP and UDP packets.

Destination Port (To)

If the protocol is TCP or UDP, then this field and the Destination Port (From) field specify the destination port range. These values must be positive. If the protocol is

ICMP, this field is not applicable. If the protocol is ALL, this field and the Destination Port (From) field specify the destination port range of the TCP and UDP packets.

ESP, AH Mode/Combination

Specifies the tunnel mode and transport mode for ESP and AH.

Client-to-Client

After you add a tunnel definition on the client, then you export the tunnel definition to another client. Import it on that other client. Then, activate both clients.

Firewall-Side Services

Activate the predefined service for the AIX IPSec client, which only allows for encryption and authentication protocols. In order to have two-way tunnel traffic, you need to create a service to allow traffic over the tunnel. Refer to Chapter 6, “Controlling Traffic Through the Firewall” on page 37 for more information.

Client-to-Firewall

If you define a tunnel between an AIX firewall and an AIX client, you have to first define the tunnel on the AIX client and then export the tunnel definition to the AIX firewall. An AIX client cannot import tunnel definitions exported from an AIX firewall, because an AIX client requires more information for the tunnel policy.

If the tunnel partner is an AIX firewall, then only choose tunnel mode and not transport mode. Otherwise, the tunnel export files will be rejected by the AIX firewall.

1. Import the tunnel definition on the firewall.
   - From the configuration client navigation tree, select Virtual Private Network.
   - On the Virtual Private Network menu click Import. FTP or tar the exported files: `fwexpmctx, fwexppolicy, fwexppolicy.3.1.1` into an empty directory on the firewall.
   - Fill in the directory name, where you FTPed or tarred the export files.
   - In the Tunnel List field, click select to select the tunnel id being imported.
   - Click OK.
2. Use the Connection Templates to create a rule.
   - From the configuration client navigation tree, select Traffic Control. Double-click the file folder to expand the view. Select Connection Templates. Double-click the file folder to expand the view. Select Rules.
   - From the Rules List, select the "VPNs in non-secure" rule and click Copy to copy it. (To quickly find this rule, click Bottom).
   - Enter a rule name. It can be whatever you like.
   - Fill in the fields with information that pertains to this specific tunnel.
   - For Direction/Control, you probably want to select: both, both, and yes.
   - Select the manual tunnel ID associated with this rule.
   - Click OK.

- Click OK before exiting the Copy IP Rule menu.

3. Use the Connection Templates to create a service.
   - From the configuration client navigation tree, select Traffic Control. Double-click the file folder to expand the view. Select Connection Templates. Double-click the file folder to expand the view. Select Services.
   - From the Services List, select the "Remote Client - AIX" and click Copy to copy it.
   - Enter a service name. It can be whatever you like.
   - Under Service Composition, click Select. On the Select a Rule menu, highlight the rule you just created, and click Apply twice.
   - Click Cancel to leave the Select a Rule menu.
   - Highlight your service and click Flow. (This will give you arrows pointing in both directions.
   - Click OK.
   - Click Close to close the Services List menu.
4. Setup the Connection.
   - From the configuration client navigation tree, select Traffic Control. Double-click the file folder to expand the view. Select Connection Setup.
   - On the Connections List, double-click NEW.
   - Fill in the Name field.
   - Select the client for the Source field.
   - Select the secure or nonsecure interface for the destination field.
   - Under Connection Services, select the new service you created for your AIX IPSec client.
   - Click OK.
   - Click OK to exit the Add a Connection menu.
5. From the configuration client navigation tree, select Virtual Private Network.
   - Select your tunnel.
   - Click Activate.
6. Activation your connection.
   - From the configuration client navigation tree, select Traffic Control. Double-click the file folder to expand the view. Select Connection Activation.
   - Click Regenerate and activate connection rules.
   - Click Execute.

Other Tunnel Actions

From the IPSec Tunnel Management menu you can also perform the following:

- Modify a Tunnel - modify an existing IPSec tunnel's attributes.
- Delete Tunnel(s) - delete IPSec tunnels from the tunnel list. Only inactive tunnels can be deleted.

- List all the Tunnel(s) and Status - list the attributes and active/inactive status for all the tunnels.
- Activate/Refresh Tunnel(s) - Activate the local end of tunnel(s). (Once the key expires, you need to refresh it).
- Deactivate Tunnel(s) - Deactivate the local end of tunnel(s).
- Export Tunnel Definition(s) - Export tunnel definitions to export files. Four files will be created under the specified directory. The files `expmctx` and `exppolicy` can be imported by another AIX client. The files `fwexpmctx`, `fwexppolicy` and `fwexppolicy.3.1.1` can be imported by an AIX firewall.

  `fwexppolicy` is associated with the migration of an IBM or manual tunnel from the Secured Network Gateway V2R2. `fwexppolicy.3.1.1` is associated with a new installation of an IBM or manual tunnel.
- Import Tunnel Definition(s) - Import tunnel definitions from export files. The client can only import tunnel definitions exported by another AIX client.

Chapter 16. 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 system log file maintained by the `syslogd` daemon.

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 IP rules. Additionally, log files can be written to at a variety of severity levels; for example, debug, info, 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 and Archiving 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 `local4` facility, configurable only through SMIT. Only root has access to the `local4` facility. Alert messages go to the `local1` facility. Administrative audit log information goes to the `local0` facility.

For report utilities to function properly, it is important that only `local4` messages appear in their input files. No other facility should be directed to the same file as `local4`, so set syslog accordingly.

If you want to see alerts on the main configuration client panel, you have to direct your alerts to a file designated as a `local1` facility. Nothing else should be designated for that file.

The following priority levels are cumulative. Emergency captures only the most severe firewall events.

- Debug
- Information
- Notice
- Warning
- Error
- Critical
- Alert
- Emergency

It is suggested that you begin with the `debug` level until your firewall procedures are stable. Then you can change to `notice` or `error` to reduce the logging activity and the size of the system log.

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 menu appears displaying the set of log facilities currently enabled.

1. Select NEW from the Log Facilities menu and click Open to add a syslog entry to those currently enabled.

   The Add Log Facilities panel appears, as shown in Figure 45.

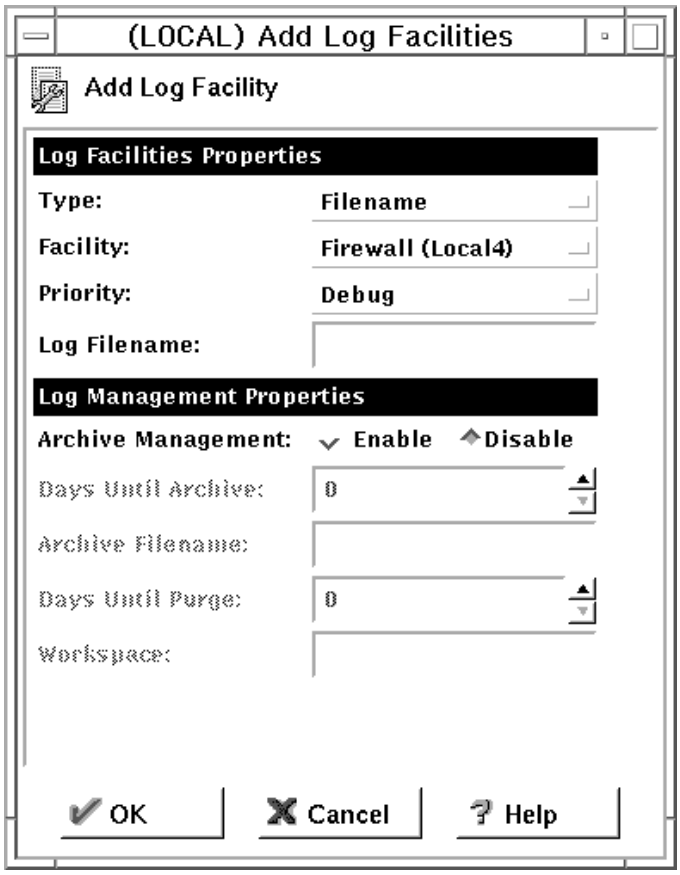

Figure 45. Add Log Facilities

   The log facility determines the type and source of information that will be logged.

2. Open the pull-down menu to select the type. Type can be either Filename, Hostname, or User ID.

   Note: If you choose hostname, you will be prompted for the TCP/IP host name of the machine that you want to send the log information to. If you specify a host hame, either DNS must be enabled on the firewall machine so that the host name can be resolved, or the host name you specify must be defined in the HOSTS file.

3. The log facility determines the type and source of information that gets logged. Open the pull-down menu to select one of the following log facilities:

   - Firewall (Local4) - general firewall logs, including filter logging
   - Log monitor (Local1) - log monitor daemon status and threshold violation warnings used to populate the Alerts Display
   - Mail - mail logs
   - Syslog - is especially useful in case the other logs fill up their file systems. Be sure to set the output "Log Filename" to `/dev/console`, or to a separate file system.
   - All Facilities

4. Open the pull-down menu to choose the priority. The logging priorities are listed in order of increasing severity. The priority you select will be the minimum level to be logged. Choose from the following log priorities: debug, information, notice, warning, error, critical, alert, or emergency.
5. Do the following:
   a. Fill in the log filename. The log filename must have an absolute path (beginning with a forward /) and the path to the file name must exist.
   b. Or, redirect the log output to another machine by entering @hostname.

      In order for this to work, you must enable the appropriate log facilities on the target system as well.
   c. Or, redirect the log output to a user ID in the local system. We recommend that you do not output `local4` to a user ID because it is not in an easily readable format and because the volume of messages sent to the user could be very high.
6. Archive management can be used with a filename type log facility only. When enabled, the log file can be compressed on a periodic basis. You can either enable or disable archive management parameters. Click the appropriate choice.
7. Select the number of days until archive. The number of days until archive must be zero or greater.
8. Enter an archive filename.
9. Select the number of days until the purge. The number of days to keep the log files must be a minimum of zero days. Log management does not include the current day when calculating the number of days to keep.
10. Enter the workspace.

    Log management requires temporary work space to run an effective log management process. The work space made available to log management should be at least equal to that of the largest log file being managed.
11. Click OK.

Change Log Facilities

1. Select the syslog entry you want to change from the Log Facilities menu and click Open.

   The Change Log Facilities panel will appear.
2. Change the desired fields. See “Add Log Facilities” on page 115 for an explanation of the fields.
3. Click OK.

Delete Log Facilities

1. Select a syslog entry from those currently enabled on the Log Facilities menu 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. This does not delete the actual log file.

Archiving Logs

With this release of firewall, it is suggested that you start with fresh log and archive files. Do not try to use log report utilities on any log files from previous firewall releases. To start a log management program to archive accumulated logs you have two options:

1. Run the `fwlogmgmt -l` command from the command line
2. Set up the `fwlogmgmt -l` command in the crontab.

The `fwlogmgmt -a` command purges the archives.

Note that when using the `fwlogmgmt -l` command, which compresses logs and puts them in the archives, if you receive message `ar0707-106`, it means that you have named a 0 length file as your archive log. Choose a different archive log name.

The most efficient or convenient means of running the log management process would be to set it up as a cron job. This periodically executes the log archiving process at a predetermined frequency. Root must set up the crontab file and determine the frequency of execution for the log management archive functions.

For example, if you want to set up the log management archiving process to run at 3:00 AM every day, type `crontab -e` and add the following line:

```
0 3 * * * fwlogmgmt -l
```

If you want to purge the archives every day at 3:00 AM, type `crontab -e` and add the following line:

```
0 3 * * * fwlogmgmt -a
```

For a more detailed crontab example, see the *IBM Firewall Reference.*

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 console when you run the `fwlogmgmt` command from the command line. If a crontab entry is used to initiate the process, then the root user is notified via standard AIX mail facility.

Report Utilities

You can use the report utility function to assist you in generating reports from the log and archive files. The purpose of report utilities is to generate tabulated files of administrative information. Tabulated means files are organized and formatted for easy mapping to relational database tables. These tables assist 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 DB2/6000. The administrator can then use the Structured Query Language (SQL), or other tools like IBM's Visualizer or Query Management Facility to query the data and generate reports.

AIX su logs, generated by the su (switch user) command, can be imported into the database in a similar fashion.

Report Utilities are installed as part of firewall install. They can also be separately installed and run on a non-firewall AIX host. The configuration client can be used to run them on a firewall. On a non-firewall, you will use SMIT or command line.

