IBM eNetwork Firewall for Windows NT

User's Guide

Version 3 Release 2.1.1

GC31-8658-01

IBM eNetwork Firewall for Windows NT

User's Guide

Version 3 Release 2.1.1

GC31-8658-01

Note

Before using this information and the product it supports, be sure to read the general information under “Appendix. Notices” on page 117.

Second Edition (June 1998)

This edition applies to Version 3 Release 2.1.1 of the IBM eNetwork Firewall for Windows NT (product number 5765-C16). This edition replaces GC31-8658-00.

Contents

About This Book

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

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

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 dialog box to dialog box.

Prerequisite Knowledge

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

An excellent book on TCP/IP that covers netstat, arp, ifconfig, ping, nslookup, DNS, sendmail, routing, and much more is *TCP/IP Network Administration*. See the *Bibliography* for more details.

An excellent book for those performing UNIX administration, that also gives and excellent overview of TCP/IP and routing, network hardware, DNS, and sendmail is the *UNIX System Administration Handbook*. See the Bibliography for more details.

Features in This Release

The IBM eNetwork Firewall for Windows NT offers a rich variety of features and includes all three firewall architectures:

1. Application proxies
   - FTP
   - HTTP, including Gopher and WAIS
   - Telnet
   - SafeMail

   HTTP, Telnet, and FTP have authentication capability.
2. Circuit-level gateway through Socks Protocol Version 5, an Internet standard
3. Filtering—an extensive and robust set of criteria on which traffic can be permitted or denied. Criteria include TCP/IP address, port, protocol, direction, adapter (secure/nonsecure), and more.

   Many predefined services make setup fast.

Socks Protocol Version 5

In addition to its simplicity and flexibility, Socks Protocol Version 5 offers these advantages:

- Easy deployment of authentication and encryption methods
- UDP association, which creates a virtual proxy circuit for traversing UDP-based proxy circuits.
- Socks V5 Watcher, which displays real-time socks performance information

Network Address Translation

With the explosive growth of the Internet, IP address depletion problem has become significant. Network address translation (NAT) provides a solution to the IP address depletion problem based upon address reuse.

The advantage of NAT is that it transparently allows a network that uses private or illegal addresses to communicate with hosts on the Internet, effectively allowing the private network to have a large address space. Furthermore, by using NAT, addresses in the private network are hidden from the external world providing an additional level of security.

Simple Administration

Through use of a Java** application, which you can administer from a remote machine, you can easily make updates to the firewall configuration. And, different administrators can be assigned different levels of authority to further control access to the firewall. This single, easy-to-understand graphical user interface (GUI) can be used to administer both the Windows NT Firewall and AIX Firewall.

Hardening of NT

When the firewall is installed, non TCP/IP protocols are disabled, unneeded system services are disabled, and local logins from nonadministrator accounts are disabled.

Strong Authentication

Support for all the popular token-based authentication mechanisms, such as SecurID, SecureNet Key, and others, is offered.

Report Utilities

Report utilities allows you to run an SQL query against the system log once it is exported to a database engine.

Alerting, Monitoring, and Logging

Extensive and detailed logging includes all firewall activity along with TCP/IP address, userids, TOD, filenames, port numbers, and so forth. A Log Monitor is included to watch for suspicious activity and alert you when thresholds are exceeded.

Isolate Multiple Networks

By using multiple Network Interface Cards (NICs) in your firewall, you can isolate multiple subnetworks.

National Language Support

National language support is offered for English, Japanese, Korean, French, simplified Chinese, traditional Chinese, Italian, Spanish, and Brazilian Portuguese.

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

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

Firewall Concepts

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

IBM Firewall Tools

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

Because the IBM Firewall is, at heart, an IP gateway, it divides the world into two or more networks: one or more nonsecure networks and one or more secure

networks. The nonsecure network is, for instance, the Internet. The secure networks are usually your corporate IP networks. Some of the tools that the IBM Firewall offers are:

- Expert filters
- Proxy servers
- Socks servers
- Specific services such as domain name service (DNS) and SafeMail

Expert Filters

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

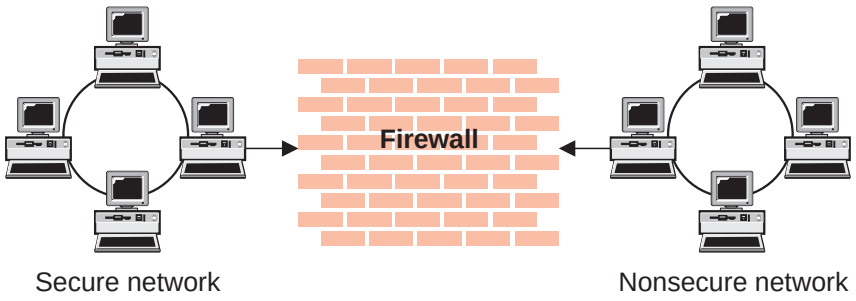

Figure 1. Firewall with Expert Filtering

Objectives of Expert Filters

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

Proxy Servers

Unlike filtering, which merely inspects packets passing through, proxy servers are applications that are part of the firewall and perform specific TCP/IP functions on behalf of a network user. The user contacts the proxy server using one of the TCP/IP applications (Telnet or FTP). The proxy server makes contact with the remote host on behalf of the user, thus controlling access while hiding your network structure from external users. Figure 2 on page 3 illustrates a proxy Telnet server intercepting a request from an external user.

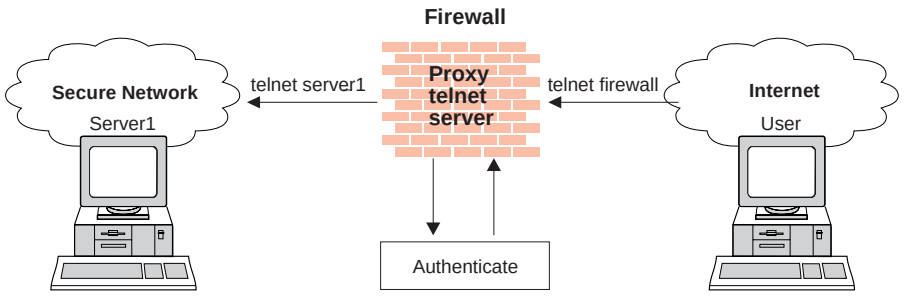

Figure 2. Firewall with a Proxy Server

The proxy services available are telnet, FTP, HTTP, WAIS, GOPHER, and HTTPS, and SafeMail.

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

Objectives of Proxy Servers

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

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

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

Socks Server

Socks is a standard for circuit-level gateways that provides address hiding but does not require the overhead of a more conventional proxy server.

The Socks server is similar to a proxy server in that the session is broken at the firewall. The difference is that socks can support all applications instead of requiring a unique proxy for each application. Transparently, the socks client starts a session with the Windows NT socks service on the IBM Firewall host then validates that the source address and user ID are permitted to establish onward connection into the nonsecure network and then creates the second session. Figure 3 on page 4 illustrates a firewall with a socks server.

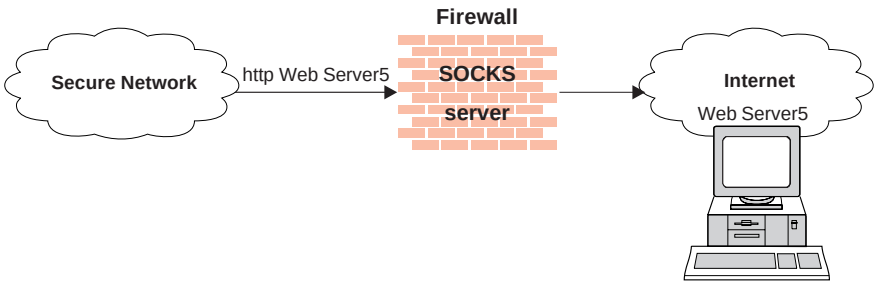

Figure 3. Firewall with a Socks Server

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

Objectives of the Socks Server

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

Domain Name Service

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

Objectives of the DNS Server

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

SafeMail

Mail is one of the primary reasons why an organization would want to access the Internet. SafeMail is an IBM mail gateway designed to hide the domain names of your internal network. The SafeMail function does not store mail on the gateway or run under the root user ID. The firewall gateway public domain name is substituted in place of the private domain names on outgoing mail so that mail appears to be coming from the firewall's address instead of the user's address. SafeMail supports Simple Mail Transfer Protocol (SMTP) and Multipurpose Internet Mail Extensions (MIME).

Using the Network Security Auditor

The Network Security Auditor scans 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 make corrections. The Network Security Auditor can be used as a periodic scanner of critical hosts or as a one-time information gathering tool. Administration of the Network Security Auditor is done through an easy-to-use command line interface. With the Network Security Auditor, you maintain vigilance over your firewall.

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

Chapter 2. Planning

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

Planning Checklist

1. Define your objective. Do you want to:
   - Access the Internet (telnet, anonymous FTP, etc.)?
   - Partition parts of your internal network?
   - Provide *external* access to your network?
2. 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?
3. Decide how you will use DNS. Refer to “Chapter 6. Handling Domain Name Service” on page 31.
4. Decide how you will use safemail. Refer to “Chapter 7. SafeMail” on page 39.
5. If you want to use socks, ensure socksified clients, such as the Netscape Navigator or the Microsoft browser are installed. For information on using socks, see “Chapter 11. Configuring the Socks Server” on page 67.
6. 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** SecureNetKey** card is to be used, purchase cards independently of the IBM Firewall.
   - If you use your own authentication method, see the chapter on Providing Your Own Authentication Methods in the *IBM eNetwork Firewall Reference*.
   - You must configure the Windows client code that implements the ability to search trusted Windows NT domains for authentication purposes, to use TCP instead of NETBIOS. NETBIOS will be disabled. The trusted Windows NT servers must have TCP/IP host names and addresses and have TCP/IP connectivity between them and the firewall. The firewall administrator needs to create connections between the firewall and the trusted NT servers in order to permit traffic to flow between the two.

     Set up this connection using the following predefined services:

     a. Domain Controller Authentication - which allows the use of the Domain Controller for user authentication

     b. NetBT Name Services broadcasts - which allows NetBIOS over TCP/IP Name Services broadcasts

     And use the NT configuration utilities to define the trust relationships.
7. 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.

8. Decide on a method for archiving log files. Archiving is an ideal candidate for a scheduled job in the Windows NT Scheduler service. See “Chapter 15. Managing Log and Archive Files” on page 103.

Network Configuration Planning Worksheet

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

Host name of firewall _____

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

IP address _____ Subnet Mask _____

IP address _____ Subnet Mask _____

IP address _____ Subnet Mask _____

IP address _____ Subnet Mask _____

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

IP address _____ Subnet Mask _____

IP address _____ Subnet Mask _____

IP address _____ Subnet Mask _____

IP address _____ Subnet Mask _____

Name of router _____

Address of router _____

Secure domain name _____

IP address of secure domain name server (DNS) _____

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

Secure Mail Server _____

Public Domain Name _____

IP address of the configuration client _____

IP address of the remote client(s) _____

Root directory of your Windows NT Firewall_____

(We refer to this as ROOTDIR throughout the documentation)

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

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

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

Setting Up the Configuration Server

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

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

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

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

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

localonly=no
The configuration can occur from any machine.

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

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

encryption=none
No encryption will occur; this is the default.

encryption=ssl
SSL encryption will occur.

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

If a configuration client cannot connect to the Firewall machine, and is on a different machine, use `fwcfgsrv cmd=list` to check that `localonly=no` is set. Also, the language used by the client and the server must match. Finally, ensure that the configuration server service is running by bringing up the services panel and

checking its status. To do this, go to the control panel, double-click the Services icon to check the status of each service. If it is not running, the service should be restarted.

Setting Up the Configuration Client (GUI)

When you install the IBM Firewall, the configuration client is automatically installed. The configuration client can also be separately installed on any Windows NT machine without the Firewall, which enables you to perform remote administration. To start the configuration client, double-click the configuration client icon in the IBM Firewall program group. When the configuration client is started, you must first log on to the Firewall using an Windows NT administrator account.

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

After the Firewall is installed, all Windows NT administrators are designated as primary firewall administrators. Use the configuration client to log on to the configuration server using a primary firewall administrator and define the additional firewall administrator usernames, if necessary. See “Chapter 12. Administering Users at the Firewall” on page 73 for information on how to define firewall administrators using the configuration client.

To set the logon timeout value for faster or slower machines, make the following change by clicking the IBM Firewall Configuration Client icon, then click **Properties**. Modify Properties by using the **Shortcut** tab. Change the parameter `timeout` to *20*, where 20 equals the number of seconds to wait for a connection to occur. 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, run `ibmfw.bat` in `ROOTDIR\cfgcli\gui` instead of using the configuration client icon. Note however, that enabling console logging can degrade performance.

Log On to the Configuration Client

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

- The user must be a firewall administrator
- The firewall administrator must have an authentication scheme defined. See “User Authentication Methods” on page 78.
- 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:

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

- Similarly, if the administrator is on the nonsecure side and using a nonsecure interface on the Firewall machine, then he or she must be defined with the appropriate authentication method for nonsecure administration. (It cannot be set to deny all).

All of the user attributes can be set through the Modify User dialog box in the configuration client or by using the command `fwuser`. All firewall administrators will have all of the above fields set appropriately after installation of the Firewall. Refer to “Chapter 12. Administering Users at the Firewall” on page 73 for more information.

Sample Logging Output for the Remote Configuration Server

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

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

Feb 03 13:52:21 1998 mr16n18: ICA2024i: User administrator successfully
authenticated using NT authentication from secure network:127.0.0.
Feb 03 13:52:21 1998 mr16n18: ICA2169i: User administrator successfully
authenticated for Remote Administration Server using NT from secure network:127.0.0.1.
```

Chapter 4. Using the Configuration Client

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

When you first install the IBM Firewall, it is initially set up to only accept requests from the configuration client on the local machine. However, you can install the configuration client on another machine and administer the Firewall remotely. See “Setting Up the Configuration Server” on page 11 for information on how to do this.

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

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

A mouse is required to use the configuration client.

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

How to Log On to the Configuration Client

1. For Logon Type, select Local if you are on the same machine as the firewall. Local is the default. Select Remote if you want to remotely access another Firewall. Remote requires that you enter a host name.
2. If you selected Remote logon, you need to enter the host name or the IP address of the firewall machine you want to log on to.
3. Select either SSL or none depending upon which encryption is used for the Firewall. For the Client, the default for Local is None and the default for Remote is SSL.
4. Enter a user name of a firewall administrator or a Windows NT administrator.
5. Enter the port number on which the server is listening. The default is 1014.
6. For Mode, select Host if you want to configure an Windows NT firewall machine that you are logging on to. With host administration, the administrator can locally or remotely update one Firewall at a time. Select Enterprise for the Enterprise Firewall Management (EFM) administration of AIX firewalls.
7. After you log on, you will see authentication messages and you might be prompted to enter a password if that is the authentication method setup for your user name. If you are prompted for a password, enter your password in the User Response field and either press Enter or click Submit. If you enter an incorrect password, you get a message. Click Close and restart the logon

process. If you are not prompted for a password, your user authentication method might be permit all. In this case you will immediately get the IBM Firewall configuration client panel.

8. After you have successfully been authenticated, you will see the main configuration panel.

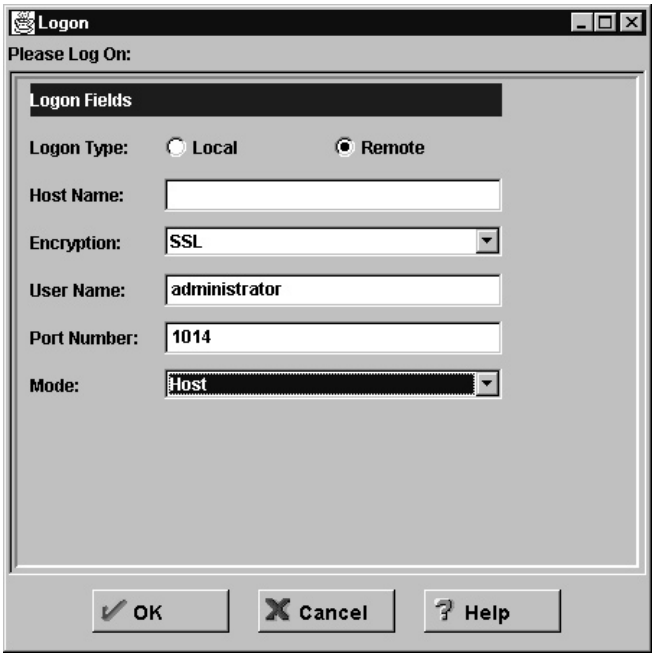

Figure 4. Configuration Client Logon Panel

The Navigation Tree

The configuration client has a collapsible tree-style navigation aid along the left side, as shown in Figure 5 on page 17.

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 5. 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 5.

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.

Other 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 the application.

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.

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 6 on page 19.

The alert records displayed are obtained from the file identified by the first `alert log` facility defined in `ROOTDIR\config\syslog.conf`. If no `alert log` facility is defined, you will see a blank display. See “Add Log Facilities” on page 103 for help in defining an `alert log` facility.

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

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

Figure 6. The Alerts Display

The Log Viewer

Clicking **Log Viewer** brings up a log viewer window, as shown in Figure 7 on page 20 . The log viewer allows you to view firewall log records. You can specify a log file and a record count (default is 25).

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

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

See "Log File Creation and Archiving Using the Configuration Client" on page 103 and "Chapter 14. Monitoring the Firewall Logging" on page 93 for more information about log files, facilities, monitoring and alerts.

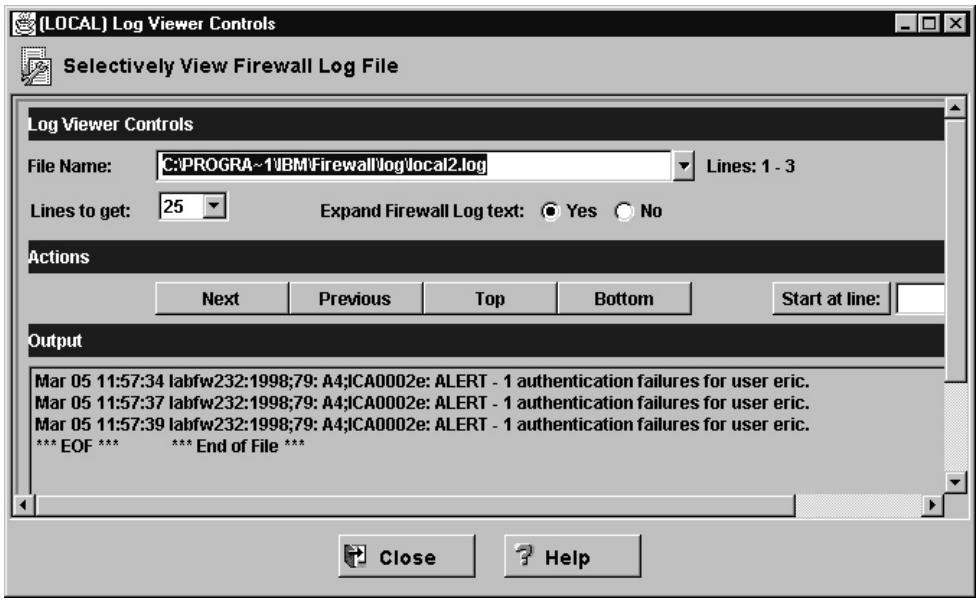

Figure 7. Log Viewer

Other Features

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

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

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

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

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

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

Copy The **Copy** button saves time when adding new items to the list. After selecting an item on the list, click **Copy** to create an item that is similar to the selected item. 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

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

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

Name Provide a name for this item. This item name must be unique for this particular function in the firewall. The name should NOT contain a pipe symbol(|), a single quote (or apostrophe) character('), or a double quote(″) character 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.

Unique Features

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

For a Windows 95 or Windows NT configuration client, the configuration client looks best with a minimum resolution of 1024 pixels x 768 pixels.

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 log on to the Firewall two or more times in quick succession using SSL, the connection will be refused. Exit and restart the configuration client.

Chapter 5. Getting Started on the IBM Firewall

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

Basic Configuration Steps

For a basic IBM Firewall setup:

1. Plan for your IBM Firewall setup. Decide in advance which functions of the firewall you want to use and how you want to use them. These sections are helpful:
   - “Chapter 1. Introducing the IBM Firewall” on page 1
   - “Chapter 2. Planning” on page 7
   - “Planning Considerations” on page 51
2. Tell the Firewall which of its interfaces are connected to secure networks. You must have a secure interface and a nonsecure interface for your firewall to work properly. From the configuration client navigation tree, open the System Administration folder and click **Interfaces** to see a list of the network interfaces on your firewall. To change the security status of an interface, select an interface and click **Change**. See “Designating Your Network Interface” on page 24 for more information.

   If you are going to connect to the Internet, contact your ISP to obtain a registered IP address for the Firewall nonsecure interface.
3. Set up your general security policy by accessing the **Security Policy** dialog in the System Administration folder. For typical Firewall configurations:
   - 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 25 for more information.
4. Set up your domain name service and mail service. Access these functions from the System Administration folder on the configuration client navigation tree. First read “Chapter 6. Handling Domain Name Service” on page 31.
5. Define key elements of your network(s) to the firewall using the **Network Objects** function in the configuration client navigation tree. Network Objects control traffic through the Firewall. Define the following key elements as network objects:
   - Secure Interface of the Firewall
   - Nonsecure Interface of the Firewall
   - Secure Network
   - Each subnet on your secure network
   - A host network object for your SDI servers and your NT domain servers, if appropriate

   See “Network Objects” on page 26 for more information.

6. Enable services on the Firewall. These are the methods by which users in the secure network can access the nonsecure network (such as socks or proxy). Which services get implemented depend on decisions you made at the planning stage. Implementing a service often requires setting up some connection configurations to allow certain types of traffic. For example, if you want to allow your secure users to surf the web on the Internet by using HTTP Proxy, not only do you need to configure the HTTP Proxy daemon on the Firewall, but you also need to set up connections to allow HTTP traffic. See “Chapter 9. Examples of Services” on page 51 for information on how to set up connections that support certain services.
7. Set up firewall users. If you are going to require authentication for functions like outbound Web access or for firewall administrators, you need to define these users to the Firewall. See “Chapter 12. Administering Users at the Firewall” on page 73 for more information.
8. If you want to use Windows NT domain passwords for authentication, you must configure the Windows client code that implements the ability to search trusted Windows NT domains for authentication purposes, to use TCP instead of NETBIOS. NETBIOS will be disabled. The trusted Windows NT servers must have TCP/IP host names and addresses and have TCP/IP connectivity between them and the Firewall. The firewall administrator needs to create connections between the Firewall and the trusted Windows NT servers in order to permit traffic to flow between the two.
9. If you will be using network address translation, first contact your ISP to obtain a registered Internet address for use with many-to-one address translation. Then, go to the *Add NAT Configuration* panel to add the registered Internet address in the *Many-to-One IP Address* field. For more information, see “Chapter 16. Translating Network Addresses” on page 109.

Following these steps should help you to get a basic firewall configuration up and running. The IBM Firewall provides other functions, such as system logs to help you ensure the security of your network. See “Chapter 15. Managing Log and Archive Files” on page 103 for more information.

Designating Your Network Interface

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

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

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

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

To identify a network interface as either secure or nonsecure:

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

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

Using the Configuration Client to Define a Security Policy

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

The IBM Firewall provides a dialog box to assist you in setting up your security policy, as shown in Figure 8.

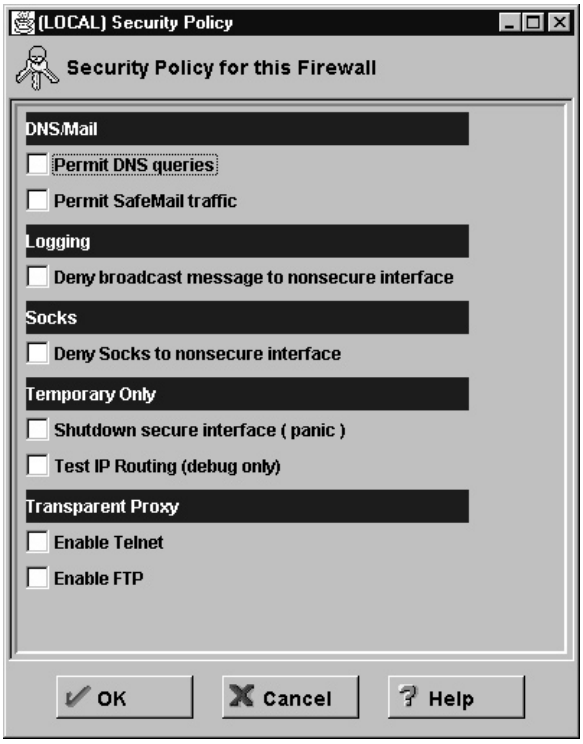

Figure 8. Security Policy

Click Help to learn more about the security policy panel.

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

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

Note that any time you select a check box that pertains to a Predefined Service and you click **OK**, you must activate these changes through the Connection Activation window. You do not need to activate the Transparent Proxy selections because these do not pertain to Predefined Services. See “Predefined Services” on page 61 for a list of predefined services.

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

- Select **Permit DNS Queries** to allow Domain Name Service resolution requests and replies.
- Select **SafeMail** to allow mail traffic to flow through the Firewall.
- Select **Deny broadcast message to nonsecure interfaces** to prevent broadcast messages from being received at the nonsecure port. If your firewall’s nonsecure interface is connected to the Internet, this service can help reduce the amount of logging on the Firewall.
- Select **Deny Socks to nonsecure 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 check box, you must save it by clicking **OK** and activate it through the Connection Activation window. **Use of this Service can cause security exposures for your Firewall. Use it with extreme caution.**
- Select **Enable Telnet** to allow Transparent Proxy Telnets.
- Select **Enable FTP** to allow Transparent Proxy FTPs.

Network Objects

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

Objects can be identified by name, icon representation, type, and description. There are several types of network objects but Host and Firewall are the most common. The default network object shipped with the IBM Firewall is ″The World″. This is a global object that encompasses all possible IP addresses. After you have filled in the network configuration worksheets (see “Network Configuration Planning Worksheet” on page 8), you are ready to build objects.

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

Using the Configuration Client to Define Network Objects

To define a single network object, select **Network Objects** from the configuration client navigation tree. The Network Objects dialog box appears. Double-click **NEW**. The **Add a Network Object** dialog box appears, as shown in Figure 9.

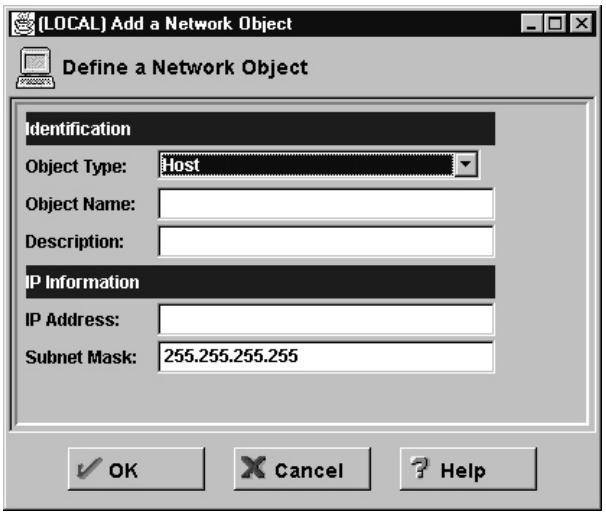

Figure 9. Add a Network Object

1. Enter the object type. Click the **Object Type** arrow to see the object types you can create. For performance reasons, it is better to create network type objects instead of host type objects. The object types you can create are:
   - Host - a particular node on your network with a mask of 255.255.255.255.
   - Network - a collective range of network addresses that is characterized by an address range and a specific subnet mask.
   - Firewall - a single machine with a firewall installed on it with a mask of 255.255.255.255. Only a firewall network object can be the target of an IBM or a manual tunnel.
   - Router - a host that routes traffic between two or more networks with a mask of 255.255.255.255.
   - Interface - a network adapter on a machine with a mask of 255.255.255.255. It does not have to be an adapter on the Firewall.
2. Fill in the object name.
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. Click **OK**.

Network Object Groups

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

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

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

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

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

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

Backing Up Your Firewall Configuration

The Firewall stores all of its configuration files in `ROOTDIR\config`. If you want to backup your firewall configuration without backing up all of the Firewall files, back up the entire contents of the `ROOTDIR\config` directory.

If you want to restore a backed up Firewall configuration, delete all of the existing files in the `ROOTDIR\config` directory and then restore the backed up versions of the files. You will have to regenerate and activate the filter rules before the restored configuration will take effect.

The key firewall configuration files are listed below. The `\config` directory on your Firewall might not contain every file listed here. Note that while most of the firewall configuration files are simple text files that can be viewed with a text editor, **manual editing of these files is not supported**.

- carriers.cfg - Pager carrier definitions
- cfgfilt.output
- explode.cfg
- filters.active - Indicates if filtering is active
- fwadpt.cfg - Definitions for network interfaces
- fwconfig.map - Contains configuration file names
- fwconns.cfg - Filter connections definitions
- fwfilters.cfg - Current active filters
- fwhttp.cfg - HTTP proxy configuration
- fwmail.conf - SafMail configuration
- fwobjects.cfg - Network objects definitions
- fwpolicy.cfg - Security Policy options
- fwrules.cfg - Filter rule template definitions

- fwservices.cfg - Services definitions
- fwsocks.cfg - Socks 5 rules from the configuration client
- fwtdefn.conf - Alert definitions
- fwtpproxy.cfg - Transparent proxy definitions
- fwusrdb.cfg - Firewall User database
- logmgmt.cfg - Archiving definitions
- modems.cfg - Modem definitions
- pager.cfg - Pager definitions
- rcsfile.cfg - Configuration Service parameters
- Socks5.conf - generated Socks 5 configuration file
- Socks5.header.cfg - User-suplied portions of generated Socks5.conf
- syslog.conf - Log facility definitions

Chapter 6. Handling Domain Name Service

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

Three domain name servers are required to accomplish this:

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

Refer to Figure 10 to see how DNS works with the IBM Firewall.

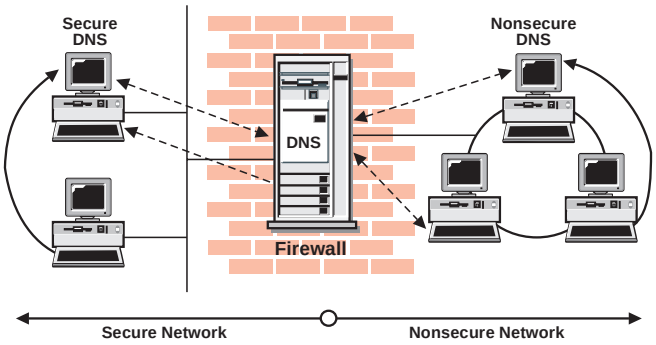

Figure 10. DNS

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

Figure 10 illustrates the Firewall's role. Anytime the Firewall needs to resolve a name for its own use, it asks the secure-side nameservers. Anytime a query is forwarded to the Firewall, it in turn forwards the query to the nonsecure nameservers.

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

Configuring DNS Using the Configuration Client

To configure DNS, select System Administration from the configuration client navigation tree. Double-click the file folder icon to expand the view. Select **Domain Name Services**. The IBM Firewall displays the current DNS configuration, which you can modify.

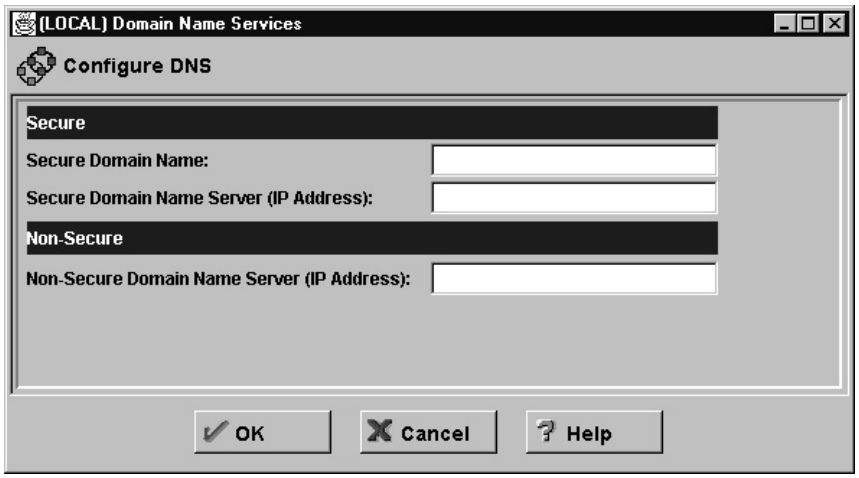

Figure 11. Domain Name Service

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

1. The **Secure Domain Name** field identifies the domain name which the Firewall will append to any unqualified hostnames.
2. The **Secure Domain Name Server** field refers to the server that resolves names and IP addresses for the hosts protected from the Internet by the IBM Firewall. You can enter dotted-decimal IP addresses, separated by spaces.
3. The **Nonsecure Domain Name Server** field refers to the server(s) provided by your service provider to resolve information that about the nonsecure network. You can enter dotted-decimal IP addresses, separated by spaces.

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

Configuring the Secure Name Server

The secure name server must be configured to forward unresolved queries to the Firewall. If you have a standard BIND implementation, add a *forwarders* statement and a *cache* statement to the *boot* file on your secure name server:

```
forwarders     aaa.bbb.ccc.ddd
cache          .              named.cache
```

Create the cache file, *named.cache*, to point to the Firewall:

```
. 99999999 IN NS firewall.private.com
firewall.private.com   99999999 IN A aaa.bbb.ccc.ddd
```

where *private.com* is the domain name used from the secure side and *aaa.bbb.ccc.ddd* is the Firewall's IP address.

In addition, you might want to add your firewall's host name to the DNS databases. This way your users can access the Firewall's Socks server, HTTP proxy, Telnet proxy, and FTP proxy using the Firewall's hostname instead of its IP address. This requires two additional steps as described in *Chapter 4* of *DNS and BIND*. See the *Bibliography* for more details about this book.

First add an A record to the domain database file:

```
firewall.private.com     IN A aaa.bbb.ccc.ddd
```

Then add a PTR record to the reverse-lookup file:

```
ddd.ccc.bbb.aaa.in-addr.arpa.   IN PTR  firewall.private.com.
```

If you do not use DNS for your secure network, your firewall must still be able to resolve its own information. Configure the firewall as described for the normal case, but list the firewall's secure interface in the **Secure Name Server** field. Then add the following line to `c:\winnt\system32\dns\boot`.

```
primary ccc.bbb.aaa.in-addr.arpa c:\winnt\system32\dns\fwnamed.rev
```

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

```
ccc.bbb.aaa.in-addr.arpa  IN SOA firewall.private.com. root.public.com. (
                          9       ; Serial
                          86400   ; Refresh after 1 day
                          300     ; Retry after 5 minutes
                          654000  ; Expire after 1 week
                          3600  ) ; Minimum TTL of 1 day
ccc.bbb.aaa.in-addr.arpa.         IN NS    firewall.private.com.
ddd.ccc.bbb.aaa.in-addr.arpa.     IN PTR   firewall.private.com.
```

Configuring the Secure Clients

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

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

Publishing Services to the Public

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

There are two ways to accomplish this. Either your service provider will list your servers as a part of their domain (and hence on their nameservers), or you must provide your own nameserver and register it with the Internet. It is by far easier for your Internet Service Provider (ISP) to provide this service for you. If you can choose this option, you need to provide them with the hostnames and IP addresses

you wish to have listed. For example, if you operate your public webserver as *www.public.com* and whose IP address is *50.100.150.200*, you need to ask your ISP to list *www.public.com at 50.100.150.200*.

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

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

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

See “Sample Configurations” for more information.

Installing Microsoft’s DNS Server

To install Microsoft’s DNS Server, go to the control panel, click **Network**, click **Services tab**, click **Add**, and select **Mircosoft DNS Server**. You will need the installation CDROM.

Troubleshooting DNS Problems

The *IBM eNetwork Firewall Reference* contains a chapter about troubleshooting the Firewall. There is a specific section in that chapter for DNS problems. This section provides suggestions for using the *nslookup* command to identify the failing segment of the DNS system.

Sample Configurations

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

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

The first example illustrates the files needed to operate the nameserver in a DMZ which is located inside the nonsecure network, as shown in Figure 12 on page 35.

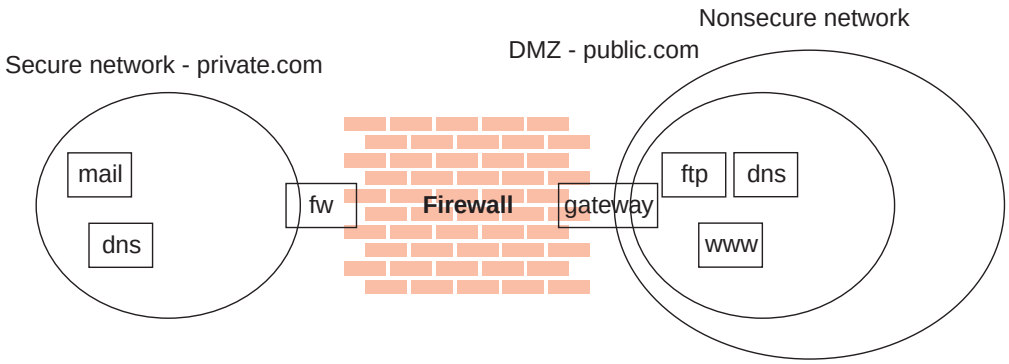

Figure 12. Nameserver in DMZ Inside Nonsecure Network

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

db.public

```
public.com.   IN SOA dns.public.com. admin.public.com.  (
                  1           ; serial number
                  10800       ; refresh after 3 hours
                  3600        ; retry after 1 hour
                  604800      ; expire after 1 week
                  86400 )     ; minimum TTL 1 day
;
; Nameservers
;
public.com           IN NS  dns.public.com.
;
; Hosts in the DMZ
;
dns.public.com.      IN A 50.100.150.202
gateway.public.com.  IN A 50.100.150.201
www.public.com.      IN A 50.100.150.200
ftp.public.com.      IN A 50.100.150.203
;
; Mail-related entries
;
public.com.          IN MX  0  gateway.public.com.
public.com.          IN CNAME gateway.public.com.
```

db.50.100.150

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

db.127.0.0

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

db.cache

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

boot

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

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

Example 2: DNS in a DMZ on a Dedicated Interface

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

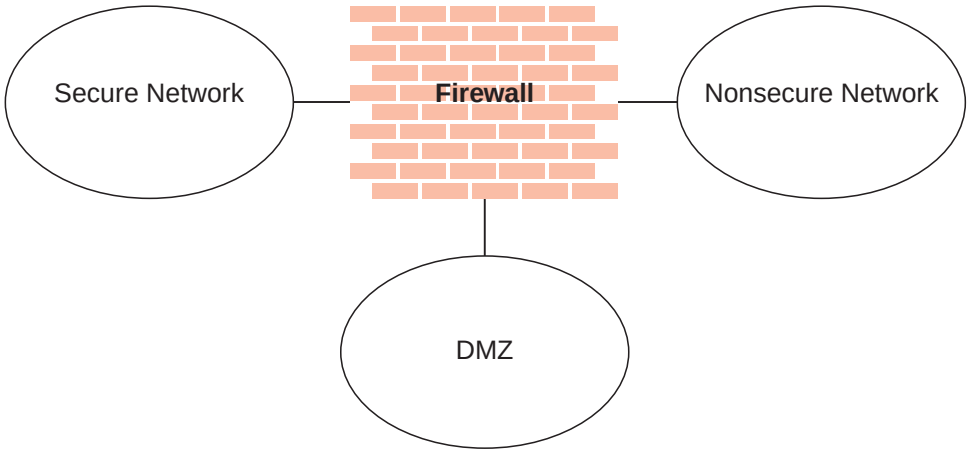

Figure 13. DNS in a DMZ on a Dedicated Interface

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

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

- DNS Client queries: --->

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

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

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

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

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

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

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

Example 3: Using the Firewall as the Secure Nameserver

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

To implement this approach, start by configuring the Firewall DNS facility using the configuration client. For the *Secure Domain Name* field, list the domain name you will be using on your secure network. For *Secure Nameserver*, list the Firewall's secure interface. For *Nonsecure Nameserver*, list the nameserver provided by your ISP, as usual. Then you must create a reverse-lookup file on the Firewall to supplement this configuration.

Create the file `c:\winnt\system32\dns\fwnamed.rev` to resemble the following example.

For this example, the Firewall's secure interface is named *fw.private.com* and its IP address is *10.100.100.1*.