For report utilities to function properly, it is important that only `local4` messages appear in their input files. No other facility should be directed to the same file as `local4` , so set syslog accordingly.

(Do not try to use report utilities on any log files from previous firewall releases.) See the *IBM 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 panel appears, as shown in Figure 46.

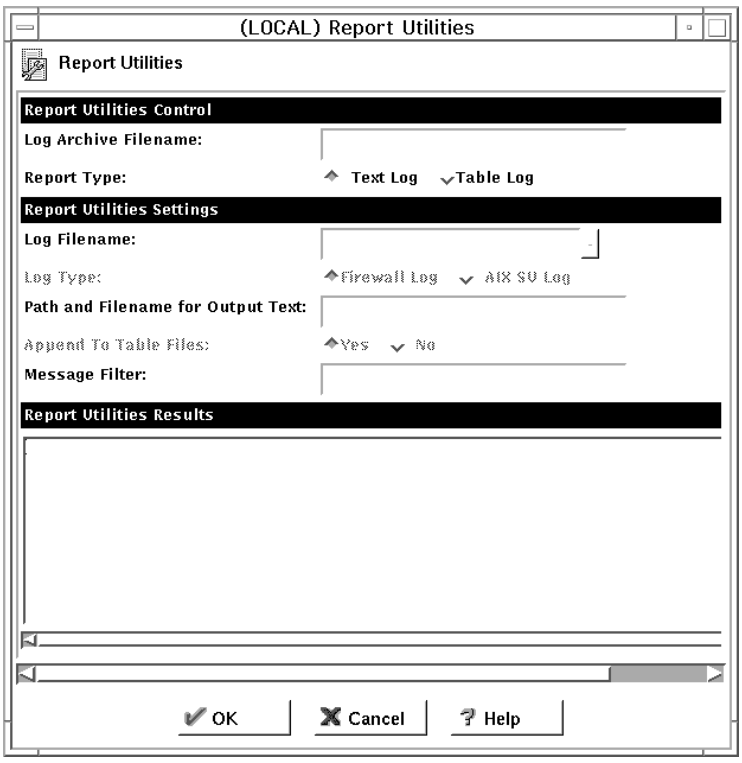

Figure 46. Report Utilities

1. The log archive filename is the archive file that contains compressed log files. Enter the archive filename that you created using Log Facilities in the Log Archive Filename field. Enter the absolute path name to the archive file. If you want to view a log file that is not archived, leave this field blank.

2. Select the Report Type. To view 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 Firewall Reference* for more information.
3. The log filename is any one of the compressed archived log files or other valid `local4` logs or the name of a su log file. If you made an entry in the log archive filename field, you can select the button in the Log Filename field to choose which log to work with. If you do not enter a log archive filename in step 1, the log file name you enter here must be the name of a valid, uncompressed firewall log file or a su file log. 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. Enter an AIX 'regular expression' in the Message Filter field. This is used to filter the set of messages for which you want to see the full text. The 'regular expression' must be one that is suitable for use with a 'grep' command. If it is not, you will get unexpected results or error messages. If you leave this field empty, all messages in the log will be placed in the Output Text file. The following are examples:

```
Regular Expression       What it Does

   ICA0                 shows log monitor threshold alert messages
   ICA3                 shows Socks messages (#ICA3000 - 3999)
   ICA[23]              shows proxy and Socks messages
   ICA2010              only shows occurences of the ICA2010 message
```

8. Clicking OK produces the requested file(s) in the specified output directory. The text file can be viewed using the firewall log viewer or a file browser of your choice. The tabulated files can be loaded into a database as described in the *IBM Firewall Reference.*

Chapter 17. Monitoring the Firewall Logging

This chapter describes how to monitor the logging in real time. An alert is generated when a configured threshold is violated.

The Log Monitor daemon, `fwlogmond`, monitors the messages sent to syslog for potential crisis situations, based upon user-defined thresholds. In the event of a threshold violation, `fwlogmond` 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 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
2. Authentication failures against any particular user ID
3. Authentication failures originating from any particular host
4. Occurrences of a message tag in the log

   The firewall administrator configures this type of threshold based on message message tags. Tags ICA2000 and ICA2001 apply to authentication failures and are ignored by the message tag threshold class.

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 `fwlogmond` daemon.

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 syslog `local1` facility configurable through the configuration client, SMIT, or the command line.
   - In the `local4` log in tabulated database format used by Report Untilities.
2. Mail to a list of users, through sendmail
3. Pager, as configured. See “Pager Notification Support” on page 123.
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.
```

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 menu, you can add, change, or delete a threshold definition. Note that delete does not mean delete from the log file. It means delete the definition.

Add Log Monitor

To add a threshold definition, select NEW from the Log Monitor Threshold Administration menu and click Open. The Add Log Monitor menu appears. Fill in the following fields:

1. Select class type by opening the pull-down menu 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 menu and click Open. The Change Log Monitor menu 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 menu and click Delete. You will be asked to confirm the deletion. Click Yes to confirm.

Pager Notification Support

This IBM Firewall function pages the system administrator when there are intrusion alerts on the firewall. To set up pager notification support, you need to configure the following three pager services.

1. Command Customization - This service must be created and modified using the configuration client. This service will contain a unique entry that defines the pager environment. See “Command Customization” on page 125 for more information on defining and customizing this service.
2. Carrier Administration - You must define a suitable carrier before connecting your modem. This service contains a list of three default carriers used in the U.S. If the carrier you are using is not one of these three, then add your carrier in this service. See “Carrier Administration” on page 127 for more information.
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 service contains a list of modems that you can choose from. See “Modem Administration” on page 128 for information on maintaining modem definitions.

Note: IBM Firewall supports the TAP communications protocol for pager notification support.

What Carriers and Modems are Supported

The carriers database file contains a list of all the carriers and their related transmission parameters. Some of the parameters besides the carrier name and modem phone numbers are:

- Dual Tone Multi Frequency (DTMF) phone number
- The maximum message length and digits for an alphanumeric pager
- The maximum number of blocks per transaction
- The maximum number of transactions per call
- The baud rate, parity, data and stop bits length

The carriers shown in Table 6 on page 124 are based in the United States and work with pager notification support.

Table 6. US Carriers

Carrier Name	Modem Phone Number
Skytel	1-800-759-6366
PageNet	1-800-720-8398
MobileCom	1-512-478-4875

Note: These carriers use TAP (Tele-AlphaNumeric Protocol). If any carrier does not use TAP, then the pager daemon will fail. Before using a particular carrier, make sure that the carrier uses TAP.

The pager code comes with six default modem definitions. These are:

1. Generic Hayes-compatible
2. US Robotics 14400 bps Sportster
3. US Robotics 9600 bps Courier
4. IBM 5853 2400 bps
5. IBM 7855 2400 bps
6. IBM 7852 28800 bps

Configuring Your Serial Port

Before using your pager, you need to configure your serial port. If you have already defined a TTY to the system and wish to use it for pager dialing, then do the following:

1. Enter SMIT on the command line
2. Select the following menu items:

```
Devices
  TTY
    Change / Show Characteristics of a TTY
```

3. Select the TTY you wish to use from the list of available TTYs.
4. Ensure the following fields are set with these values:

```
Enable LOGIN    = disable
BAUD rate       = 9600
BITS per character = 8
Number of STOP BITS = 1
```

5. Click Enter.

If you have not previously defined a TTY, then perform the following steps:

1. Enter SMIT on the command line.
2. Select the following menu items:

```
Devices
  TTY
    Add a TTY
```

3. Select tty rs232 Asynchronous Terminal from the list of TTY types.
4. Select the desired serial port from the list of available serial ports.
5. Set the following fields with these values:

```
PORT number   = (desired port number)
Enable LOGIN  = disable
BAUD rate     = 9600
BITS per character = 8
Number of STOP BITS = 1
```

6. Click Enter.

Configuring Pager Notification Support

Pager Setup is used to configure the command customization file and to maintain carriers and modems. If you are using pager, you should use Pager Setup to customize your pager environment before using Log Monitor.

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
- Assign a priority
- Define a pager type and ID
- Write a pager message

Command Customization Settings

When you select Pager Setup from the navigation tree you get a Pager Setup menu with Command Customization Settings similar to the menu shown in Figure 47 on page 126.

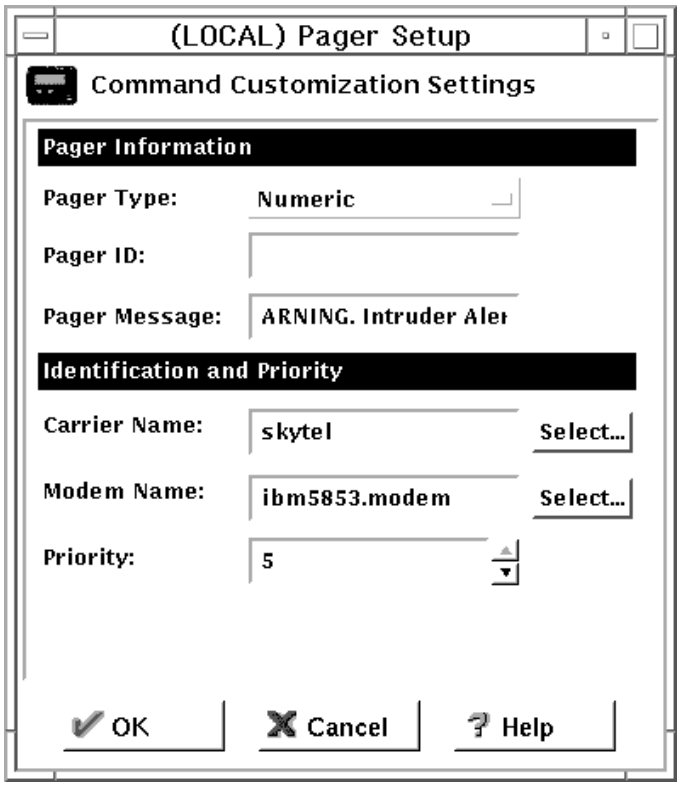

Figure 47. Pager Setup

Type or select values in the entry fields to be added.

1. Enter the pager type or open the pull-down menu to select from the list. Valid values are Numeric or Alpha (alphanumeric).
2. Enter the pager ID. This is usually a unique PIN number assigned to your pager by your carrier company.
3. Enter the pager message. 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 for alphanumeric pagers or your message might be truncated. Do not use a colon (:). If you do, it will be replaced by a blank space character.
4. If there is no carrier name, click Select to define a carrier. You will get the Pager Carrier Administration menu. See “Carrier Administration” on page 127 for details on how to fill in this menu.
5. If there is no modem name, click Select to define the modem. You will get the Pager Modem Administration menu. See “Modem Administration” on page 128 for details on how to fill in this menu.
6. Enter the priority for sending the page or use the slide control to select a priority. The highest priority is 5 (default) and the lowest priority is -1.
7. Click OK.

Change Command Customization

When you select Pager Setup from the navigation tree you get the Pager Setup menu 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 menu or the Pager Modem Administration menu 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 menu.

If no customization entry exists, then pager notification support will not be able to send a page.

Carrier Administration

From the Pager Setup menu, go to the carrier name field and click Select. You get a Pager Carrier Administration menu similar to the menu shown in Figure 48.

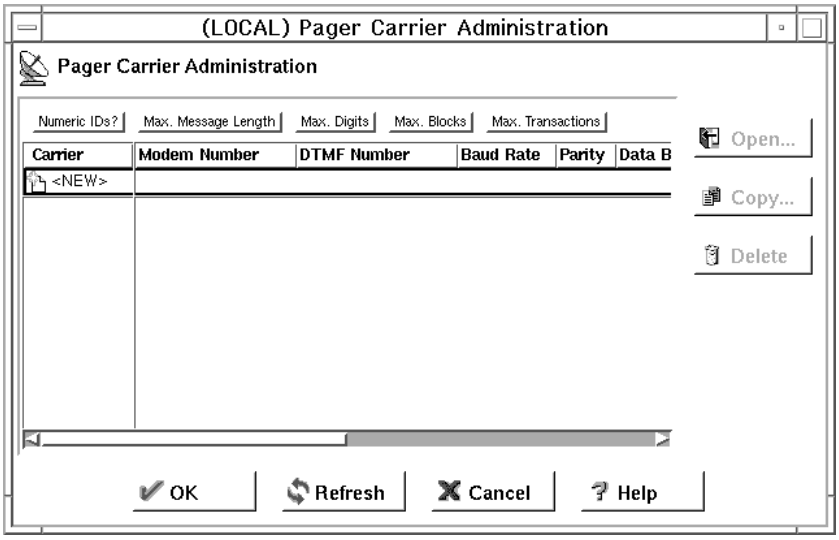

Figure 48. Pager Carrier Administration

Add a Carrier

To add a new carrier select NEW on the Pager Carrier Administration menu and click Open. Type or select values in the appropriate entry fields:

1. Enter the carrier name.
2. Enter the modem phone number. The digits of this phone number can be separated by a hyphen for clarity.
3. Enter the Numeric ID field value. Click (Yes) for numeric pagers or (No) for alphanumeric pagers. This field determines whether or not the paging carrier allows numeric IDs to be used during a data connection on the modem line.

4. Enter the Dual Tone Multi Frequency (DTMF) phone number. It is the carrier-wide touch-tone input number. This field is optional and can be blank.
5. Enter the Alphanumeric Pager field value. Click (Yes) for alphanumeric pagers and (No) for numeric pagers.
6. Enter the maximum message length.
7. Enter the maximum digits for the alphanumeric pager. (The length of the pager ID must be less than the maximum digits specified in this field).
8. Enter the maximum blocks per transaction.
9. Enter the maximum transactions per call.
10. Enter the baud rate. Open the pull-down menu and choose a value from the list.
11. Click Even, Odd, or None for the parity field.
12. Choose the default data bits; click either 7 or 8.
13. Choose the default stop bits; click either Yes or No.
14. Click OK.

Change Carrier

1. Select the carrier you want to change from the Pager Carrier Administration menu and click Open.
2. Refer to “Add a Carrier” on page 127 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 menu and click Delete.
2. You will be asked to confirm the deletion. Click YES to confirm.

   Note: The carrier database file must always contain at least one carrier. If no carriers are defined, then pager notification support will fail.

Modem Administration

From the Pager Setup menu, go to the modem name field and click Select. You get a Pager Modem Administration menu similar to the menu shown in Figure 49 on page 129.

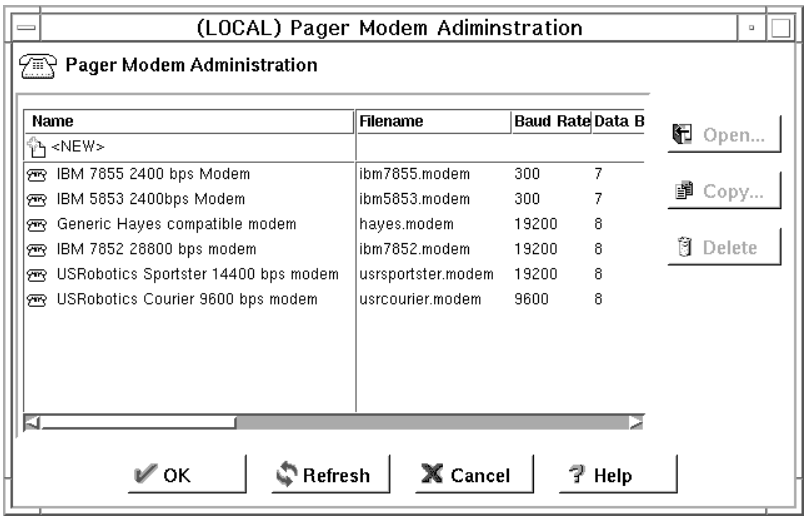

Figure 49. Pager Modem Administration

You can add, change, or delete various modems using this menu.

Add a Modem

To add a new modem definition file, select NEW from the Pager Modem Administration menu and click Open. On the Add Modem menu, type or select values in the entry fields.

1. Enter the modem filename. This must end with a .modem extension.
2. Enter the modem name.
3. Enter the initialization string. The characters in this field are sent to the modem in command mode to initialize the modem. The initialization string selected should set your modem to do the following:
   - Upon drop of Data Terminal Ready (DTR), the modem should hang up, not reset and return to the command mode.
   - Give verbal response codes to commands. These responses should correspond to those in the Valid Command Response and Valid Connect Response fields of the modem file in use.
   - Set the modulation speed either at the DTE speed (or the speed of the last command) or set to automatically detect the other modem's speed. No matter which modulation speed is used, the DTE speed must not be changed. If your modem is configured to auto-adjust the modulation rate, and the DTE is sending commands to your modem at 1200 baud (the rate given in your Paging Carrier's database record), but the modem actually connects at 300 baud, your modem will be expected to buffer the speeds so that the DTE can remain at 1200 baud.
   - Your modem might be set to echo characters while in command mode. If so, the modem file must indicate this in the Does Modem Echo Local field.
   - The modem should not echo characters while in connect mode.
   - The initialization string should include the command to hang-up the modem and disable Auto-Answer.

4. Enter the command mode string. This field should contain the set of characters that should be sent while in connect mode. This forces the modem into command mode without hanging up.
5. Enter the command terminator. This field indicates the character that should be appended to the end of all command sequences to force the modem to accept the command. Normally this is just a carriage return. (If you are using a backward slash, put another backward slash before it, for example, use \\r for \r.)
6. Enter the hangup command. This field should contain the command to force your modem to hang-up after dialing. The default that works with most modems is ATH0.
7. Enter the valid command response. This field should contain the string that allows your modem to accept commands. Normally OK is sufficient.
8. Enter the valid connection response. This field should contain the string that your modem will output when a carrier has been detected and a connection has been made. Most modems use CONNECT.
9. Enter outside line number. This field should contain the outside line number used to access the outside exchange. Usually, this will be followed by a "p" to notify the modem about a temporary pause. If you do not have an outside line #, use "p" only or enter the number followed by "p" as in the example 9p.
10. Click Yes or No. If Yes, the modem will echo local characters while in connect mode.
11. Enter the dial command. This is the command sent to the modem in command mode and followed by the Outside Line # field and the paging carrier's phone number. The default is ATDT which works for most modems.
12. Enter the dial pause. This field should contain the character used in a dial string to force your modem to wait for a short period of time (about 1 second) before continuing with the dial string. This is normally a comma (,).
13. Enter the dial number. This field should contain the character used in the dial string to force your modem to dial the touch tone corresponding to the # sign. This is normally just the pound sign (#) itself.
14. Enter the dial *. This field should contain the character used in the dial string to force your modem to dial the touch tone corresponding to the * sign. This is normally just the asterisk (*) itself.
15. Enter Return to command mode after dial. This field should contain the character to append the dial string in order to force the modem back into command mode after completing the dial string. The default is a semicolon(;), which works with most modems.
16. Enter the default baud rate. This field should contain the default baud rate for the modem. Open the pull-down menu to choose from a list of valid values.
17. Enter the default data bits. This field should contain the default data bits for the modem. Click either 7 or 8.
18. Enter the default stop bits. This field should contain the default stop bits for the modem. Click either 1 or 2.
19. Enter the default parity. This field should contain the default parity for the modem. Click either Even, Odd, or None.

20. Enter the default device. This field should contain the default device. This device file must exist under the /dev directory and should match with your configured serial port.
21. Click OK.

Change Modem

1. Select a modem name from the Pager Modem Administration menu and click Open to change a modem definition file.

   On the Change Modem menu you will see a list of fields you can change for the modem definition. Refer to “Add a Modem” on page 129 for explanations of these fields.
2. Click OK.

Delete Modem

1. Select a modem name from the Pager Modem Administration menu 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 syslog 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 syslog files, see Chapter 16, “Managing Log and Archive Files” on page 115.

Chapter 18. Using the File System Integrity Checker

Use the file system integrity checker to monitor changes to vital firewall or system files. If those files are inadvertently or maliciously modified, the security of the entire internal network may be compromised. The IBM Firewall maintains a database which contains:

1. A list of files considered sensitive.
2. The MD5 checksum of each file
3. The MD5 checksum of each file's access control list, which contains:
   - Attributes (setuid, setgid, and sticky bits)
   - Base permissions
     – owner's ID and mode
     – group's ID and mode
     – other's ID and mode
   - Extended permissions

The file integrity checker uses the AIX command `aclget` for permissions data.

When executed, the checker compares the current system status against the database. In the event of a discrepancy, the checker sends an alert listing the files that have been changed. You are notified of file modification, creation, and permission changes only.

`/etc/security/fwfschck.db.list` contains the list of sensitive files, which is used to generate the database. You can add additional files to this list.

Configuring File System Integrity Checker Using the Configuration Client

Select System Administration from the configuration client navigation tree. Double-click the file folder icon to expand the view. Select File System Integrity Checker. The File System Integrity Checker panel appears, as shown in Figure 50 on page 134.

1. To execute the standard mode and run the checker, click Check System Files Against Last Saved Database Copy.

   You will see the results displayed on the menu.

2. To update the database to reflect the current system status, click Update Database to Reflect Current System Files.

   The updated files are displayed on the menu.

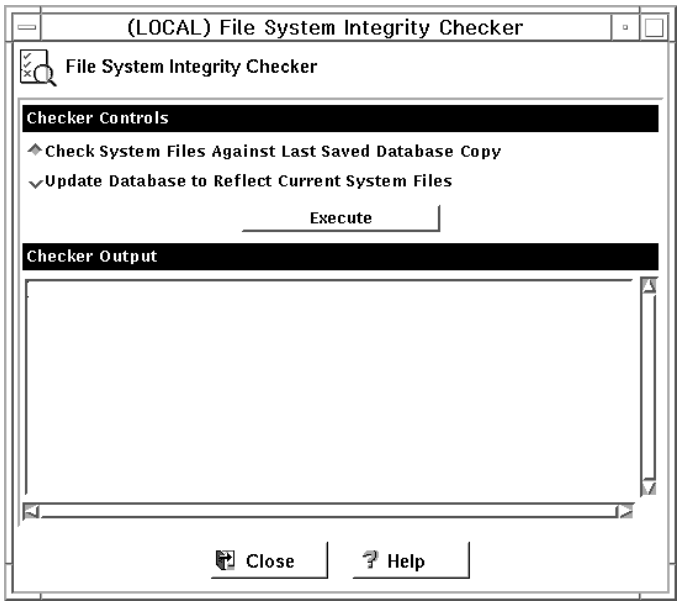

Figure 50. File System Integrity Checker

Setting Up the File System Integrity Checker as a Cron Job

Run the `fwschk` command on a regular basis. You can run the checker from the configuration client or the command line, but it is more convenient to automate it, so that the system runs at predefined times. As user root, type `crontab -e` at the command line to edit cron entries.

The following example causes the system to run the checker every day at 3:30 AM, sending output to the log file.