```
100.100.10.in-addr.arpa.   IN SOA fw.private.com. admin.fw.private.com. (
                  1           ; serial number
                  10800       ; refresh after 3 hours
                  3600        ; retry after 1 week
                  604800      ; expire after 1 week
                  86400  )    ; minimum TTL 1 day
1.100.100.10.in-addr.arpa.        IN NS fw.private.com.
1.100.100.10.in-addr.arpa.        IN A  fw.private.com.
```

Then add the following line to `c:\winnt\system32\dns\boot`:

```
primary 100.100.10.in-addr.arpa     fwnamed.rev
```

In this scenario, your clients must be configured to indicate the Firewall (10.100.100.1) as their DNS server. Your Firewall will assist with the resolution of external information, but there will be no resolution of secure-side information. This means that any secure-side client that wants to connect to the configuration server or any of the proxy servers on the Firewall, must refer to the Firewall by IP address, not by hostname.

Chapter 7. SafeMail

The IBM Firewall SafeMail gateway provides a gateway for SMTP traffic. It relays messages from the secure mailserver(s) to the nonsecure side, hiding sensitive domain names as it goes. It relays messages from the nonsecure side in to the secure mail domain and insulates the secure network from attacks.

Although SafeMail does not perform content screening, SafeMail does provide a user exit through which content screening can be performed. For more information see “The SafeMail User Exit” on page 40.

SafeMail relays messages in real time from the sender to the receiver. This is to avoid the risks and complexity involved with maintaining a message queue on the Firewall. This necessitates certain configuration requirements upon the adjacent mail domains. In some cases, these requirements will not be practical for a particular installation. In such a case, any of several SMTP servers can be purchased separately and installed in place of SafeMail. If you choose to install a full SMTP server, configure it with security in mind. See “Using an SMTP Server Instead of SafeMail” on page 41 for more information.

Configuring SafeMail Using the Configuration Client

To configure SafeMail, select System Administration from the configuration client navigation tree. Double-click the file folder icon to expand the view. Select **SafeMail**. The IBM Firewall displays the list of configured mail servers and domains. You must configure one entry for each private-side mail domain being configured.

1. To add a domain, select **NEW** and click **Open**. The **Add Mail Server** dialog box appears.
2. The **Secure Domain Name** field contains the name by which the mail domain being described is known to users on the secure side of the firewall.
3. The **Secure Mail Server Name** field contains the host name or dotted-decimal IP address of the mail server to which this entry applies. This server must be on one of the secure networks. You can list only a single mailserver for a given domain.
4. The **Public Domain Name** field contains the name by which the mail domain being described is known to users on the nonsecure side of the firewall. This name will be substituted in place of the secure domain name, in order to hide the topography of the secure network.
5. Click **OK**.

Change a Mail Configuration Entry

To change a mail configuration entry, select an entry in the list and click **Open**. The **Change Mail Server Configuration** dialog box appears.

The **Secure Domain Name** field is disabled, but you can change the other fields, as described in “Configuring SafeMail Using the Configuration Client”.

Notes:

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

2. If you have *not* previously configured SafeMail and you specify a secure mail server here, this mail server is added to the configuration.

Delete a Mail Configuration Entry

To delete a SafeMail configuration entry, select an entry in the list and click **Delete**. You will get a delete warning. Click **OK** to delete or **Cancel**, if you change your mind.

Configuring the Secure Servers

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

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

Configuring the Public Domain

The only configuration necessary in the nonsecure network is to list your Firewall as the mail exchanger for your network. Ask your service provider to add the necessary information to their DNS servers. See "Chapter 6. Handling Domain Name Service" on page 31 for additional specifics regarding the mechanics involved.

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

The SafeMail User Exit

SafeMail provides a user exit by which an installation can tailor SafeMail to reject potentially malicious traffic. See the *IBM eNetwork Firewall Reference* for a detailed description of the Software Developers Kit provided for this purpose.

This feature allows you to create a function, *UsrCheck()*, that gets called each time SafeMail receives a packet from the sender. The function is passed a structure that contains several fields related to the state of the system. This structure includes a

unique session ID, the sending and receiving servers' IP addresses, indicators for commands previously received, and a plain text buffer containing the packet being analyzed.

Some of the types of checks that can be implemented in this function are:

- *banned* host lists
- scanning for disallowed character sequences, such as inappropriate language or project code names
- examination of embedded quoted strings
- message length restrictions

If desired, the user exit can also be used to implement an interface to a third-party content-screening product.

If the user exit function decides that a message should not be processed, the function returns a reason code back to SafeMail. SafeMail will immediately reject the connection to the sending SMTP server. At the same time, a message will be written to the firewall log, including the reason code returned by the user exit.

When writing a user exit, keep in mind that this function is called for every packet received. Take care to write it as effectively as possible, to avoid negatively impacting the performance of the system. Also, keep in mind that this function will run in a multithread environment, and therefore must be written in a thread-safe manner. You can write the user exit with any compiler which supports multithreaded operation and can use the *_cdecl* linkage convention. Sample makefiles are provided for IBM Visual Age C++ and for Microsoft Visual C++.

Using an SMTP Server Instead of SafeMail

Disabling SafeMail

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

Configuring an SMTP Server

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

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

It is also important that your mail server hide information about the secure network. According to the rules of SMTP, each server that forwards a piece of mail should insert a *Received:* header line. These header lines can be used by an attacker to

map your secure network. SafeMail strips all these headers when it processes a note; configure your SMTP server to do the same. Also, SafeMail rewrites all private-side hostnames to the public domain name. This removes even more information that could be used to map your network.

Sample Logging Output for SafeMail

The following is a sample of the logging output for SafeMail.

```
Feb 03 13:46:11 1998 mr16n18: ICA2163i: safemaild started.

Feb 03 13:41:14 1998 mr16n18: ICA2177i: SafeMail connection 0xd71e7a19
received from RACK3BD.
Feb 03 13:41:21 1998 mr16n18: ICA2179i: SafeMail has forwarded 215575
bytes for connection 0xd71e6118 from 9.67.144.52 to 9.67.131.250.
Feb 03 13:41:21 1998 mr16n18: ICA2178i: SafeMail session 0xd71e7a19
has been established from 9.67.144.52 to 9.67.131.250.
Feb 03 13:41:23 1998 mr16n18: ICA2177i: SafeMail connection 0xd71e831a
received from RACK3BD.
Feb 03 13:41:36 1998 mr16n18: ICA2177i: SafeMail connection 0xd71e901b
received from RACK3BD.
Feb 03 13:41:56 1998 mr16n18: ICA2179i: SafeMail has forwarded 215567
bytes for connection 0xd71e7a19 from 9.67.144.52 to 9.67.131.250.
Feb 03 13:41:56 1998 mr16n18: ICA2178i: SafeMail session 0xd71e831a
has been established from 9.67.144.52 to 9.67.131.250.
Feb 03 13:41:56 1998 mr16n18: ICA2178i: SafeMail session 0xd71e901b
has been established from 9.67.144.52 to 9.67.131.250.
Feb 03 13:41:57 1998 mr16n18: ICA2179i: SafeMail has forwarded 346
bytes for connection 0xd71e901b from 9.67.144.52 to 9.67.131.250.
Feb 03 13:41:57 1998 mr16n18: ICA2179i: SafeMail has forwarded 358
bytes for connection 0xd71e831a from 9.67.144.52 to 9.67.131.250.
```

Log messages indicate the following:

- ICA2177 - indicates the start of a new connection.
- ICA2179 - indicates successful termination.
- ICA2178 - indicates contact has been made with the receiving SMTP server.
- ICA2181 - indicates that SafeMail rejected the session. See the *IBM eNetwork Firewall Reference* for reason codes.
- ICA2180 - indicates the session end.
- ICA2182 - indicates that the user exit decided that the session should be rejected.

Chapter 8. Controlling Traffic Through the Firewall

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

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

Using the Configuration Client to Build Connections

You use the components of the configuration client illustrated in Figure 14 on page 44 to create network objects, rule templates, services, and connections.

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

Services
Are built of one or more rule templates. Defines the type of IP traffic that is permitted or denied between a source and destination object. For example, you could construct a service to permit Telnet or deny Ping. (One of the FTP services is comprised of eight rule templates). The IBM Firewall comes with a set of default services. You cannot delete these preloaded default services but you can modify certain fields. However, if these predefined services do not meet your needs you can add to services by using the rule templates to create new rules. See “Defining Services” on page 63 for more information.

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

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

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

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

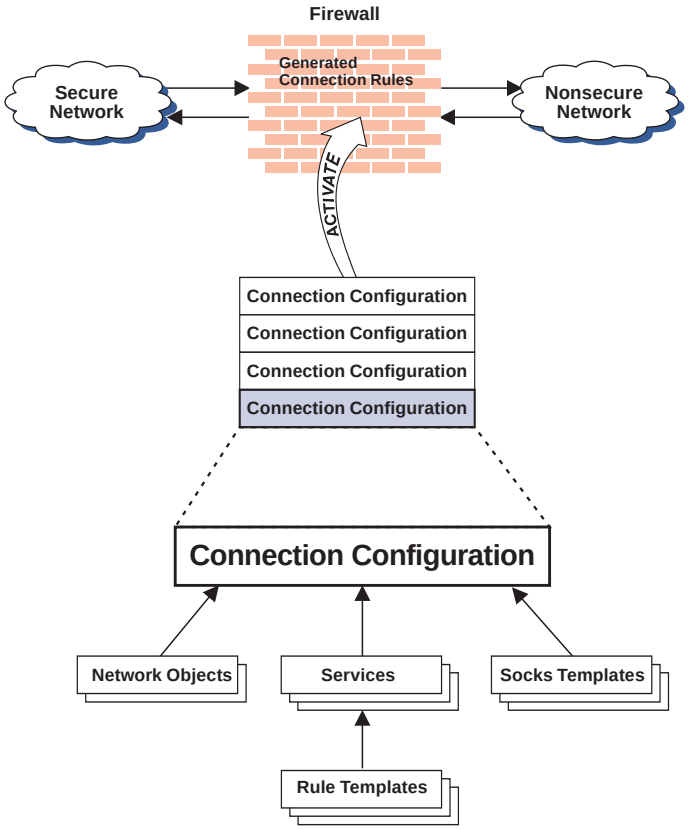

Figure 14. Building Connections

Building Connections Using Predefined Services

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

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

To build a connection, select Traffic Control from the configuration client navigation tree and double-click the file folder icon to expand the view. Select **Connection Setup**. The **Connections List** dialog box appears. Select **NEW** and click **Open**. The **Add a Connection** dialog box appears, as shown in Figure 15 on page 45.

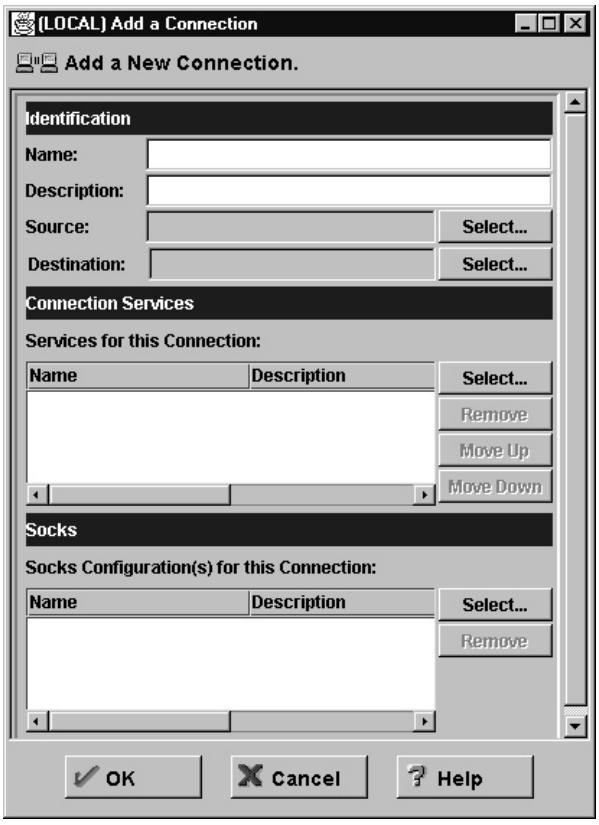

Figure 15. Add a Connection

1. Fill in a name for the connection.
2. Fill in a description of the connection.
3. For the source field, click **Select** and choose a network object from the **Network Object** dialog list.
4. For the destination field, click **Select** and choose a network object from the **Network Object** dialog list.
5. To choose the services for this connection, click **Select** and choose the type of traffic you wish to control between the endpoints.
6. Choose one or more services from the list to add the service to the Connection.
7. You can reorder the list by selecting a service and clicking **Move Up** or **Move Down**. See “Ordering Connections”.
8. You can remove a service by selecting it and clicking **Remove**.
9. Use **Socks Configuration for this Connection**. Follow steps 5–7 to make connections for Socks.
10. After you have everything defined, click **OK**.
11. Activate all of your connections. See “Connection Activation” on page 46.

Ordering Connections

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

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

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

Connection Activation

Note: Before you activate connections, make sure 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
The Firewall builds the generated connection rules from the connection configuration and activates that rule set.

Deactivate Connection Rules
The Firewall is now protected by the default rules.

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

Validate Rule Generation
The rules you have created are either valid or invalid.

Enable Connection Rules Logging
The Firewall logs selected traffic to the `firewall log` facility.

Disable Connection Rules Logging
Stops the Firewall logging.

The **Connection Activation** dialog box appears, as shown in Figure 16 on page 47.

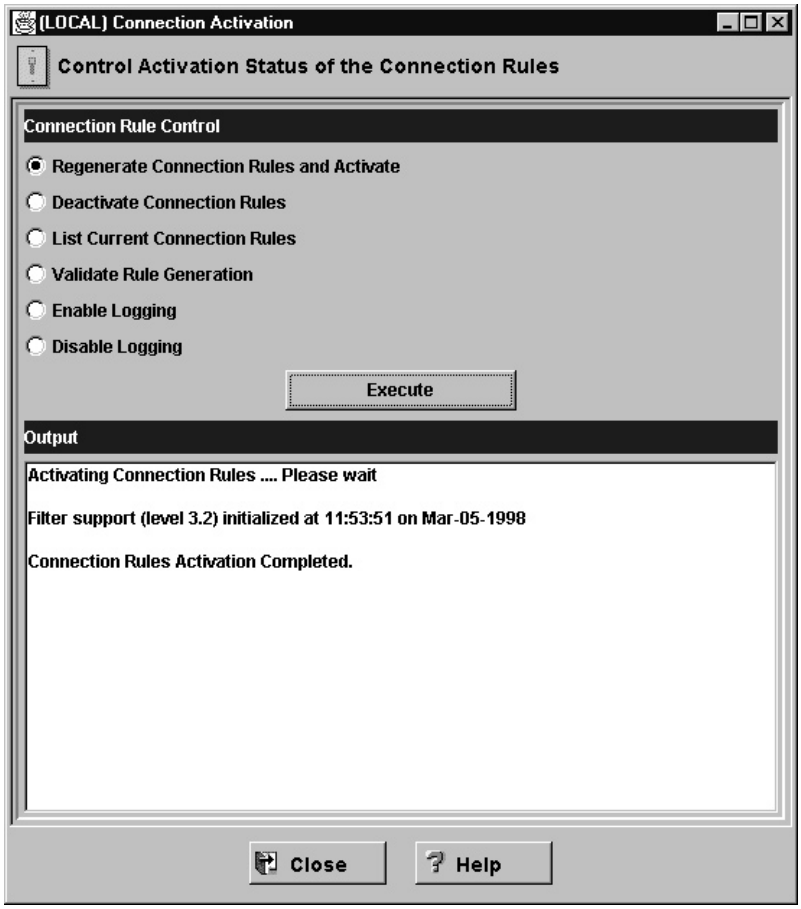

Figure 16. Connection Activation

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

Sample Logging Output when Regenerating and Activating Connection Rules

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

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

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

Feb 03 13:46:55 1998 mr16n18: ICA1037i: #:1 permit 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0
udp eq  53 eq 53 both local both l=n f=y t=0 e=none a=none
Feb 03 13:46:55 1998 mr16n18: ICA1037i: #:2 permit 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0
udp gt  1023 eq 53 both local both l=n f=y t=0 e=none a=none
Feb 03 13:46:55 1998 mr16n18: ICA1037i: #:3 permit 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0
udp eq  53 gt  1023 both local both l=n f=y t=0 e=none a=none
Feb 03 13:46:55 1998 mr16n18: ICA1037i: #:4 permit 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0
tcp     gt  1023 eq 25 both local both l=y f=y t=0 e=none a=none
Feb 03 13:46:55 1998 mr16n18: ICA1037i: #:5 permit 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0
tcp/ack eq  25 gt 1023 both local both l=y f=y t=0 e=none a=none
Feb 03 13:46:55 1998 mr16n18: ICA1037i: #:6 permit 9.67.144.49 255.255.255.240
9.67.130.154 255.255.248.0 all any 0 any 0 both both both l=y f=y t=0 e=none a=none
```

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

Determining the Rule States

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

1. The configuration is not active.

   You have not yet used the configuration client to activate the configuration or you have deactivated the configuration. This is the state of the configuration when you first install the IBM Firewall and boot your system or deactivate filter rules. Default filters are in place to protect your network from intrusion when you first install the Firewall.

   Firewall Access:

   - The default filter configuration permits all local inbound traffic and permits all outbound traffic.

2. The configuration is active but has errors.

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

   Firewall Access:

   - Permit all local inbound traffic.
   - Permit all outbound traffic.

3. The configuration is active and valid. Note that there may have been some warnings, most notably, duplicate filter rules.

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

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

   Firewall Access:

   - Access determined by the configuration file.

     Each packet that is received by, or is about to be sent by, any network interface is examined and its contents compared against each rule in the generated connection rules. When a match is found, the action (permit or deny access) on that rule is carried out.

- If no rules match the packet, there is an implied “deny all” rule that denies access.

Chapter 9. Examples of Services

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

Planning Considerations

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

As described in "Chapter 8. Controlling Traffic Through the Firewall" on page 43, endpoints are represented to the Firewall by network objects. If you have not already done so, you should complete the network planning worksheet in "Chapter 2. Planning" on page 7 and create the network objects necessary to represent your network.

The examples in this chapter use the following network objects:

Secure Interface
: The secure interface of the Firewall.

Nonsecure Interface
: The nonsecure interface of the Firewall.

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

The World
: The nonsecure network.

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

If the firewall acts as a proxy, then the firewall will perform the necessary work on behalf of the secure user and the nonsecure host(s) will never know that the secure host exists. If the firewall is to route the traffic, then the secure host and the nonsecure host will speak directly to each other.

If you will use the Firewall as a proxy, then the endpoints of your communication will include the firewall, as shown in Figure 17 on page 52.

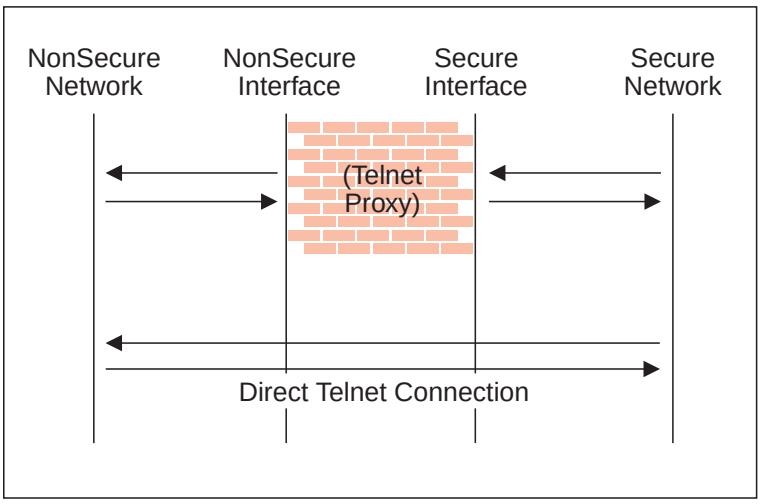

Figure 17. Telnet Proxy and Direct Telnet Connection

Example of Telnet Proxy

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

As described in Figure 17, 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 Filtered Telnet

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

Table 2. Filtered Telnet

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

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

Example of Proxy HTTP

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

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

Note, though, that when these additional protocols are used, they are wrapped in HTTP between the client and the proxy. Therefore the communication would be similar to the diagram in Figure 18.

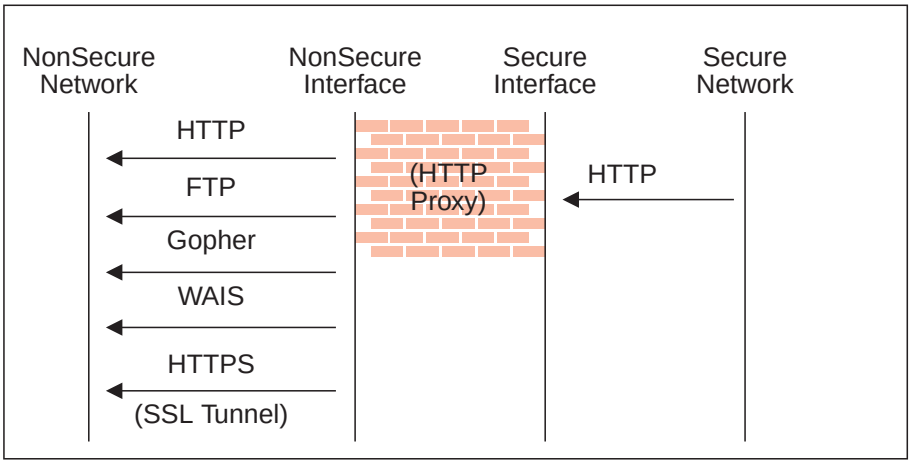

Figure 18. Proxy HTTP

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 13. Configuring Proxy Servers” on page 85.

Example of Socks

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

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

Table 4. Socks

Source Object	Destination Object	Services Required
Secure Network	Secure Interface	Socks 1/2
NonSecure Interface	The World	Select from... • HTTP proxy out 2/2 • FTP proxy out 2/2 • Telnet proxy out 2/2 (Any second-half proxy service for which you wish to provide support)
Secure Network	The World	In the Socks Configuration window, select from... • 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 11. Configuring the Socks Server” on page 67 .

Hints for DNS

Very little communication will take place efficiently if you do not provide DNS resolution. See “Chapter 6. Handling Domain Name Service” on page 31 for details on configuring DNS. Do not forget to enable ″Permit DNS Queries″ in your Security Policy.

Hints for Nonsecure Socks Clients

The Security Policy panel contains a check box 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.

If you want to allow clients to enter your network from the nonsecure network, you *must not* turn on this checkbox.

Chapter 10. Customizing Traffic Control

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

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

Using the Configuration Client to Create Rule Templates

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

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

   The IBM Firewall displays an **Add IP Rule** dialog box, as shown in Figure 19 so that you can define a rule.

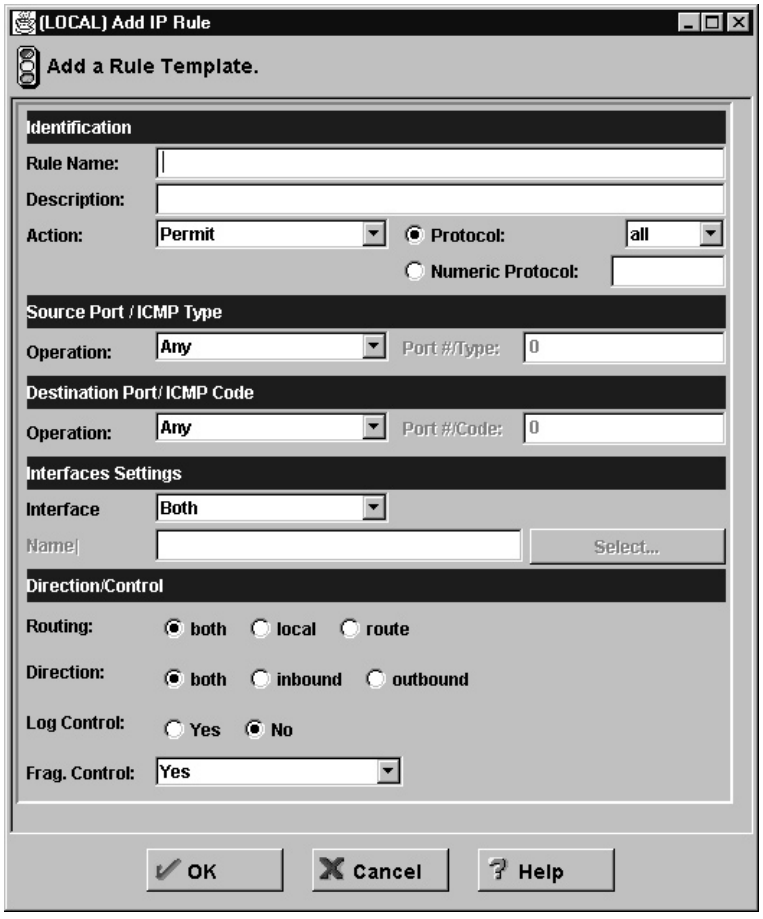

Figure 19. Add IP Rule

3. Enter the Rule Name.

4. Enter the Rule Description. This field is optional.
5. Click the action arrow and choose to either permit or deny access to the Firewall.
6. Click the protocol arrow and select from the following list:

   all Any protocol will match this rule.

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

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

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

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

   ospf The packet protocol must be open shortest path first protocol (ospf) to match this rule. When ospf is specified as the protocol, the source port operation and source port value is used for the ospf record type value. Filtering can also be performed on the ospf type. A type value of **any** can be specified and the destination port fields must be specified as **any 0**. Anything else is ignored.

   ipip The packet protocol must be IP-in-IP protocol (IPIP) to match this rule. When IPIP is specified, the port fields must be specified as **any 0**.

   esp The packet protocol must be encapsulating security protocol used by the virtual private network for sending encapsulated IP packets to match this rule.

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

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

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

   The values for operation are:

   - Any

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

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

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

Here are some of the ICMP types and codes:

Type	Code and Description
0	0 - Ping reply
8	0 - Ping request
3	1 - Host unreachable
3	3 - Port unreachable
5	1 - Redirect for host

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

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

   secure
   For packets coming or going on the secure interface

   nonsecure
   For packets coming or going on the nonsecure interface

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

10. If you choose specific for the interface type, the name of the specific interface will appear in the Name field.
11. Click the desired routing:

    both Applies to all traffic.

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

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

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

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

12. Click the desired direction:

    both For packets going out from or into the selected interface

    inbound
    For packets coming into the selected interface from the network

    outbound
    For packets going out from the selected interface to the network

13. If you choose Yes for the Log Control field, every packet that matches that rule is recorded in the `firewall log` with priority level `Error`. If this parameter is not specified, the default is no.
14. Click the **Fragment Control** arrow to choose the desired fragment control. For IP packet information to match a rule fragmentation control specification, the control is interpreted as follows:

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

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

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

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

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

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

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

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

Change IP Rule Configuration Entry

To modify an IP rule that you have created:

1. Double-click on an existing rule in the **Rules List**. The **Modify IP Rule Configuration** dialog box appears.
2. Modify the appropriate fields as described in “Chapter 10. Customizing Traffic Control” on page 57 and click **OK** to apply the changes.

Delete Rule Configuration Entry

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

Predefined Services

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

The preloaded default services are:

All non-secure
Deny all traffic across nonsecure interface

All permit
Permit all traffic(for debugging purposes only)

All permit, in one direction
Permit all traffic(for debugging purposes only)

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

All shutdown
Deny all packets (shutdown or debug)

Anti Spoofing
Deny inbound nonsecure packets with secure source address

Broadcasts
Deny broadcast messages to nonsecure interface

Config Client non-secure
Permit use of the configuration client from nonsecure network

Config Client secure
Permit use of the configuration client from secure network

CU-SeeMe
CU-SeeMe Video on default ports 7649 and 7648

DNS queries
(SECURITY POLICY) Permit DNS queries

DNS transfers
Permit DNS zone transfers (for secondary name servers)

Domain Controller Authentication
Allows use of Domain Controller for user authentication

FTP proxy in 1/2
Permit FTP inbound from nonsecure network to Firewall

FTP proxy in 2/2
Permit FTP inbound from Firewall to secure network

FTP proxy out 1/2
Permit FTP outbound from secure network to Firewall

FTP proxy out 2/2
Permit FTP outbound from Firewall to nonsecure network

Gopher proxy in 2/2
Permit gopher from Firewall to secure network

Gopher proxy out 2/2
Permit gopher from Firewall to nonsecure network

HTTP deny non-secure
Deny HTTP to nonsecure interfaces

HTTP direct out
Permit HTTP from secure network directly to nonsecure network

HTTP proxy in 2/2
Permit HTTP from Firewall to secure network

HTTP proxy out 1/2
Permit HTTP (port 8080) from secure network to the Firewall

HTTP proxy out 2/2
Permit HTTP from Firewall to nonsecure network

HTTPS direct out
Permit HTTPS (SSL) from secure network to nonsecure network

HTTPS proxy out 2/2
Permit HTTPS (SSL tunnel) from Firewall to nonsecure network

IDENTD
Permit user identification with Socks protocols

Mail (SECURITY POLICY) Permit Mail traffic through the Firewall

NetBT Name Services broadcasts
Allows NetBIOS over TCP/IP Name Services broadcasts

Ping Permit Ping outbound secure network to anywhere

SDI authentication
Permit connection to SecurID ACE server in the secure network

Socks 1/2
Permit use of Socks from secure network to the Firewall

Socks deny non-secure
Deny Socks from nonsecure adapters

Socks in 1/2
Permit use of Socks from nonsecure network to the firewall

Telnet direct out
Permit Telnet outbound from secure network to nonsecure network

Telnet proxy in 1/2
Permit Telnet inbound from nonsecure network to the Firewall

Telnet proxy in 2/2
Permit Telnet in from the Firewall to the secure network

Telnet proxy out 1/2
Permit Telnet out from secure network to Firewall

Telnet proxy out 2/2
Permit Telnet out from Firewall to nonsecure network

VDOLIVE Direct In
Permit nonsecure client to secure server

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

VDOLIVE Direct Out
Permit secure client to nonsecure server

WAIS proxy in 2/2
Permit WAIS (z39.50) from the Firewall to the secure network

WAIS proxy out 2/2
Permit WAIS (z39.50) from the Firewall to the nonsecure network

Defining Services

After you have defined a rule(s), you need to add the rule(s) to a service. Select Traffic Control from the configuration client navigation tree and double-click on Connection Templates, then select Services. The Services List dialog box appears. Double click NEW to get the Add Service dialog box, as shown in Figure 20 on page 64 .