```
30 3 * * * /bin/fwfschk -l
```

If the file system integrity checker fails, it logs a message that is by default in the log monitor thresholds.

Chapter 19. Supporting the RealAudio Protocol

RealAudio protocol is a special protocol developed by Progressive Networks, which supports live and on-demand audio from the Internet. In the recommended configuration, the protocol requires two connections. The first connection is a TCP connection from the RealAudio player to the RealAudio server. After this connection is established, the RealAudio server optionally establishes a UDP channel back to the player. If the RealAudio server is TCP only, no further action is required by the firewall. In the scenario where UDP is used, the UDP connection is dynamic in the sense that the destination port number is dynamic.

The IBM Firewall supports the RealAudio protocol by monitoring and identifying these RealAudio TCP connections. Once a connection is identified, a dynamic filter rule for a UDP packet will be defined. This filter rule will be removed once the RealAudio TCP connection is closed. This is transparent to the RealAudio user. No extra configuration or knowledge is needed.

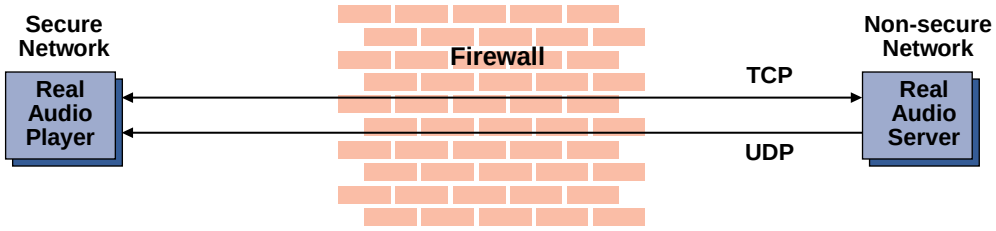

Figure 51. RealAudio Connections through the IBM Firewall. Once a RealAudio TCP connection is detected, the back channel UDP packet from the RealAudio server to the RealAudio player will be permitted to pass through the firewall as long as the TCP connection is active.

Configuring RealAudio Using the Configuration Client

To configure RealAudio, select System Administration from the configuration client navigation tree. Double-click the file folder icon to expand the view. Select Real Audio. The Real Audio menu appears.

1. Fill in the server port number for RealAudio. The RealAudio default server port number is 7070. However, you can reconfigure it to any valid TCP port number.
2. Fill in the maximum concurrent sessions allowed for RealAudio. The default is 10. It can be any non-negative integer.
3. Click OK.

RealAudio Web site

You can find more information on RealAudio at: **http://www.realaudio.com**.

Chapter 20. Using the Network Security Auditor

Use the Network Security Auditor to scan your network for security holes or configuration errors. The Network Security Auditor scans your servers and firewalls for a list of problems or vulnerabilities, such as open ports and other exposures, and compiles a list so you can correct problems. The Network Security Auditor can be used as a periodic scanner of critical hosts or as a one-time information gathering tool. With the Network Security Auditor, you maintain vigilance over your firewall. The Network Security Auditor is enabled to audit hosts on the local Class C network.

Figure 52 shows a Network Security Auditor sample output.

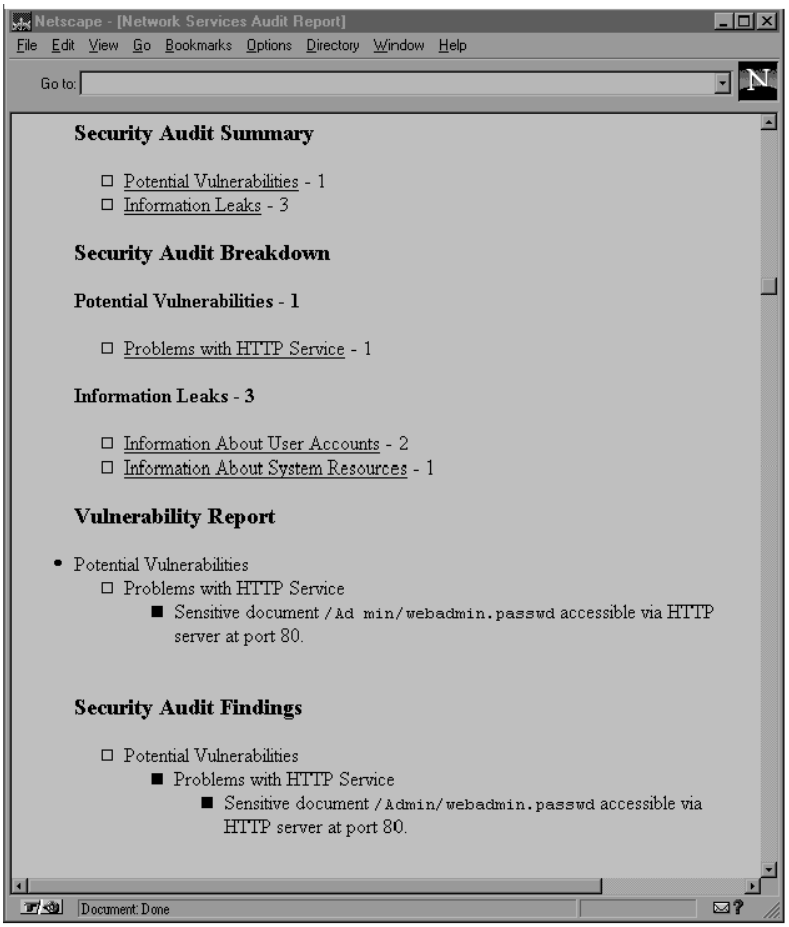

Figure 52. Network Security Auditor Sample Output

Features of the Network Security Auditor include:

- Scanning TCP and UDP ports
- Recognizing servers on non-standard ports
- Reporting dangerous services, known vulnerabilities, obsolete server versions, and servers or services in violation of customized site policy
- Generating reports in HTML for easy browsing

The results of the audit can be stored in a database for use in future report generation, or for immediate report generation.

Using the Network Security Auditor GUI

When you type `nsauditor` on your command line, you bring up the Network Security Auditor GUI, which allows you to perform the following:

Set up the Network Security Auditor
Prepare the Network Security Auditor to execute a new scan
Produce reports from previous Network Security Auditor scans
View the Man Pages

What the Network Security Auditor Can Do

Network Security Auditor allows an administrator to find out what network services are being offered, while at the same time reporting any of these services that are inherently insecure, or that contain known vulnerabilities. Network Security Auditor does not depend on advance knowledge about where network services should be found. Instead, this information is used only as a hint.

Network Security Auditor attempts to verify that the service is indeed active on the expected port. If the service does not behave properly, Network Security Auditor is often able to determine what actual service is on the port. Once the service has been ascertained, all vulnerability checks for that service are performed. In addition, Network Security Auditor is able to determine services that are on ports that have no predefined standard service. This means that Network Security Auditor, for example, is able to locate and test HTTP servers that are on any TCP port.

Currently, Network Security Auditor is able to locate and recognize the following network services for TCP:

finger
FTP
HTTP
IMAP
NNTP
POP
SMTP
SSH
SSLv2
telnet

Currently, Network Security Auditor is able to locate and recognize the following network services for UDP:

SunRPC (any)
FSP

In addition to the TCP services that can be recognized on any port, Network Security Auditor verifies the following services for TCP:

auth
chargen
daytime

discard
echo
netbios-ssn
printer
quotd
rexec
rlogin
rsh
time
X11
writesrv (AIX)
cppbrowse (AIX)

In addition to the UDP services that can be recognized on any port, Network Security Auditor verifies the following services for UDP:

chargen
daytime
discard
echo
FSP
Kerberos
netbios-ns
RIP
SNMP
syslog
talk
TFTP
time
timesync
SRC (AIX)
bootp

Network Security Auditor verifies the following services for SunRPC:

nfsmount
portmap
rusers
bootparam

The verification that these services are indeed active allows Network Security Auditor to find unauthorized network services. In support of this, Network Security Auditor allows the administrator to define policies. Policies can be defined for what TCP/UDP ports should be visible or active (for checking filtering rules), as well as what network services are allowed to be active. Policy violations are grouped together and reported separately.

In addition to service recognition, Network Security Auditor attempts to determine the vendor version of the server providing the service. Known vulnerable versions will be flagged during the scan. Vulnerable version information can be easily extended by the administrator as new information comes out.

As mentioned, once the service has been recognized, all the security checks for the service are then performed. This means that HTTP security checks will be per-

formed on all HTTP services found on a machine, no matter what port they are on. Following is a list of some of the checks currently performed:

HTTP	Dangerous CGI programs (phf, etc... configurable) Dangerous files (passwd files, etc... configurable)
SMTP	Dangerous commands (configurable), Dangerous alias (configurable) EXPN/VRFY information leak, Remote execution of commands
FTP	Guest flag check, Anonymous FTP - writable files
telnet	ENV opt (dynamic linker bug), weak passwords
rlogin	-f root
rsh	Bypass password (hosts.equiv or .rhosts problems)
rexec	weak passwords
SunRPC	Report any dangerous SunRPC services (configurable)
bootparamd	Get NIS domainname
NFS	Report exported filesystems
SMB	Report Share list, Flag filesystems shared to everyone
SNMP:	Guessable community string (read or write)
TFTP:	Allow system files to be grabbed?
Kerberos	Leak principles and realm?
X11	Open X server?
IMAP/POP	Buffer overflow

Network Security Auditor presents the information in an easy to view format. Network Security Auditor comes with four different report templates. It is possible for the user to define their own report templates to suit their taste.

Reports can be generated as HTML documents. When viewed with a browser, the report can be navigated through links within the report. In addition, the report can contain links to external information sources, such as CERT advisories.

Targets

Network Security Auditor runs in two modes.

1. Scanning mode
2. Reporting mode

When in scanning mode, targets indicate systems that should be scanned. Targets can be specified as a hostname, IP address in dotted-quad form, an IP address range, or a IP network/netmask form. Only targets that are on the same subnet as the host running NSA can be scanned.

The hostname form's behavior changes if a hostfile has been specified. In this case, the hostname must match an entry in the current hostfile. It can match either the tag field or the hostname field in the hostfile. If a hostfile is specified, and there are no targets indicated for it, then all entries in the hostfile are targeted.

When in reporting mode, a report will be generated for recorded scans selected by the target. The syntax for the target is host@run-specifier. Both host and run-specifier are optional. If no host is specified, then all hosts in the current database are selected. If no run-specifier is specified, then the last scan of the selected hosts is selected. This means that if no targets are specified, the last scan of all hosts in the current database will be selected.

Scans can be selected by either direct reference, or by the date of the scan. Direct references are done as a numerical index. The index value of 0 (zero) selects the last scan. Positive indices index from the first scan to the last, thus an index value of 1 (one) selects the first scan. Negative indices index back from the last scan. An index value of -1 (negative 1) selects the next to last scan. Index values out of range (positive or negative) will select the last scan.

Date specifications for selecting scans are broad. Almost any form of specifying a date can be used. You will probably need to include the date specification in quotes to protect it from the shell. All scans indicated by the date will be selected. Thus, a specification of '@May 1997', will select all scans for all hosts performed in May 1997. A specification of 'myhost@last week' will select all scans of myhost performed the previous week (a week being defined as starting on Sunday morning).

Options

Network Security Auditor has many options for specifying run time options. The ordering of options is important. Configuration class options should be specified first. Options affecting the scan of an individual host, apply to all hosts specified after that option. Note that an option that affects scan type definitions will not affect the scan type definitions of hosts listed before that option.

Configuration Options

The following are configuration options for the Network Security Auditor:

directory (basedir) Specify the directory that contains configuration files. The default is `/etc/nsa`.

configfile (configfile)
Specify configuration file. If the filename does not start with a slash (/), then the file is looked for in the base directory. The default is config in the base directory.

scannerdata (scandatafile)
Specify data file containing information used by scanners. The default is scannerdata in the base directory.

scannerdefs (scandefsfile)
Specify file containing definitions of scanner types. If the filename does not start with a slash (/), then the file is looked for in the base directory. The default is scannerdefs in the base directory.

rpcnamefile (rpcnamefile)
Specify a Sun RPC name/number mapping file. If the filename does not start with a slash (/), then the file is looked for in the base directory. The default is rpc in the base directory.

xmitnames (namefile)
Specify a file containing a mapping of names to speeds for the xmitrate and xmithost options. The default is `xmitrates` in the base directory.

referdata (referencefile)
Specify the file that contains mappings of references keys to the actual location of the data. The default is references in the base directory.

services (servicesfile)
Specify a TCP/UDP name/number mapping file. If the filename does not start with a slash (/), then the file is looked for in the base directory. The default is services in the base directory. Note that this file does not have the same format as the `/etc/services` file on UNIX systems.

messagecat (messagecatalog)
Specify the message catalog containing messages used in report generation. The default is messages.cat in the base directory.

templates (searchpath)
Specify a colon (:) separated list of directories to search for report templates. If elements of the path do not start with a slash (/), then the base directory is prepended to the path.

policies (searchpath)
Specify a colon (:) separated list of directories to search for policy definition files. If elements of the path do not start with a slash (/), then the base directory is prepended to the path.

paramfile (paramfile)
This option is deprecated. It will be changed to portmap. This will be used to specify what services are on a port.

General Option

The following is a general option for the Network Security Auditor:

database (findingsdb)
Specify a database in which to record security audit findings. The database option applies to hosts which are specified after it. If no database is specified, an in-memory database is used, and a report is automatically generated.

Scan Options

The following are scan options for the Network Security Auditor:

hostfile (hostfile) Specify a host configuration file.

scantype (scantype)
Specify the type of scan to perform. Scan types are defined in the scannerdefs file. The scantype defaults to the name default.

policy (policyname) Specify a site policy configuration file. If the filename is not absolute, then the policy search path is searched. The default is no policy.

t, +t, tcpports, +tcpports (portlist)
Change the TCP ports to be scanned of the current scan type. If the dash (-) form of the option is used, then the port list overrides the current definition. If the plus (+) form of the option is used, then the port list extends the current definition.

u, udpports, +udpports (portlist)
Change the UDP ports to be scanned of the current scan type. If the dash (-) form of the option is used, then the port list the current definition. If the plus (+) form of the option is used, then the port list extends the current definition.

rpcprogs, +rpcprogs (proglist)
Change the Sun RPC programs to be checked of the current scan type. If the dash (-) form of the option is used, then the program list overrides the current definition. If the plus (+) form of the option is used, then the program list extends the current definition.

options, +options (optlist)
Change the options defined within the current scan type. If the dash (-) form of the option is used, then the option overrides the current definition. If the plus (+) form of the option is used, then the options extend the current definition.

xmitrate (number|name)
Specify the maximum number of ports to scan per second. The default is 800. The value can be a name defined in the XMITNames file.

xmithost (number|name)
Specify the maximum number of ports to scan per second for an individual host. The default is 200. The value can be a name defined in the XMITNames file.

maxfd (number) Specify the maximum number of file descriptors to use for performing the port scans. The default is to use all available file descriptors.

tcptimeout (seconds)
Specify the number of seconds to wait when attempting to create a connection to a TCP port. After this amount of time, the operation will be aborted. The default is 10 seconds.

udptimeout (seconds)
Specify the number of seconds to wait for a response from a UDP port. After this amount of time, the port will be probed again, up to the UDP retry count.

udpretry (number) Specify the number of retries when probing a UDP port. The default is 10. Specifying a value of zero means to only send one probe, a value of 1 will cause two probes (the original plus one retry).

tcpsrcport (port) Specify the source port to use when creating TCP connections. In some situations, using a source port of 20 (ftp-data) will provide successful connections that would otherwise be filtered.

concurrent (number)
Specify the number of hosts that can be scanned at the same time. The default is 4.

pingfirst Try sending an ICMP Echo Request to the host and wait for a reply before scanning the host. If no reply is received, the host is not scanned. The request is transmitted ten times before giving up.

random Scan ports in a random order. By default ports are scanned in ascending order. Note that this option will cause more memory to be used.

adaptive The scan rates for hosts will adapt based on retransmissions. This can improve the scanners accuracy. However, some machines are so slow, that it will cause the scanner to take forever. This currently only works if UDP ports are being scanned.

norecord Do not write findings to a database. This option is only valid when used in combination with the hostfile option, without an overriding specification of a database.

noresolve or -n Do not attempt to resolve IP addresses into names. The IP address itself will be used as the hostname.

Report Options

The following are report options for the Network Security Auditor:

format (reporttemplate)
Specify the report template to use for generating the report. The default is standard. Report templates are searched for in the report template search path. There are currently four report templates available:

standard Generates a full report of all findings.

summary Generate report with counters for various categories of findings.

vulnerable Generates only the potential vulnerabilities section.

policy Generates only the policy violations section.

report Force a report to be generated, even when the result of the scan is recorded to a database on disk.

separate Create a separate report for each host that was scanned. This option is only useful with the outfile option.

outfile (filename) Specify the name of file to write the report to. The default is to write to the screen. If the separate option is specified, then the filename is used as extension for the output files.

mailto (mailaddr) If specified, the report is e-mailed to this address. If the report is to be separated, then multiple e-mail messages will be generated. This option is ignored if the outfile option is specified.

html Generate HTML output. The default is to generate a formatted report.

references Add a reference section to the report. The report will contain links to external information, as well as an additional extended information section.

Configuraton Files

The Network Security Auditor uses several configuration files. These are found in the directory specified by the "Directory" option and defaults to `/etc/nsa`. The various files and their contents are described in the following sections.

NSA Configuration File

The following are options for the Network Security Auditor configuration file.

Directory Specify the directory that contains configuration files. The default is `/etc/nsa`.

MessageCatalog Specify the message catalog containing messages used in report generation. The default is messages.cat in the base directory.

ReportCatalog Specify a colon (:) separated list of directories to search for report templates. If elements of the path do not start with a slash (/), then the base directory is prepended to the path.

PolicyCatalog Specify a colon (:) separated list of directories to search for policy definition files. If elements of the path do not start with a slash (/), then the base directory is prepended to the path.

MaxFiles Specify the maximum number of file descriptors to use for performing the port scans. The default is to use all available file descriptors.

RPCNameFile Specify a Sun RPC name/number mapping file. If the filename does not start with a slash (/), then the file is looked for in the base directory. The default is rpc in the base directory.

ScannerDefs Specify file containing definitions of scanner types. If the filename does not start with a slash (/), then the file is looked for in the base directory. The default is scannerdefs in the base directory.

ScannerData Specify data file containing information used by scanners. The default is scannerdata in the base directory.

Concurrent Specify the number of hosts that can be scanned at the same time. The default is 4.

XMITRate Specify the maximum number of ports to scan per second. The default is 800. The value can be a name as defined in the XMITNames file.

XMITHost Specify the maximum number of ports to scan per second for an individual host. The default is 200. The value can be a name as defined in the XMITNames file.

TCPTimeOut Specify the number of seconds to wait when attempting to create a connection to a TCP port. After this amount of time, the operation, will be aborted. The default is 10 seconds.

UDPTimeOut Specify the number of seconds to wait for a response from a UDP port. After this amount of time, the port will be probed again, up to the UDP retry count.

UDPRetry Specify the number of retries when probing a UDP port. The default is 10. Specifying a value of zero means to only send one probe, a value of 1 will cause 2 probes (the original plus one retry).

Format Specify the report template to use for generating the report. The default is standard. Report templates are searched for in the report template search path.

MailTo Specify an address to e-mail reports to. Since the only way to turn this off from the command line is to use the outfile option, use of this in configuration files should be limited.

Policy Specify a site policy configuration file. If the filename is not absolute, then the policy search path is searched. The default is no policy.

PingFirst Try sending an ICMP Echo Request to the host and wait for a reply before scanning the host. If no reply is received, the host is not scanned. The request is transmitted 10 times before giving up.

ScanType Specify the type of scan to perform. Scan types are defined in the scannerdefs file. The scantype defaults to the name default.

HTML Generate HTML output. The default is to generate a formatted report.

References Add a reference section to the report. The report will contain links to external information, as well as an additional extended information section.

Scanner Definitions File

This configuration file can contain lines with the following keywords:

tcpports Specify ranges of TCP ports to scan. Range specifications can be numeric, or use the TCP service names.

udpports Specify ranges of UDP ports to scan. Range specifications can be numeric, or use the UDP service names.

rpcprogs Indicate RPC programs that should be checked. Not implemented yet. All RPC programs for which a scanner exists are checked.

options Indicate optional checks to perform. Currently implemented options are ftp-walk-tree, ip-options, and ip-source-route.

Site Policy Definitions File

Site policy definitions allow you to define what should or should not be found on a network. If a particular type (server, service, etc) has neither allow nor deny clauses, then that type defaults to allow all. If a deny clause exists, but no allow, then anything not in the deny clause is allowed. If an allow clause exists, but no deny, then only things listed in the allow clause are allowed. If both exist, the deny clauses take precedence over the allow clauses. In this case, anything not listed is implicitely allowed.

(allow|deny) server (serverlist)

Indicate servers which are allowed or denied. Servers are any SunRPC server defined in the RPC name file, as well as any standard server name. Examples of standard server names are SMTP, FTP, telnet, HTTP, finger, rshell, and rlogin. See the following example:

```
deny server rlogin,rshell,pcnfs,rex,ypbind,ypserv
```

(allow|deny) service (servicelist)

Specify specific services that are allowed or denied. Currently defined services are anonymous-ftp, smtp-expn-vrfy, and nfs-world-export.

(allow|deny) (tcp|udp) visible (portrange)

Specify what TCP or UDP ports should be visible. A visible port is any that accepts or rejects a connection. A port is considered not visible if it does not respond. This currently does not work well for UDP ports.

(allow|deny) (tcp|udp) active (portrange)

Specify what TCP or UDP ports should be active. An active TCP port is one which accepts a connection request. UDP ports will be considered active if they respond to traffic (does not work correctly yet).

Scanner Control Data File

The scanner control data file contains additional control data for scanning.

Report Templates

Additional report templates may be defined in the report template directory.

Chapter 21. Translating Network Addresses

With the explosive growth of the Internet, IP address depletion becomes a problem. Network address translation (NAT) provides a solution to the IP address depletion problem by allowing addresses inside your secure IP network to be reused by any other IP network.

Network address translation translates secure IP addresses to temporary external registered addresses from the address pool in order to communicate with the outside world. IBM provides network address translation for any TCP or UDP application, which does not require changes in the data transferred. In addition, FTP is supported by providing translation of the IP address in the port command.

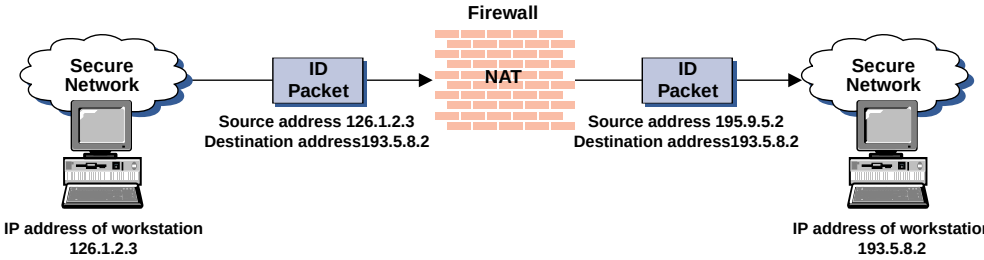

Figure 53. Network Address Translation

NAT-CU-SeeMe does not work with network address translation.

NAT Configuration File

The NAT configuration file `/etc/security/fwnat.cfg` controls the translation of IP addresses in a secure IP address space to IP addresses in an unsecure IP address space. The NAT configuration file `/etc/security/fwnat.cfg` can contain up to 512 of entries listed below. (Reserve, Translate, and Exclude are used to configure dynamic clients. Map is used to define servers).

- Reserve Registered Addresses - A reserve registered address entry defines a set of registered IP addresses that can be used for outbound connections. When a secure host sends a packet to a nonsecure network, a registered IP address is allocated from the reserved registered address pool. This unique registered IP address is used to transport an IP frame between the IBM Firewall and machines outside of the secure network. You can have multiple series of multiple address ranges. The following is an example of a reserve registered address followed by the mask and the timeout value:

  ```
  RESERVE 195.9.5.0 255.255.255.0 30
  ```

- Translate Secured IP Addresses - A translate secured IP address entry defines a set of secure network addresses that require NAT to perform IP address translation. The following is an example of a translate secure IP address:

  ```
  TRANSLATE 126.1.2.0 255.255.255.0
  ```

- Exclude Secured IP Addresses - 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. The following is an example of an exclude secure IP address:

```
EXCLUDE 128.1.2.0 255.255.255.0
```

- Map Secured IP 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 registered IP addresses in the map secure IP address entries can overlap the IP address space specified by the reserve registered IP address entries. The following is an example of a static address translation:

  ```
  MAP 126.1.2.6 195.9.5.6
  ```

Configuring Network Address Translation Using the Configuration Client

1. From the configuration client navigation tree, double-click the Address Translation file folder icon to expand the view. Double-click the NAT file folder icon to expand the view.
2. Select NAT Setup to configure the Network Address Translation module.

   The Network Address Translation List appears, as shown in Figure 54.

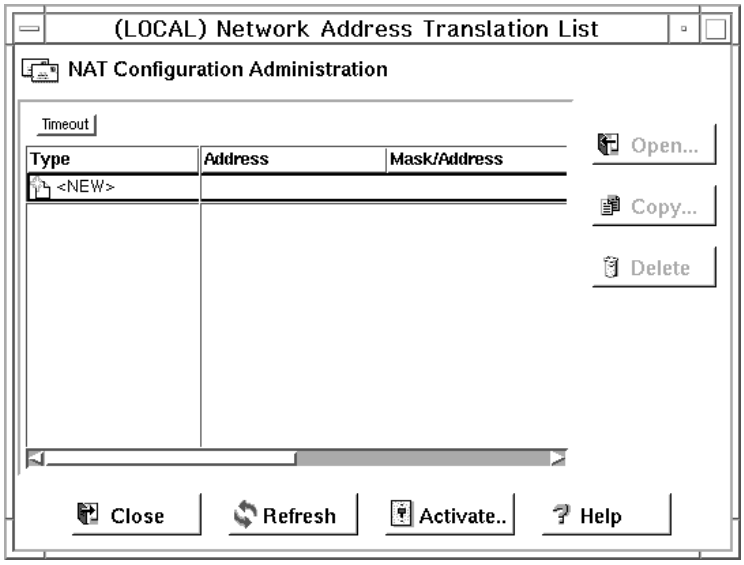

Figure 54. Network Address Translation List

3. Network Address Translation entries contained in the NAT configuration file are displayed on this menu. You can also add, change, or delete NAT entries.

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 dialog screen appears, as shown in Figure 55 on page 151.

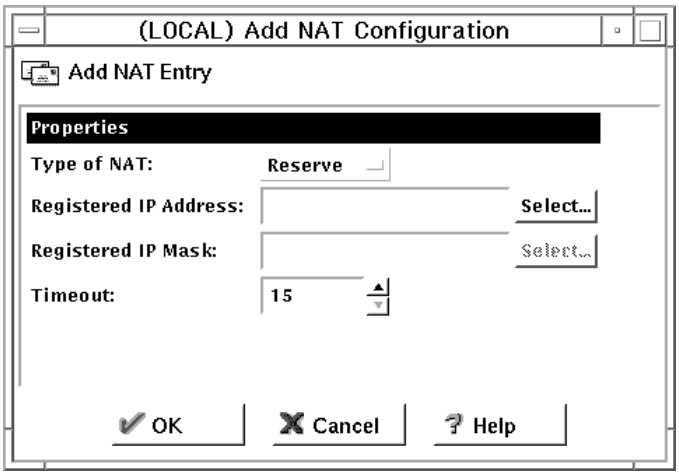

Figure 55. Add NAT Configuration

2. From the Add NAT dialog screen, open the pull-down menu in the Type of NAT field and select from the following:
   - Reserve Registered Network Address: Adds the IP addresses specified to the registered address pool.
   - Translate Secured Network Address: Specifies a range of secure IP addresses that require network address translation.
   - Exclude Secured Network Address: Specifies a range of secure IP addresses that should be excluded from network address translation.
   - Map Secured Network Address: Defines a one-to-one secure-to-registered IP address static translation.

Reserve Registered Network Address

If you selected Reserve from the Add NAT screen, enter the following values:

Registered IP Address
Specify a dotted-decimal IP address that identifies a range of registered IP addresses to be added to the registered address pool.

Choose a network object by clicking Select to get the Select Network Object menu. Select a network object and click OK. The network object is added to the Network Object field on the Add NAT Configuration menu.

Registered IP Address Mask
Specify a mask, like a subnet mask that specifies the bits in the registered IP address used to add a range of IP addresses to the registered address pool. Bits in these masks that are set to 0 indicate that bit positions that have 0 or 1 are reserved registered IP addresses. So, specifying 255.255.255.255 in the mask indicates that only one registered address is added to the registered address pool, whereas a mask of 255.255.0.0 causes class B IP addresses to be added to the registered address pool.

Timeout Value

Enter the number of minutes an address translation can remain idle before NAT can free the registered IP address. This timeout value only applies to the address translation that uses a registered IP address within the range of IP addresses specified by this entry.

The default is 15 minutes and is the minimum value allowed. If you enter a value of less than 15 minutes, the value will be processed as 15 minutes.

Translate Secured Network Address

If you selected Translate from the Add NAT screen, enter the following values:

Secured IP Address

Specify a dotted-decimal IP address that identifies a range of secure IP addresses that require network address translation.

Choose a network object by clicking Select to get the Select Network Object menu. Select a network object and click OK. The network object is added to the Network Object field on the Add NAT Configuration menu.

Secured IP Address Mask

Specify a mask, like 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.

Exclude Secured Network Address

If you selected Exclude from the Add NAT screen, enter the following values:

Secured IP Address

Specify a dotted-decimal IP address that identifies a range of secure IP addresses that should be excluded from network address translation.

Choose a network object by clicking Select to get the Select Network Object menu. Select a network object and click OK. The network object is added to the Network Object field on the Add NAT Configuration menu.

Secured IP Address Mask

Specify 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.

Map Secured Network Address

If you selected Map from the Add NAT Configuration menu, enter the following values:

Value	Description
Secured IP Address	A dotted-decimal IP address that should be translated into a specified registered IP address. You can choose a network object by clicking Select to get the Select Network Object menu. Select a network object and click OK. The network object is added to the Network Object field on the Add NAT Configuration menu.
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 menu. Select a network object and click OK. The network object is added to the Network Object field on the Add NAT Configuration menu.

Change NAT Entry

Select an existing NAT entry from the NAT Configuration menu and click Open to change Network Translation entries in the NAT configuration file.

Delete NAT Entry

1. Select an existing NAT entry from the NAT Configuration menu and click Delete to remove a Network Translation entry from the NAT configuration file.

   A confirmation menu appears.

2. Select Yes or No.

NAT Activation

1. From the configuration client navigation tree, double-click the Address Translation file folder icon to expand the view. Double-click the NAT file folder icon to expand the view.
2. Select NAT Activation and a menu similar to the one shown in Figure 56 on page 154 appears.

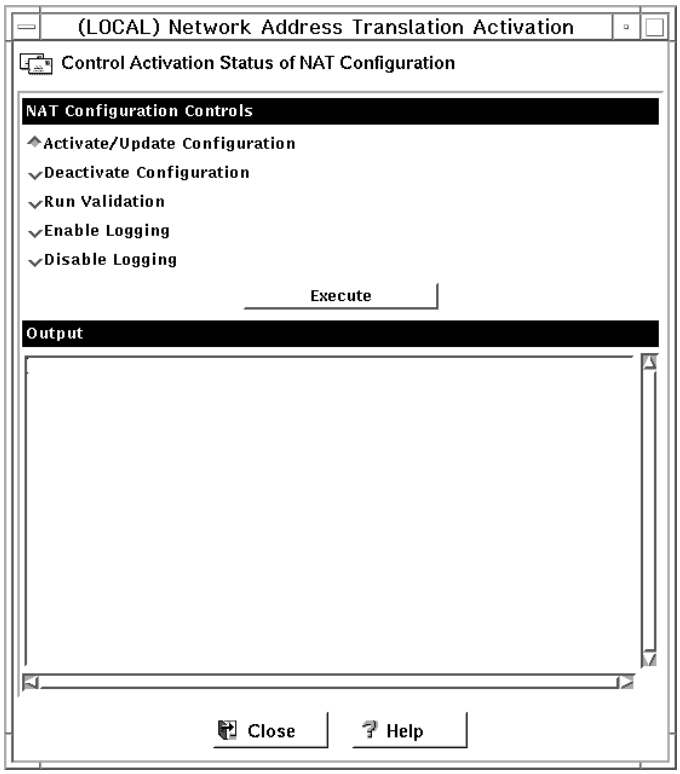

Figure 56. NAT Activation

3. You can select any of the following and then click Execute:
    - Validate network address translation entries contained in a specified NAT configuration file.
    - Activate/Update the configuration to display Network Address Translation entries currently used by the NAT module.
    - Deactivate NAT to disable network address translation.
    - Enable Logging to enable network address translation logging.
    - Disable Logging to disable network address translation logging.

Logging takes place at the initial mapping phase and is not dynamic. It does not show each actual mapping in the kernel upon request. You must enable logging before activation. If not, you could lose log entries. If you refresh the log or erase it, you will loose log entries.

Chapter 22. SNMP

The Simple Network Management Protocol (SNMP) has been widely used in the TCP/IP environment for network management. It can also be used to monitor IBM Firewall server status and generate traps. There are a significant number of SNMP managers existing in customer environments that can be used to monitor the resources and components without introducing the overhead of a management framework and requiring new application programs. Therefore, using SNMP with the IBM Firewall is a natural extension of management of IBM Firewall servers.

SNMP support in the IBM Firewall environment consists of two parts:

1. IBM Firewall Subagent
2. IBM Firewall Management Information Base (MIB)

See the *IBM Firewall Reference* for information on the MIB. The MIB is located in `/etc/fwmib.defs` and must be imported into your network management station. Refer to your network management station documentation for information on how to do this.

To perform SNMP queries from the local firewall, you must have `bos.net.tcp.server` installed for the `snmpinfo` command.

Configuring SNMP Using the Configuration Client

There is a default filter rule upon installation that denies all SNMP traffic. For the IBM Firewall to be managed by an SNMP manager, a predefined filter service to permit a specified SNMP manager IP address can be used.

Upon installation, the SNMP daemon and SNMP Firewall subagent are not started.

Note: It is recommended that you configure DNS before using the SNMP subagent on the firewall.

If you click start on the configuration client to activate the SNMP Firewall subagent, the daemon will automatically be activated. If the SNMP Firewall subagent is active and the machine is brought down, rebooting the machine will start the SNMP Firewall subagent automatically. Issuing a reboot will not start the subagent if the agent was not activated previously. When an SNMP manager is deleted or added to the IBM Firewall, the daemon will be refreshed if it is running.

To configure SNMP, select System Administration from the configuration client navigation tree. Double-click the file folder icon to expand the view. Select SNMP. Double-click the file folder to expand the view. Select Manager. The Add SNMP Manager panel appears. Select Address or Hostname and fill in the remaining fields. Click OK to add an SNMP manager.

From the configuration client navigation tree, select SNMP. Double-click the file folder to expand the view. Select Subagent. The SNMP Sub Agent Configuration panel appears, as shown in Figure 57 on page 156.

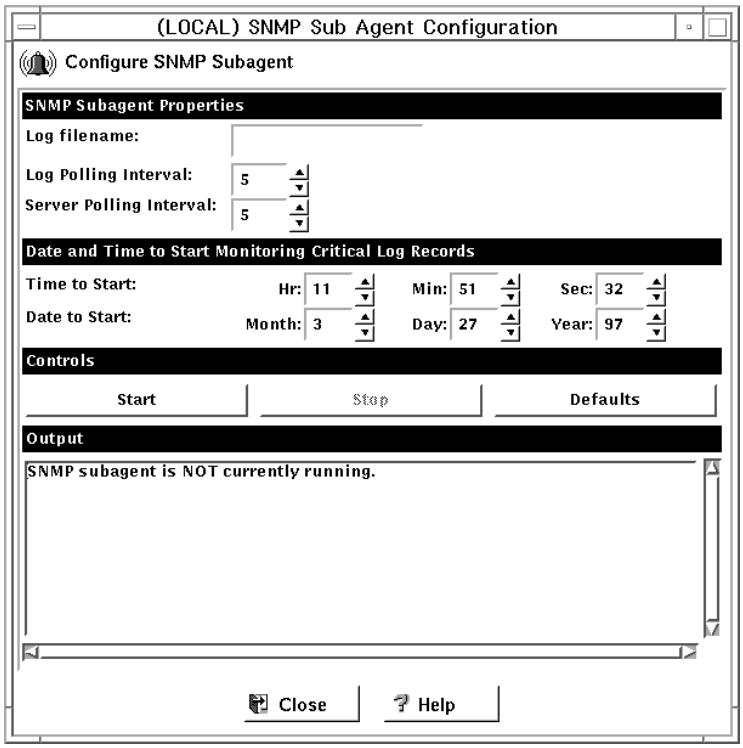

Figure 57. SNMP Sub Agent Configuration

1. Log Filename specifies the name of the critical syslog to be polled by the subagent. This string should be an absolute path to a file. This field defaults to the `local4` log file specified in `/etc/syslog.conf`. When a critical log entry is detected, a trap is sent to the SNMP trap monitor.
2. Log Polling Interval is the frequency with which the critical syslog file is polled for its status. The value of this field must be an integer between 0 and 1440 (24 hours). A value of 0 disables the thread. The default value is 5 minutes. When a critical log entry is detected, a trap is sent to the SNMP trap monitor.
3. Server Polling Interval is the frequency with which the Firewall server daemons are polled for their status. The value of this field must be an integer between 0 and 1440 (24 hours). A value of 0 disables the thread. The default value is 5 minutes.

   Specifically, the following daemons are checked:

   - inetd
   - fwpagerd
   - fwmaild
   - named
   - phttpd
   - sockd

   If the status of any of these daemons has changed from the last poll, a trap is sent to the SNMP trap monitor.
4. Time to Start indicates the time at which to begin monitoring (and trapping) critical log records. The default is the time the subagent is started. Thus, if you would like the monitoring to start at a later time, after you start the subagent, you can customize the time values according to your desired start time.

5. Date to Start indicates the date on which to begin monitoring (and trapping) critical log records. The default is the date the subagent is started. Thus, if you would like the monitoring to start at a later date, after you start the subagent, you can customize the date values according to your desired start date.
6. Click Start to start the subagent with the displayed operational settings.
7. Click Stop to stop the subagent.
8. Click Defaults to return the operational setting values displayed on this screen, to their default values.

For information on trappable events, see the *IBM Firewall Reference.*

Appendix A. Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make them available in all countries in which IBM operates. 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. Subject to IBM's valid intellectual property or other legally protectable rights, any functionally equivalent product, program, or service may be used instead of the IBM product, program, or service. The evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, are the responsibility of the user.

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
USA

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:

Site Counsel
IBM Corporation
P.O. Box 12195
3039 Cornwallis
Research Triangle Park, NC 27709-2195
USA

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 document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement.

This document is not intended for production use and is furnished as is without any warranty of any kind, and all warranties are hereby disclaimed including the warranties of merchantability and fitness for a particular purpose.

IBM is required to include the following statements in order to distribute portions of this document and the software described herein to which contributions have been made by the University of California and NEC Systems Laboratory.

This product includes software developed by the University of California, Berkeley and its contributors.

Portions Copyright © 1993, 1994 by NEC Systems Laboratory.

This product contains code licensed from RSA Data Security Incorporated.

Trademarks

The following terms are trademarks of the IBM corporation in the United States or other countries or both:

AIX	AIXwindows
AIX/6000	Common User Access
DB2	HACMP
IBM	OS/2
RS/6000	RISC System/6000

Microsoft, Windows and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Java and HotJava are trademarks of Sun Microsystems, Inc.

Other company, product, and service names, which may be denoted by a double asterisk (**), may be trademarks or service marks of others.

Bibliography

- Cheswick, William R. and Bellovin, Steven M., *Firewalls and Internet Security*, Addison-Wesley Professional Computing Series, 1994.
- Garfinkle, Simson and Spafford, Gene, *Practical UNIX Security* O'Reilly & Associates, Inc., 1991.
- *AIX/6000 General Concepts and Procedures for IBM RISC System/6000*, GC23-2202.

Glossary

This glossary contains technical terms that are used in the documentation for many of the IBM networking software products. It includes IBM product terminology as well as selected terms and definitions from:

- The *American National Standard Dictionary for Information Systems*, ANSI X3.172-1990, copyright 1990 by the American National Standards Institute (ANSI). Copies may be purchased from the American National Standards Institute, 11 West 42nd Street, New York, New York 10036. Definitions are identified by the symbol (A) after the definition.
- The ANSI/EIA Standard—440-A, *Fiber Optic Terminology*. Copies may be purchased from the Electronic Industries Association, 2001 Pennsylvania Avenue, N.W., Washington, DC 20006. Definitions are identified by the symbol (E) after the definition.
- The *Information Technology Vocabulary* developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC1/SC1). Definitions of published parts of this vocabulary are identified by the symbol (I) after the definition; definitions taken from draft international standards, committee drafts, and working papers being developed by ISO/IEC JTC1/SC1 are identified by the symbol (T) after the definition, indicating that final agreement has not yet been reached among the participating National Bodies of SC1.
- The *IBM Dictionary of Computing*, New York: McGraw-Hill, 1994.
- Internet Request for Comments: 1208, *Glossary of Networking Terms*
- Internet Request for Comments: 1392, *Internet Users' Glossary*
- The *Object-Oriented Interface Design: IBM Common User Access Guidelines*, Carmel, Indiana: Que, 1992.

The following cross-references are used in this glossary:

Contrast with: This refers to a term that has an opposed or substantively different meaning.

Synonym for: This indicates that the term has the same meaning as a preferred term, which is defined in its proper place in the glossary.

Synonymous with: This is a backward reference from a defined term to all other terms that have the same meaning.

See: This refers the reader to an entry that provides more information, to a term that is the expanded version of an abbreviation or acronym, or to a more preferred term.

See also: This refers the reader to terms that have a related, but not synonymous, meaning.

Deprecated term for: This indicates that the term should not be used. It refers to a preferred term, which is defined in its proper place in the glossary.

A

adapter. A part that electrically or physically connects a device to a computer or to another device.

address. In data communication, the unique code assigned to each device, workstation, or user connected to a network.

Administrative Domain. A collection of hosts and routers, and the interconnecting networks, managed by a single administrative authority.

AIX. Advanced Interactive Executive.

AIX operating system. IBM's implementation of the UNIX operating system. The RS/6000 system, among others, runs the AIX operating system.

API. Application programming interface.

application-level gateway. In a firewall, a proxy server that performs a requested service for a client. Contrast with *circuit-level gateway*.

application program interface. See *application programming interface (API)*.

application programming interface (API). The set of programming language constructs or statements that can be coded in an application program to obtain the specific functions and services provided by an underlying operating system or service program.

ASCII (American National Standard Code for Information Interchange). The standard code, using a coded character set consisting of 7-bit coded characters (8 bits including parity check), that is used for information interchange among data processing systems, data communication systems, and associated equipment. The ASCII set consists of control characters and graphic characters. (A)

authentication. (1) In computer security, verification of the identity of a user or the user's eligibility to access an object. (2) In computer security, verification that a message has not been altered or corrupted. (3) In computer security, a process used to verify the user of an information system or protected resources.

B

Berkeley Software Distribution (BSD). Pertaining to any of the series of UNIX specifications or implementations distributed by the University of California at Berkeley. The mnemonic “BSD” is usually followed by a number to specify the particular version of UNIX that was distributed (for example, BSD 4.3). Many vendors use BSD specifications as standards for their UNIX products.

BSD. Berkeley Software Distribution.

button. (1) A mechanism on a pointing device, such as a mouse, used to request or initiate an action or a process. (2) A graphical device that identifies a choice. (3) A graphical mechanism that, when selected, performs a visible action. For example, when a user clicks on a list button, a list of choices appears.

C

circuit-level gateway. In a firewall, a proxy server that redirects a client's request through the firewall to the intended server. Contrast with *application-level gateway*.

click. To press and release a button on a pointing device without moving the pointer off of the object or choice.

client. A user.

command. A request from a terminal for the performance of an operation or the execution of a particular program.

command prompt. A displayed character or string of characters that indicates that a user may enter a command to be processed.

connection. (1) In data communication, an association established between functional units for conveying information. (I) (A) (2) In TCP/IP, the path between two protocol applications that provides reliable data stream delivery service. In the Internet, a connection extends from a TCP application on one system to a TCP application on another system.

D

daemon. A program that runs unattended to perform a standard service. Some daemons are triggered automatically to perform their task; others operate periodically.

DATABASE 2 (DB2). An IBM relational database management system.

Data Encryption Standard (DES). In computer security, the National Institute of Standards and Technology (NIST) Data Encryption Standard, adopted by the U.S. government as Federal Information Processing Standard (FIPS) Publication 46, which allows only hardware implementations of the data encryption algorithm.

datagram. (1) In packet switching, a self-contained packet, independent of other packets, that carries information sufficient for routing from the originating data terminal equipment (DTE) to the destination DTE without relying on earlier exchanges between the DTEs and the network. (I) (2) See *IP datagram*, *packet*, *segment*, and *User Datagram Protocol (UDP)*.

DB2. DATABASE 2.

default. Pertaining to an attribute, condition, value, or option that is assumed when none is explicitly specified. (I)

DES. Data Encryption Standard.

directory. (1) A table of identifiers and references to the corresponding items of data. (I) (A) (2) A type of file containing the names and controlling information for other files or other directories. (3) A listing of the files stored on a disk or diskette.

distinguished name. (1) In systems management, the name of an object formed from the sequence of the relative distinguished names (RDNs) of the object and each of its superior objects. Because each object has exactly one superior object (except the global root, which has none), each object has only one distinguished name. (2) The abstract syntax of a distinguished name or a value of this type of abstract syntax.

DNS. Domain Name System.

domain. See *Administrative Domain* and *domain name*.

domain name. In the Internet suite of protocols, a name of a host system. A domain name consists of a sequence of subnames separated by a delimiter character. For example, if the fully qualified domain name (FQDN) of a host system is `ralvm7.vnet.ibm.com`, each of the following is a domain name:

- ralvm7.vnet.ibm.com
- vnet.ibm.com
- ibm.com

Domain Name System (DNS). In the Internet suite of protocols, the distributed database system used to map domain names to IP addresses.

drive. A peripheral device, especially one that has addressed storage media.

E

EFM. See *Enterprise Firewall Manager*.

electronic mail (e-mail). (1) Correspondence in the form of messages transmitted between user terminals over a computer network. (T) (2) The generation, transmission, and display of correspondence and documents by electronic means. (A)

e-mail. Electronic mail.

encapsulation. In communications, a technique used by layered protocols by which a layer adds control information to the protocol data unit (PDU) from the layer it supports. In this respect, the layer encapsulates the data from the supported layer. In the Internet suite of protocols, for example, a packet would contain control information from the physical layer, followed by control information from the network layer, followed by the application protocol data.

Enterprise Firewall Manager (EFM). A component of the IBM Firewall that allows an organization to manage the configuration of multiple firewalls from a central location. This term may also refer to a machine on which this component is installed or to an IBM Firewall that is configured to be the EFM.

F

File Transfer Protocol (FTP). In the Internet suite of protocols, an application layer protocol that uses TCP and Telnet services to transfer bulk-data files between machines or hosts.

filter. (1) A device or program that separates data, signals, or material in accordance with specified criteria. (A) (2) See also *IP filter*.

finger. In the Internet suite of protocols, a program that displays information about the current users of a local or remote system. The finger usually displays the user's full name, last login time, idle time, terminal line, and terminal location (where applicable).

firewall. In communication, a functional unit that protects and controls the connection of one network to other networks. The firewall (a) prevents unwanted or unauthorized communication traffic from entering the protected network and (b) allows only selected communication traffic to leave the protected network.

FQDN. Fully qualified domain name.

FTP. File Transfer Protocol.

fully qualified domain name (FQDN). In the Internet suite of protocols, the name of a host system that includes all of the subnames of the domain name. An example of a fully qualified domain name is ralvm7.vnet.ibm.com. See also *host name*.

functional unit. An entity of hardware or software, or both, capable of accomplishing a specified purpose. (I) (A)

G

gateway. A functional unit that connects two networks or subnetworks having different characteristics, such as different protocols or different policies concerning security or transmission priority.

Gopher. In the Internet suite of protocols, a distributed information service that makes available hierarchical collections of information. A single Gopher client can access information from any accessible Gopher server. The Gopher client provides the user with a menu-driven interface.

graphical user interface (GUI). A type of computer interface consisting of a visual metaphor of a real-world scene, often of a desktop. Within that scene are icons, representing actual objects, that the user can access and manipulate with a pointing device. Contrast with *command line interface (CLI)*.

GUI. Graphical user interface.

H

hacker. (1) A computer enthusiast who uses his or her knowledge and means to gain unauthorized access to protected resources. (T) (A) (2) A computer enthusiast.

HACMP. See *high-availability cluster multiprocessing*.

handle. (1) In the Advanced DOS and OS/2 operating systems, a binary value created by the system that identifies a drive, directory, and file so that the file can be found and opened. (2) In the AIX operating system, a data structure that is a temporary local identifier for an object. Allocating a handle creates it. Binding a handle makes it identify an object at a specific location.

hardening. The process of disabling nonsecure software on the machine where the IBM Firewall is being installed.

high-availability cluster multiprocessing (HACMP). An application service that enables up to eight RS/6000 servers to access the same data in parallel. This optimizes application execution and scalability and protects against unplanned outages and server downtime.

host. In the Internet suite of protocols, an end system. The end system can be any workstation.

host address. See *IP address*.

host name. In the Internet suite of protocols, the name given to a machine. Sometimes, “host name” is used to mean *fully qualified domain name (FQDN)*; other times, it is used to mean the most specific subname of a fully qualified domain name. For example, if `ralvm7.vnet.ibm.com` is the fully qualified domain name, either of the following may be considered the host name:

- `ralvm7.vnet.ibm.com`
- `ralvm7`

I

ICMP. Internet Control Message Protocol.

internet. A collection of networks interconnected by a set of routers that allow them to function as a single, large network. See also *Internet*.

Internet. The internet administered by the Internet Architecture Board (IAB), consisting of large national backbone networks and many regional and campus networks all over the world. The Internet uses the Internet suite of protocols.

Internet Control Message Protocol (ICMP). The protocol used to handle errors and control messages in the Internet Protocol (IP) layer. Reports of problems and incorrect datagram destinations are returned to the original datagram source. ICMP is part of the Internet Protocol.

Internet Protocol (IP). A connectionless protocol that routes data through a network or interconnected networks. IP acts as an intermediary between the higher protocol layers and the physical network. However, this protocol does not provide error recovery and flow control and does not guarantee the reliability of the physical network.

Internet service provider (ISP). An organization that provides access to the Internet.

interoperability. The capability to communicate, execute programs, or transfer data among various functional units in a way that requires the user to have little or no knowledge of the unique characteristics of those units. (T)

IP. Internet Protocol.

IP address. The 32-bit address defined by the Internet Protocol, standard 5, Request for Comments (RFC) 791. It is usually represented in dotted decimal notation.

IP datagram. In the Internet suite of protocols, the fundamental unit of information transmitted through an internet. It contains source and destination addresses, user data, and control information such as the length of the datagram, the header checksum, and flags indicating whether the datagram can be or has been fragmented.

IP filter. In the Internet suite of protocols, a set of rules based on IP addressing that control whether one host can access another host through a firewall.

IP tunnel. A mechanism for data encapsulation across an IP network.

ISP. See *Internet Service Provider*.

J

Java. An object-oriented programming language for portable interpretive code that supports interaction among remote objects. Java was developed and specified by Sun Microsystems, Incorporated.

L

LAN. Local area network.

local area network (LAN). (1) A computer network located on a user's premises within a limited geographical area. Communication within a local area network is not subject to external regulations; however, communication across the LAN boundary may be subject to some form of regulation. (T) (2) A network in which a set of devices are connected to one another for communication and that can be connected to a larger network.