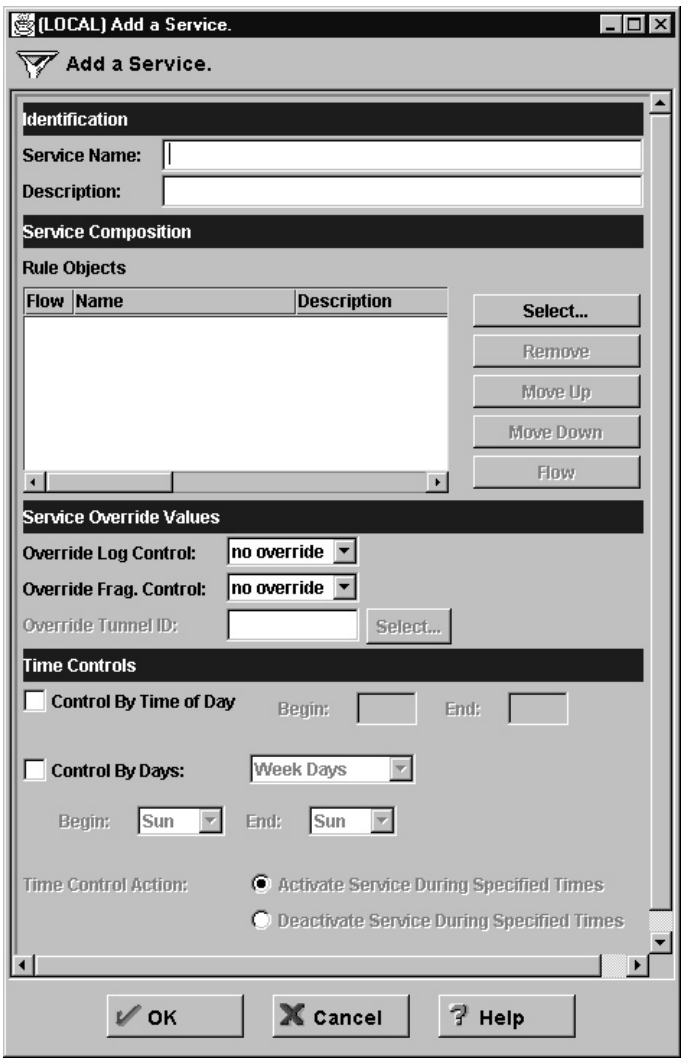

Figure 20. Add a Service

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 time controls allow you to associate a time range with each service. Therefore, this service will only be valid in a specified time period. If there is no time specification for a service, that service is valid all the time.

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

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

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

   Choose **Deactivate Service During Specified Times** if you want this service to be deactivated during the specified times. This service will be activated during the times outside of those specified.
6. Click **Select** to choose the rules that comprise this service.
7. 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.
   ```
8. When a packet is received, the IBM Firewall compares the information in the packet to the rules in the rules configuration file starting at the top of the file. It stops comparing when the first match is found and performs the action contained in the rule.

   Once you have added a series of rules to the service, you can change their order. Select a rule from the **Service Objects** list and click the **Move Up** or

Move Down buttons to reposition the rule. Or you can remove a rule by clicking **Remove**. The configuration client displays a refreshed list of rules. Click **OK** to save your changes.

Chapter 11. Configuring the Socks Server

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

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

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

The Socks server is similar to the proxy server. But while the proxy server actually performs the TCP/IP function at the Firewall, the Socks server just identifies the user and redirects the function through the Firewall. The actual TCP/IP function is performed at the client workstation, not at the firewall. This saves processing in the Firewall. The users in the secure network can use the many TCP/IP products that support the socks standard. Figure 21 illustrates the Socks server intercepting an HTTP request from a client within the secure network.

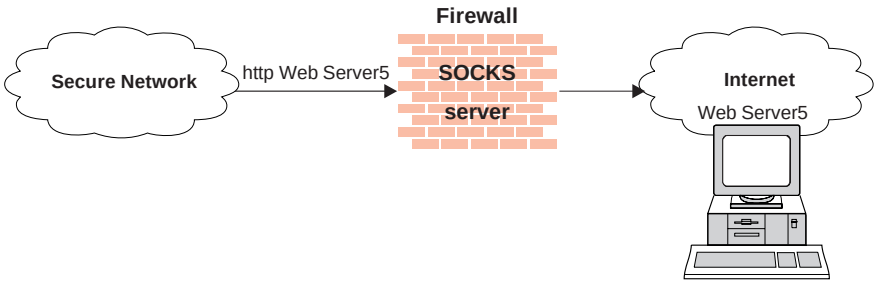

Figure 21. The Socks Server

The Socks server effectively hides your internal IP addresses from the outside world.

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

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

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

1. The most permissive profile does not enable outbound authentication and permits any user, whether using a Socks Protocol Version 4 or Socks Protocol Version 5 client to connect. In this scenario inbound connections are denied.
2. The migration profile allows Socks Protocol Version 4 users to pass unauthenticated, but requires Socks Protocol Version 5 users to authenticate. Inbound Socks Protocol Version 4 connections are denied and inbound Socks Protocol Version 5 connections are required to authenticate. This is the default profile.
3. The most secure profile requires that all users use Socks Protocol Version 5 clients and provide valid authentications.

When the Firewall is installed, the socks server 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 54, for an example of how to set up a socks service.

Protocols Supported by Socks Protocol Version 5 Server

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

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

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

Configuring the Socks Server Using the Configuration Client

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

Add a New Socks Rule

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

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

   The **Add a Socks Rule** dialog box appears, as shown in Figure 22.

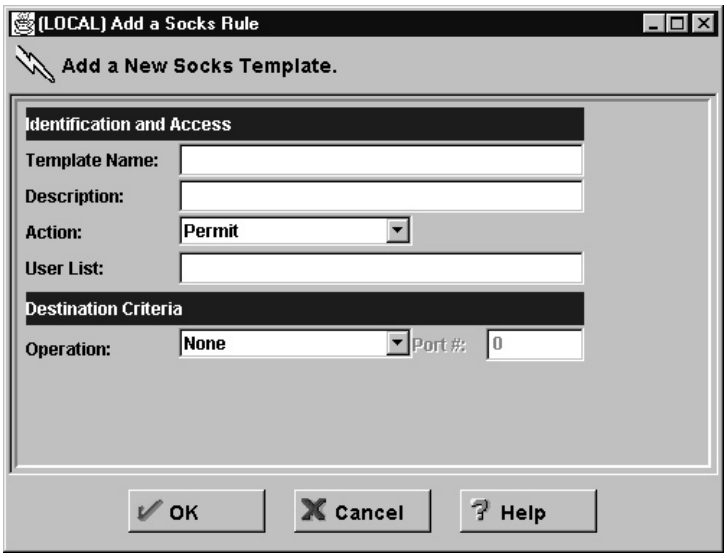

Figure 22. Add a Socks Rule

2. In the **Template Name** field, enter the name of the socks entry. This name must be unique and should not contain a pipe symbol(|), a single quote (or apostrophe) character (') or a double quote(″) character as these are used as file delimiters. Use of these characters will result in unreliable data.
3. Fill in a description.
4. Click the Action arrow and choose to either permit or deny access from a source to a destination.

   When a datagram comes into the socks server, the server compares the datagram specifications to each rule in the configuration file starting with the first rule until it finds a rule that matches exactly. Then it stops searching and performs the relevant action (either permit or deny access) on that rule. If no match is found, access is denied automatically.
5. In the **User List** field, you can enter a user ID or a list of user IDs. If you enter a list, separate the entries with commas. Do not use spaces, tabs, the pipe symbol (|), or double quotes(″) in the user list.
   - The user list is limited to 396 characters.
   - User IDs must be IDs of users on the requesting host, not those on the destination host or socks server host.
   - A user ID can consist of 1 to 8 characters, including:
     – a through z
     – A through Z
     – 0 through 9

– _ (underscore)

6. A user ID should not contain the following characters pipe symbol (|) double quote character(″).
7. 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.
8. In the **Operation** field, enter the logical operation to be performed on the port number:

   eq Equal to

   neq Not equal to

   lt Less than

   gt Greater than

   le Less than or equal to

   ge Greater than or equal to

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

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

Modify a Socks Rule

1. Double-click on an entry on the **Socks** dialog box.

   The **Modify a Socks Rule** dialog box appears.
2. Change the appropriate fields as described in “Add a New Socks Rule” on page 69 , and click **OK**.

Delete a Socks Rule

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

Activate Connection Rules

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

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

Sample Logging Output for Socks

The following is a sample of the logging output for Socks.

```
Feb 03 13:46:20 1998 mr16n18: ICA3123i: Sockd server starting
Feb 03 13:47:31 1998 mr16n18: ICA3010i: Session start
Feb 03 13:47:31 1998 mr16n18: ICA3011i: Session start
Feb 03 13:49:15 1998 mr16n18: ICA3007i: Too many threads
Feb 03 13:58:31 1998 mr16n18: ICA3015i: Session termination
```

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

Socks-Server Chaining

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

To set up socks-server chaining with the Socks server, edit the `socks5.header.cfg` file. This file resides in the Firewall's `config` subdirectory. Add the following:

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

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

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

```
no proxy 10.0.0.0/255.0.0.0 - - -
socks4      0/0   - socks.private.com 1080
```

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

Chapter 12. Administering Users at the Firewall

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

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

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

Adding a User to the IBM Firewall

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

Types of Users

The IBM Firewall divides users into three categories:

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

Firewall Administrators
Can use the Firewall proxy services, but they can also configure the Firewall by using the configuration client and by logging on to the Firewall from a remote host. Like proxy users, firewall administrators cannot perform local logins to the Firewall machine.

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

Primary Firewall Administrators
Have the same capabilities as firewall administrators. They can also perform local logins to the Firewall machine. Primary firewall administrators can create and modify definitions for other firewall administrators.

Types of Databases

There are two types of user databases.

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

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

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

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

Windows NT User database

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

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

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

Using the Configuration Client to Add a User

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

1. From the configuration client navigation tree, select Users. The **User Administration** dialog box appears.
2. Select **New** from the **User Administration** dialog box and click **Open**. The **Add User** dialog box appears, as shown in Figure 23 on page 75.

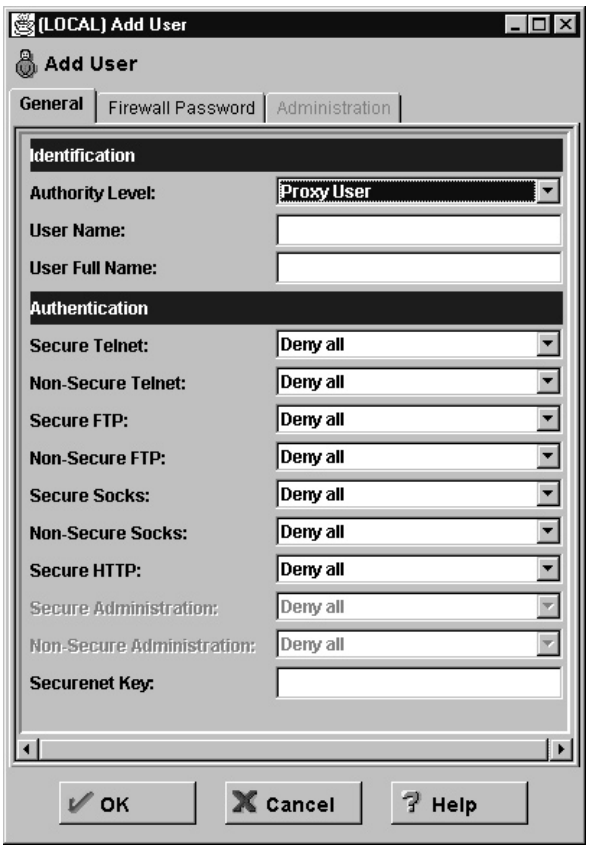

Figure 23. Add User

3. Provide this information:

Authority Level

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

Socks/Proxy User

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

Firewall Administrator

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

Primary Firewall Administrator

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

The primary firewall administrator is defined by creating a user in the NT database and making that user a member of the NT

administrators group. Modify the fwdfadm record to define the attributes for the primary firewall administrator.

User Name

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

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

a through z

A through Z

0 through 9

_ (the underscore)

User names are not case sensitive.

The Firewall comes with two preinstalled users:

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

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

   The administrator can view fwdfuser or change the assigned authentication method using the configuration client or the command line. However, fwdfuser cannot be deleted and must always exist at the firewall. In addition, firewall password and SNK are not valid authentication types for fwdfuser. For more information, refer to the *IBM eNetwork Firewall Reference.*

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

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

User Full Name

Specifies a description of the user.

The following fields refer to authentication methods. Click the arrows to select from the list of authentication methods. They are explained in "User Authentication Methods" on page 78.

Secure Telnet

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

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

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

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

Secure Socks
Specifies the Socks V5 authentication method for Socks client connections coming from the secure side of the firewall. Click the arrow to select from a list of choices. They are explained in "User Authentication Methods" on page 78.

Non-Secure Socks
Specifies the Socks V5 authentication method for Socks client connections coming from the nonsecure side of the firewall. Click the arrow to select from a list of choices. They are explained in "User Authentication Methods" on page 78.

Secure HTTP
Specifies a user ID/password type of authentication on outbound HTTP proxy requests. Click the arrow to select from a list of choices. They are explained in "User Authentication Methods" on page 78.

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

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

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

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

SecureNet Key
Specifies the character sequence to be entered by a remote user who has an 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**.

See "Authentication Methods" on page 81 for more information.

User Authentication Methods

The choices for user authentication are:

Deny All
The user is denied access.

Permit All
No authentication is needed.

NT Logon Password
NT logon password is less secure than the firewall password. However, if users are already defined in an Windows NT domain, you can use the Windows NT logon password so the user does not need multiple passwords.

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

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

Set up this connection using the following predefined services:

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

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

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**.

See “Authentication Methods” on page 81 for more information.

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.

See “Authentication Methods” on page 81 for more information.

User-Supplied Authentication 1, 2, and 3
Authentication is supplied by the user. You can install up to three user-supplied authentication methods on the Firewall. For information on how to create and compile a subroutine for user-supplied authentication, refer to the *IBM eNetwork Firewall Reference*.

Firewall 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.

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

Require User to Change

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

Lock Password

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

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

Notes:

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

   Warning Days Before Expiration (days)

   Number of days before a password expires in which the Firewall will give the user the option to change the password.

   Maximum Weeks Before Expiration

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

   Maximum Weeks Before Lockout

   Number of weeks in which the password is not used before it is locked out.

   Maximum Login Retries Allowed

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

   Passwords Before Reuse

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

   Weeks Before Password Reuse

   Number of weeks passwords are kept in the password history list. The password cannot be changed to any password that is currently in the history list.

   Minimum Length

   Minimum number of characters in a password.

Minimum Alphabetic Characters
Minimum number of alphabetic characters in a password.

Minimum Other Characters
Minimum number of non-alphabetic characters in a password.

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

Minimum Different Characters
Minimum number of different characters in the password.

Click the **Firewall Password** tab to customize these values for each user, as shown in Figure 24.

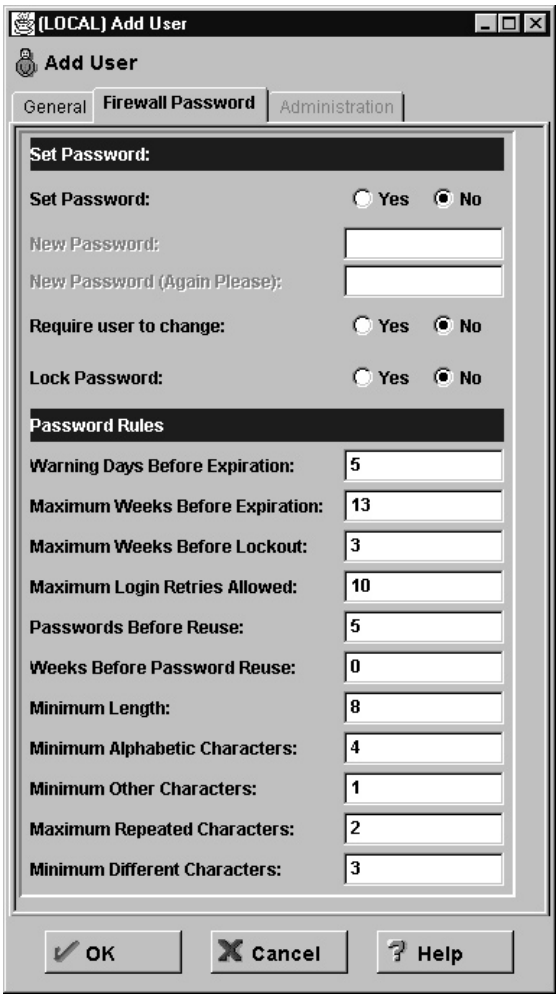

Figure 24. Firewall Password Tab

Changing a User's Access

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

1. Select the user you want to change from the **Users** dialog box and click **Open**.

2. When the **Modify User** dialog box appears, change the appropriate fields. See “Adding a User to the IBM Firewall” on page 73 for a list of user attributes that you can change.
3. When you have made the changes, click **OK**.

Deleting a User from the IBM Firewall

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

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

Administrator Authority Level by Function

Only the *primary firewall administrator* can create and modify administrators and determine which firewall functions they will have authority over. For example, you can limit a particular administrator to just having the authority to perform the Users and Log Monitor functions.

On the **Add User** dialog box, select Firewall Administrator for the **Authority Level** field. See “Adding a User to the IBM Firewall” on page 73 for more details on completing the **Add User** dialog box.

Then, select the **Administrator** tab at the top of the **Add User** dialog box. Select which functions the administrator is authorized to use.

Authentication Methods

The following are various user authentication methods.

Deny All

The IBM Firewall prohibits access to the server.

Permit All

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.

Firewall Password

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

```
Password:
```

Enter your firewall 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 users use FTP initially, SecurID card authentication will fail because FTP does not have the option to allow a password change. Users must use telnet the first time they try to do SecurID card authentication through which they will create a PIN. Users can use that PIN subsequently for later authentications like FTP, HTTP, and so forth.

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 an Assurenet Pathways SecureNet Key card. When you initialize the SNK card, use the following:

- Display format (hexadecimal)
- ERASE capability (on or off)
- Single digit challenge capability (off)

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
   ```

SecurNetKey (SNK) has been renamed Defender Handheld Token** (DHT) by AXENT** Technologies.

NT Logon Password

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

User-Supplied Authentication 1, 2, and 3

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

Chapter 13. Configuring Proxy Servers

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

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. Also, it is not a caching proxy. Nothing is stored on the firewall on behalf of an HTTP request.

Persistent Sessions

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

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

If determined incorrectly, this could result in a display indicating traffic on the connection when there is none. An example of this is the animated icon of a browser that continually runs even though the complete page has been loaded. Click **Stop** to halt the animation. See "Maximum Persistent Requests" on page 87 and "Persistent Connection Timeout" on page 87 for information on these parameters.

Configuring HTTP Proxy Using the Configuration Client

To configure HTTP Proxy, do the following:

1. 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.
2. Activate filters
3. Add a connection. See "Example of Proxy HTTP" on page 53 for an example of how to set up a connection on the nonsecure side of your network.
4. To configure HTTP Proxy, select HTTP from the configuration client navigation tree. The IBM Firewall displays the **HTTP Proxy** dialog box, as shown in Figure 25 on page 86.

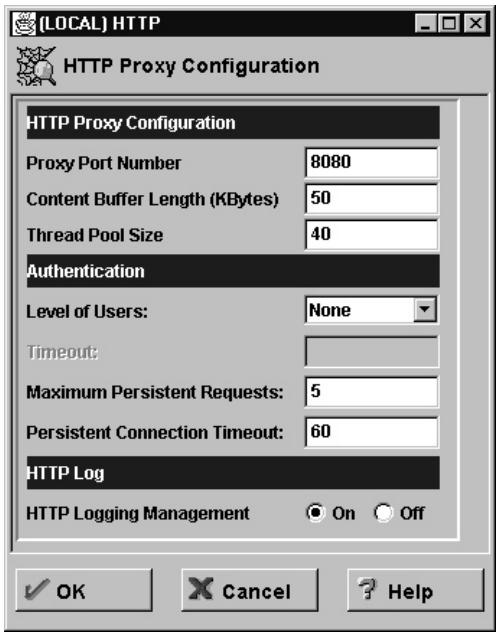

Figure 25. HTTP

5. To stop the proxy, select `my computer/control panel/services`. Choose the IBM Firewall HTTP Proxy and click *Stop*.

   The executable phttpd is a system service that starts automatically when the system is started.

Configure the parameters on the **HTTP Proxy** dialog box. If you change any parameters, the Firewall HTTP proxy service will stop and start again. Active proxy users will have their requests terminated until the proxy restarts (a few seconds of time).

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 8088.

The default filter rules are set to disallow inbound, nonsecure traffic on port 8080, but allow secure traffic on that same port. The proxy will reject only nonsecure proxy requirements. The default is `8080`. 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.

Max Content Buffer Length

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

Specify the value in kilobytes (K). The default is `50K`.

Thread Pool Size

Use this parameter to set the fixed number of threads that you want to have active at one time. 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. Specify a whole number like 60, for example. The default is `200`.

Level of Users

This parameter tells the proxy what level of users to authenticate. Specify the value as either all, new, or none. The default is none. The values are:

all All browsers will be sent the proxy authenticate response to indicate that the browser should prompt the user for a userID and password. If the browser does not support the proxy authenticate response, the error page displays indicating this. If the browser supports it, the userID and password prompt will be displayed.

new Is used as a migration aid. It will only send back a 407 proxy authenticate response, to tell the browser to issue a userid/password prompt, to a client browser that identifies itself as an HTTP/1.1 browser. You can set a switch in Internet Explorer 4.0 so that it will broadcast requests with the HTTP/1.1 identifier. Netscape and others identify themselves as HTTP/1.0 requests.

none Does not check browser requests. Does not prompt for any userid/password.

Timeout

This parameter tells the proxy how much time to wait for a client request before requiring the user to reauthenticate himself. A user is authenticated from the specific IP address and userID given at the time of the original authentication for this period of idle time. Specify the time in minutes. The default is 60 minutes.

As long as the user is actively browsing, this time window will not expire.

Maximum Persistent Requests

This parameter indicates the maximum number of requests that a proxy can receive on an HTTP/1.1 persistent connection. This is a performance tool that directly impacts the authentication timeout. While in a persistent session, no test of the authentication of a user is done until the persistent session ends. Specify the value as a whole number, for example 25. The default is 5.

Persistent Connection Timeout

This parameter indicates the time in seconds to keep an HTTP/1.1 persistent connection with a client browser once an HTTP/1.1 compliant browser starts a session with the proxy. This is a performance tool that directly impacts the authentication timeout. While in a persistent session, no test of the authentication of a user is done until the persistent session ends. Specify the time in seconds. The default is 60.

HTTP Logging Management

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

Browser Configuration

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

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

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

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

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:

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

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

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

For more information, see “Example of Proxy HTTP” on page 53.

Methods Supported

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

- FTP
- Gopher
- HTTP
- HTTPS
- WAIS

Sample Logging Output for HTTP Proxy

The following is a sample of the logging output for HTTP Proxy authenticated get requests.

```
Mar 06 14:04:50 1998 fire3: ICA2140i: httpd --> HTTP Proxy authentication UNSUCCESSFUL
for user <Unknown>, on 9.67.140.162, thru secure network ... RC:30.
Mar 06 14:04:50 1998 fire3: ICA2099i: httpd --> Status: 407 from client
9.67.140.162, who requested "GET http://9.67.128.69/ HTTP/1.1" for 0 bytes.
Mar 06 14:05:05 1998 fire3: ICA2024i: User fred successfully authenticated
using NT authentication from secure network:9.67.140.162.
Mar 06 14:05:05 1998 fire3: ICA2169i: User fred successfully authenticated
for HTTP Server using NT from secure network:9.67.140.162.
Mar 06 14:05:05 1998 fire3: ICA2140i: httpd --> HTTP Proxy authentication
SUCCESSFUL for user (fred), on 9.67.140.162, thru secure network ...RC:1.
Mar 06 14:05:05 1998 fire3: ICA2099i: httpd --> Status: 200 from client
9.67.140.162, who requested "GET http://9.67.128.69/HTTP/1.1" for 2693 bytes.
Mar 06 14:05:05 1998 fire3: ICA2140i: httpd --> HTTP Proxy authentication
SUCCESSFUL for user (fred), on 9.67.140.162, thru secure network ...RC:1.
Mar 06 14:05:06 1998 fire3: ICA2099i: httpd --> Status: 200 from client
9.67.140.162, who requested "GET http://9.67.128.69/Admin/lgsplash.gif HTTP/1.1"
for 211 bytes.
Mar 06 14:05:10 1998 fire3: ICA2140i: httpd --> HTTP Proxy authentication
SUCCESSFUL for user fred, on 9.67.140.162, thru secure network ...RC:1.
Mar 06 14:05:10 1998 fire3: ICA2099i: httpd --> Status: 200 from client
9.67.140.162, who requested "GET http://9.67.128.69/Admin/lgmast.gif HTTP/1.1"
for 211 bytes.
```

Logging activity is explained as follows:

- ICA2099i - shows a return code of 407 and means that the authentication failed for that get request.

  The browser then asks the user for some authentication. The browser asks for a userid and password.
- ICA2140i - the authentication was successful for user fred.

The authentication occurs on each get request for every element on the Web page.

FTP

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

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

   The proxy server will ask for your user name:

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

   ```
   login: jane_doe
   ```

The server validates your identity depending on the authentication scheme selected when your user name was added to the Firewall (see “Adding a User to the IBM Firewall” on page 73). See “Authentication Methods” on page 81 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. This is probably a different user name and password from those you used to FTP to the Firewall.

The default timeout value for login is 60 seconds and for idle proxy is 7200 seconds. To change the default timeout values see “Overriding Timeout Values in FTP and Telnet Proxies” on page 91.

Transparent FTP

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

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

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

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

   ```
   USER:
   ```

3. Enter your user name at the nonsecure network:

   ```
   USER: username@remote_site_host_name
   ```

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

   ```
   password:
   ```

5. Enter your password.

The default timeout value for login is 60 seconds and for idle proxy is 7200 seconds (two hours). To change the default timeout values see “Overriding Timeout Values in FTP and Telnet Proxies” on page 91.

Telnet

Use the telnet proxy to login to the firewall proxy server. You can use either the host name or Internet address. Then, after your credentials are authenticated, you use the telnet command at the Firewall to log in to the intended host. For example, let’s use telnet from inside the secure network, through the Firewall with the host name of `telnet_gw`, to access your ultimate destination, `forhost.network.outside.com`.