login. The procedure by which a user begins a terminal or communication session.

M

Management Information Base (MIB). (1) A collection of objects that can be accessed by means of a network management protocol. (2) A definition for management information that specifies the information available from a host or gateway and the operations allowed. (3) In OSI, the conceptual repository of management information within an open system.

menu. (1) A list of options displayed to the user by a data processing system, from which the user can select an action to be initiated. (T) (2) In text processing, a list of choices displayed to the user by a text processor from which the user can select an action to be initiated. (T) (3) A list of choices that can be applied to an object. A menu can contain choices that are not available for selection in certain contexts. Those choices are indicated by reduced contrast.

message. An assembly of characters and sometimes control codes that is transferred as an entity from an originator to one or more recipients. A message consists of two parts: envelope and content. (T)

MIB. Management Information Base.

modem (modulator/demodulator). (1) A functional unit that modulates and demodulates signals. One of the functions of a modem is to enable digital data to be transmitted over analog transmission facilities. (T) (A) (2) A device that converts digital data from a computer to an analog signal that can be transmitted on a telecommunication line, and converts the analog signal received to data for the computer.

multihomed host. In the Internet Protocol (IP), a host that is connected to more than one network.

N

name resolution. In Internet communications, the process of mapping a machine name to the corresponding Internet Protocol (IP) address. See also *Domain Name System (DNS)*.

NAT. See *network address translation*.

National Computer Security Association (NCSA). An independent organization that strives to improve computer security by working with and fostering interaction among its members and constituents, which include computer users; product developers and vendors in the computer and communication industry; and computer and information security experts.

NCSA. See *National Computer Security Association*.

network. (1) An arrangement of nodes and connecting branches. (T) (2) A configuration of data processing devices and software connected for information interchange. (3) A group of nodes and the links interconnecting them.

network address translation (NAT). In a firewall, the conversion of secure IP addresses to external registered addresses. This enables communication with external networks but masks the IP addresses that are used inside the firewall.

Network Security Auditor. In an IBM Firewall, a program that scans a list of hosts and reports weak spots and potential security exposures for each system.

nonsecure interface. For security gateways, the physical layer connection between the gateway and a nonsecure network. Contrast with *secure interface*.

nonsecure network. A set of nodes that are not controlled by a single administrative party. Contrast with *secure network*.

O

octal. (1) Pertaining to a selection, choice, or condition that has eight possible different values or states. (I) (A) (2) Pertaining to a fixed-radix numeration having a radix of eight. (I) (A)

octet. A byte that consists of 8 bits. (T)

P

packet. In data communication, a sequence of binary digits, including data and control signals, that is transmitted and switched as a composite whole. The data, control signals, and, possibly, error control information are arranged in a specific format. (I)

parameter. A variable that is given a constant value for a specified application and that may denote the application. (I) (A)

partitioned data set (PDS). A data set in direct access storage that is divided into partitions, called members, each of which can contain a program, part of a program, or data.

path. The route used to locate files; the storage location of a file. A fully qualified path lists the drive identifier, directory name, subdirectory name (if any), and file name with the associated extension.

PDS. Partitioned data set.

PDU. Protocol data unit.

port. (1) An access point for data entry or exit. (2) A connector on a device to which cables for other devices such as display stations and printers are attached. (3) In the Internet suite of protocols, a 16-bit number used to communicate between TCP or the User Datagram Protocol (UDP) and a higher-level protocol or application. Some protocols, such as File Transfer Protocol (FTP) and Simple Mail Transfer Protocol (SMTP), use the same well-known port number in all TCP/IP implementations. (4) An abstraction used by transport protocols to distinguish among multiple destinations within a host machine. (5) Synonymous with *socket*.

port number. In Internet communications, the identification of an application entity to the transport service.

PostScript. A standard specified by Adobe Systems, Incorporated, that defines how text and graphics are presented on printers and display devices.

protocol. A set of semantic and syntactic rules that determine the behavior of functional units in achieving communication. (I)

protocol data unit (PDU). A unit of data specified in a protocol of a given layer and consisting of protocol control information of this layer, and possibly user data of this layer. (T)

protocol suite. A set of protocols that cooperate to handle the transmission tasks for a communication system.

proxy server. A server that receives requests intended for another server and that acts on the client's behalf (as the client's proxy) to obtain the requested service. A proxy server is often used when the client and the server are incompatible for direct connection (for example, when the client is unable to meet the security authentication requirements of the server but should be permitted some services).

R

RealAudio system. A client/server-based media delivery system developed by Progressive Networks. The RealAudio system supports live and on-demand audio over the Internet and can be used by news, entertainment, sports, and business organizations to create and deliver multimedia over the Internet.

reduced instruction-set computer (RISC). A computer that uses a small, simplified set of frequently used instructions for rapid execution.

Request for Comments (RFC). In Internet communications, the document series that describes a part of the Internet suite of protocols and related experiments. All Internet standards are documented as RFCs.

RFC. Request for Comments.

RISC. Reduced instruction-set computer.

S

SafeMail. An IBM proprietary mail gateway.

secure interface. For security gateways, the physical layer connection between the gateway and a secure network. Contrast with *nonsecure interface*.

secure network. A set of nodes that are controlled by a single administrative party. Contrast with *nonsecure network*.

Sendmail. In the UNIX operating system, the mail server that uses the Simple Mail Transfer Protocol (SMTP) to route mail from one host to another on the network.

server. (1) A functional unit that provides shared services to workstations over a network; for example, a file server, a print server, a mail server. (T) (2) In a network, a data station that provides facilities to other stations; for example, a file server, a print server, a mail server. (A)

session. In network architecture, for the purpose of data communication between functional units, all the activities which take place during the establishment, maintenance, and release of the connection. (T)

Simple Mail Transfer Protocol (SMTP). In the Internet suite of protocols, an application protocol for transferring mail among users in the Internet environment. SMTP specifies the mail exchange sequences and message format. It assumes that the Transmission Control Protocol (TCP) is the underlying protocol.

Simple Network Management Protocol (SNMP). In the Internet suite of protocols, a network management protocol that is used to monitor routers and attached networks. SNMP is an application layer protocol. Information on devices managed is defined and stored in the application's Management Information Base (MIB).

SMIT. System Management Interface Tool.

SMTP. Simple Mail Transfer Protocol.

SNMP. Simple Network Management Protocol.

socket. (1) An endpoint for communication between processes or application programs. (2) Synonym for *port*.

socket interface. A Berkeley Software Distribution (BSD) application programming interface (API) that

allows users to easily write their own communication application programs.

socks server. A circuit-level gateway that provides a secure one-way connection through a firewall to server applications in a nonsecure network.

spoofing. A hacker's technique of using someone else's IP address to gain access to a network.

SQL. Structured Query Language.

Structured Query Language/Data System (SQL/DS). An IBM relational database management system.

subdirectory. A directory contained within another directory in a file system hierarchy.

subnet. (1) In TCP/IP, a part of a network that is identified by a portion of the IP address. (2) Synonym for *subnetwork*.

subnet address. In Internet communications, an extension to the basic IP addressing scheme where a portion of the host address is interpreted as the local network address.

subnetwork. (1) Any group of nodes that have a set of common characteristics, such as the same network ID. (2) Synonymous with *subnet*.

System Management Interface Tool (SMIT). An interface tool of the AIX operating system for installing, maintaining, configuring, and diagnosing tasks.

T

TCP. Transmission Control Protocol.

TCP/IP. Transmission Control Protocol/Internet Protocol.

Telnet. In the Internet suite of protocols, a protocol that provides remote terminal connection service. It allows users of one host to log on to a remote host and interact as directly attached terminal users of that host.

Time Sharing Option (TSO). An option of the MVS operating system that provides interactive time sharing from remote terminals.

Transmission Control Protocol (TCP). A communications protocol used in the Internet and in any network that follows the U.S. Department of Defense standards for internetwork protocol. TCP provides a reliable host-to-host protocol between hosts in packet-switched communications networks and in interconnected systems of such networks. It uses the Internet Protocol (IP) as the underlying protocol.

Transmission Control Protocol/Internet Protocol (TCP/IP). A set of communications protocols that support peer-to-peer connectivity functions for both local and wide area networks.

TSO. Time Sharing Option.

tunnel. See *IP tunnel*.

U

UDP. User Datagram Protocol.

UNIX operating system. An operating system developed by Bell Laboratories that features multiprogramming in a multiuser environment. The UNIX operating system was originally developed for use on minicomputers but has been adapted for mainframes and microcomputers. The AIX operating system is IBM's implementation of the UNIX operating system.

user. (1) Any person or any thing that may issue or receive commands and messages to or from the information processing system. (T) (2) Anyone who requires the services of a computing system.

User Datagram Protocol (UDP). In the Internet suite of protocols, a protocol that provides unreliable, connectionless datagram service. It enables an application program on one machine or process to send a datagram to an application program on another machine or process. UDP uses the Internet Protocol (IP) to deliver datagrams.

V

virtual private network (VPN). A network comprised of one or more secure IP tunnels connecting two or more networks.

VPN. See *virtual private network*.

W

World Wide Web (WWW). A network of servers that contain programs and files. Many of the files contain hypertext links to other documents available through the network.

WWW. World Wide Web.

Index

Special Characters

A

B

C

D

E

F

G

H

I

K

L

M

N

O

P

R

S

T

U

V

W

Communicating Your Comments to IBM

IBM Firewall for AIX
User's Guide
Version 3.1.1

Publication No. GC31-8419-00

If you especially like or dislike anything about this book, please use one of the methods listed below to send your comments to IBM. Whichever method you choose, make sure you send your name, address, and telephone number if you would like a reply.

Feel free to comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this book. However, the comments you send should pertain to only the information in this manual and the way in which the information is presented. To request additional publications, or to ask questions or make comments about the functions of IBM products or systems, you should talk to your IBM representative or to your IBM authorized remarketer.

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.

If you are mailing a readers' comment form (RCF) from a country other than the United States, you can give the RCF to the local IBM branch office or IBM representative for postage-paid mailing.

- If you prefer to send comments by mail, use the RCF at the back of this book.
- If you prefer to send comments by FAX, use this number:

  1-800-227-5088(US and Canada)

- If you prefer to send comments electronically, use this network ID:
  - USIB2HPD@VNET.IBM.COM
  - USIB2HPD at IBMMAIL

Make sure to include the following in your note:

- Title and publication number of this book
- Page number or topic to which your comment applies.

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

IBM Firewall for AIX
User's Guide
Version 3.1.1

Publication No. GC31-8419-00

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	□	□	□	□	□

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	□	□	□	□	□
Complete	□	□	□	□	□
Easy to find	□	□	□	□	□
Easy to understand	□	□	□	□	□
Well organized	□	□	□	□	□
Applicable to your tasks	□	□	□	□	□

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. _____

Readers' Comments — We'd Like to Hear from You
GC31-8419-00

Cut or Fold
Along Line

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

GC31-8419-00

Cut or Fold
Along Line

Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

GC31-8419-00

Spine information:

IBM Firewall for AIX User's Guide *Version 3.1.1*