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

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

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

   ```
   login:
   ```

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

   ```
   login: jane_doe
   ```

The server validates your identity depending on the authentication scheme selected when your user name was added to the Firewall (see “Adding a User to the IBM Firewall” on page 73). See “Authentication Methods” on page 81 for information about how users are authenticated by proxy servers.

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

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

```
ENTER DESIRED HOST:
```

Type

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

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

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

Transparent Telnet

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

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

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

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

   ```
   Login:
   ```

3. Enter your user name at the nonsecure network:

   ```
   Login@remote_host
   ```

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

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

Overriding Timeout Values in FTP and Telnet Proxies

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

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

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

```
telnet
proxyTimeout=7200
loginTimeout=60

ftp
proxyTimeout=7200
loginTimeout=60
```

Chapter 14. Monitoring the Firewall Logging

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

The IBM Firewall monitors the messages sent to the firewall log for potential crisis situations, based upon user-defined thresholds. In the event of a threshold violation, the Firewall delivers an alert, in a manner specified by the firewall administrator.

Threshold Definitions

A threshold consists of count and time parameters — if a count (number of specific 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

All thresholds can be configured using the configuration client or the command line interface. Any changes to the threshold definitions are picked up automatically by the IBM Firewall.

Alert Messages

When a threshold has been reached, the IBM Firewall generates an alert message. Delivery of the alert message can take any of the following four forms:

1. Entry in a log file:
   - Through the `alert log` facility configurable through the configuration client or the command line.
   - In the `firewall log`
2. Send an e-mail message to a list of users
3. Pager, as configured. See “Pager Notification Support” on page 95.
4. Execution of a user-defined command, with the alert message as the first parameter

The alert message contains information relevant to the particular threshold violation. For example:

```
ICA0001e: ALERT — 20 authentication failures.
ICA0002e: ALERT — 10 authentication failures for user root.
ICA0003e: ALERT — 15 authentication failures from host 56.67.78.89
ICA0004e: ALERT — Tag ICA1234e with 3 log entries.
```

Alert messages and other messages originated by the Log Monitor are not monitored.

Configuring Log Monitor Using the Configuration Client

This section describes how to use the configuration client to configure the real-time log monitor. Select System Logs from the configuration client navigation tree. Double-click the file folder icon to expand the view. Click **Log Monitor Thresholds**.

From the **Log Monitor Threshold Administration** dialog box, you can add, change, or delete a threshold definition.

Add Log Monitor

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

1. Click the **Class type** arrow to choose from the list of class types. Class types are:
   - Mail notification
   - Execute command
   - Per User Authentication Failure Threshold
   - Total Authentication Failure Threshold
   - Per Host Authentication Failure Threshold
   - Message Threshold
2. If you selected class type: Mail Notification, enter an e-mail address. You can define multiple mail notification classes.

   All threshold violation messages are sent to the specified e-mail address.
3. If you selected class type: Execute Command, fill in a command filename.

   The log monitor will execute this command with the alert message as its first parameter. You can only define one execute command class.
4. If you selected class type: Message Threshold, fill in a message tag, a standard tag from the IBM Firewall log messages that you want to be monitored.
5. If you selected one of the threshold classes, fill in the threshold count field.

   The threshold count is the maximum number of failed events allowed within the specified time period.
6. If you selected one of the threshold classes, fill in the threshold time field.

   The threshold time is the number of minutes beginning with the first occurrence of an event.
7. If you selected one of the threshold classes, click Yes or No to indicate whether you want pager notification to be active.
8. Filling in a comment is optional.
9. Click **OK**.

Change a Threshold Definition

To change a threshold definition, select the item to be changed from the **Log Monitor Threshold Administration** dialog box and click **Open**. The **Change Log Monitor** dialog box appears.

1. Enter the changes you want for the threshold count and threshold time fields.

The threshold count is the maximum number of failed authentication messages to be detected within the specified time period. The threshold time is the number of minutes beginning with the first occurrence of a message.

2. Click **OK**.

Delete a Threshold Definition

To delete a threshold definition, select the item to be deleted from the **Log Monitor Thresholds** dialog box and click **Delete**. You will be asked to confirm the deletion. Click **Yes** to confirm. Note that delete does not mean delete from the log file. It means delete the definition.

Pager Notification Support

The Firewall can page a system administrator by sending a message to the administrator's beeper when there are intrusion alerts on the Firewall. To set up pager notification support, you need to configure the following three pager components.

1. Command Customization - This component must be created and modified using the configuration client. It sets defaults for the pager command, which is used by the log monitor and can be used from the command line. This component will contain a unique entry that defines the pager environment. See "Command Customization" on page 96 for more information on defining and customizing this component.
2. Carrier Administration - You must define a suitable carrier before connecting your modem. This component contains a list of default carriers used in the U.S. If the carrier you are using is not one of these, then add your carrier in this component. See "Carrier Administration" on page 97 for more information.

   Validate the existing phone numbers for the carriers by getting these numbers from your carriers. When talking with your carriers, be sure to get the carrier's modem phone number and other settings that are valid for the particular service you have purchased.
3. Modem Administration - Before connecting your modem, you must create suitable modem definitions. These definitions will contain all relevant modem information that pager notification support will use. This component contains a list of modems that you can choose from. You can add to this list, however some modems might not be compatible with your carrier's support. See "Modem Administration" on page 99 for information on maintaining modem definitions.

Note: IBM Firewall supports the Tele-AlphaNumeric Protocol (TAP) communications protocol for pager notification support.

What Carriers and Modems are Supported

The carriers database file contains a list of the carriers and related transmission parameters. You can add other carriers. Some of the parameters, besides the carrier's name and modem phone number are:

- The maximum message length for an alphanumeric pager and the maximum digits for a numeric pager
- The baud rate, parity, data and stop bits length

Before using a particular carrier, make sure that the carrier uses the TAP protocol.

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

- IBM MOD 448 14400 bps
- IBM 5853 2400 bps
- IBM 7852 28800 bps
- IBM 7855 2400 bps
- Generic Hayes compatible
- US Robotics Courier 9600 bps
- Zoom V.34

Configuring Pager Notification Support

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

Before starting, you need to get the correct modem phone numbers, pager ID, and modem parameters from your carrier.

To configure pager notification support, select System Administration from the configuration client navigation tree. Double-click the file folder icon to expand the view. Select **System Logs**. Double-click the file folder icon to expand the view. Select **Pager Setup**.

Command Customization

When you select **Pager Setup** you can select a carrier and modem to use and write a pager message.

Command Customization Settings

When you select **Pager Setup** from the navigation tree you get a **Pager Setup** dialog box with Command Customization Settings similar to the dialog box shown in Figure 26.

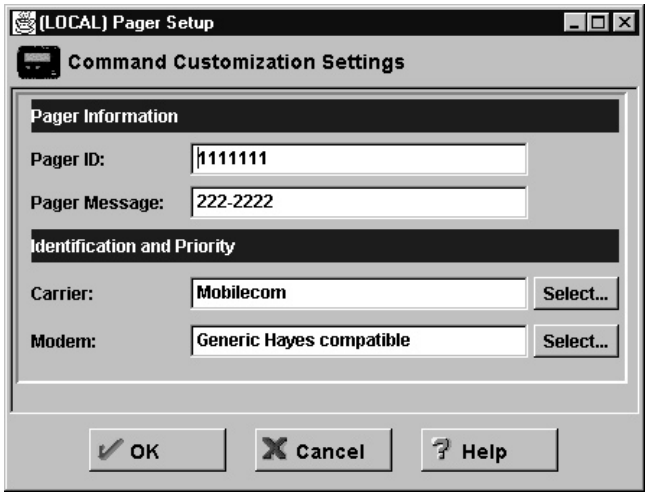

Figure 26. Pager Setup

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

1. Enter the pager ID. This is usually a unique PIN assigned to your pager by your carrier company.
2. Enter the pager message. This is a string containing the default message the user wants to send. For numeric pagers, this must be a number only. For alphanumeric pagers, this can be a text message. Do not exceed the maximum message length specified in your carrier setup or your message might be truncated. Do not use a colon (:). If you do, it will be replaced by a blank space character.
3. If there is no carrier name, click **Select** to define a carrier. You will get the **Pager Carrier Administration** dialog box. See “Carrier Administration” for details on how to fill in this panel.
4. If there is no modem name, click **Select** to define the modem. You will get the **Pager Modem Administration** dialog box. See “Modem Administration” on page 99 for details on how to fill in this panel.
5. Click **OK**.

Change Command Customization

When you select Pager Setup from the navigation tree you get the **Pager Setup** dialog box with Command Customization Settings.

1. Type or select values in the entry fields to modify the values of the existing customization entry fields.
2. Click **OK**.

Delete Command Customization

1. You can delete an entry on the **Pager Carrier Administration** dialog box or the **Pager Modem Administration** dialog box by selecting an item from the list and double-clicking **Delete**.

   You will be asked to confirm the deletion.
2. Click **Yes** to confirm the deletion or **No** to return to the **Pager Setup** dialog box.

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

Carrier Administration

From the **Pager Setup** dialog box, go to the carrier name field and click **Select**. You get a **Pager Carrier Administration** dialog box similar to the one shown in Figure 27 on page 98.

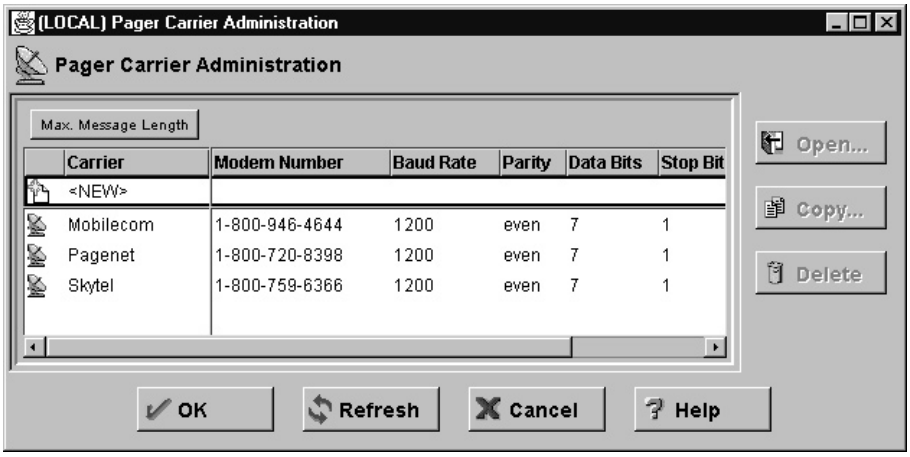

Figure 27. Pager Carrier Administration

Add a Carrier

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

1. Enter the carrier name. This can be anything as long as it is unique and provides enough information for you to recognize which carrier it is.
2. Enter the carrier phone number, which is the phone number for a modem at the carrier company, as opposed to their voice paging or other service number. It must be the right modem number for regional or national coverage and for an numeric or alpha pager, as required by the paging device and the service you have contracted.
3. Enter TAP for paging method; the only value allowed.
4. Enter the password if the carrier allows or requires one.
5. Enter the maximum message length for an alphanumeric pager and the maximum digits for a numeric pager.
6. Enter the baud rate. Click the arrow and choose a value from the list.
7. Click **Even**, **Odd**, or **None** for the parity field.
8. Choose the default data bits; click either **7** or **8**.
9. Choose the default stop bits; click either **1** or **2**.
10. Click **OK**.

Change Carrier

1. Select the carrier you want to change from the **Pager Carrier Administration** dialog box and click **Open**.
2. Refer to “Add a Carrier” for an explanation of the fields you can change. The carrier name itself cannot be changed. This field will be disabled.
3. Make your desired changes.
4. Click **OK**.

Delete Carrier

1. Select the carrier you want to delete from the **Pager Carrier Administration** dialog box and click **Delete**.
2. You will be asked to confirm the deletion. Click **Yes** to confirm.

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

Modem Administration

Your modem manual will contain relevant information about how to initialize your modem. You might need to coordinate modem settings with your carrier. In general, only Hayes-compatible modems, that use the standard modem commands are supported.

From the **Pager Setup** dialog box, go to the modem name field and click **Select**. You get a **Pager Modem Administration** dialog box similar to the one shown in Figure 28.

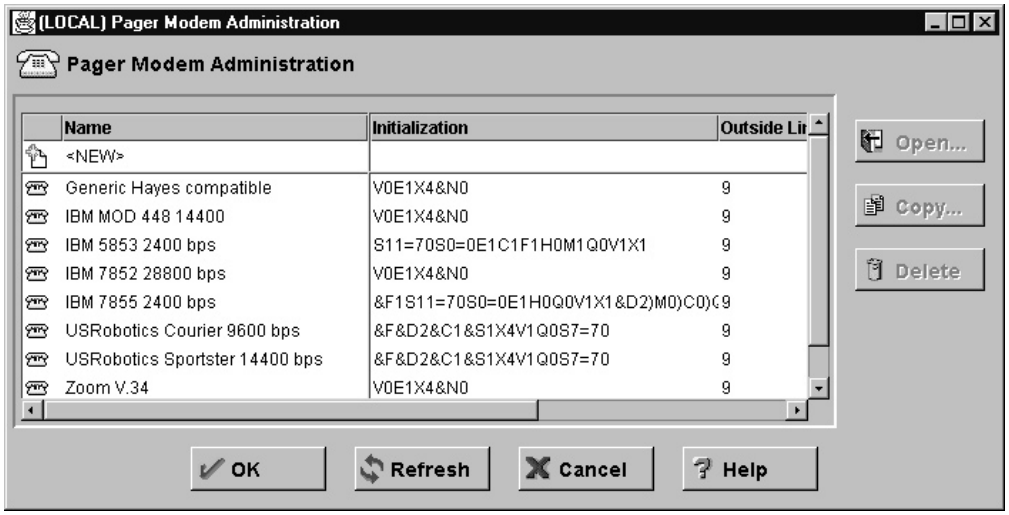

Figure 28. Pager Modem Administration

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

Add a Modem

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

1. Enter the modem name. This can be anything as long as it is unique among the other definitions and provides enough information for you to recognize which modem it is.
2. Enter the COM Port number, which defines the serial COM Port to which the modem is attached. Enter a number less than 10. While the modem must be hardware-configured to this port, it must not be defined to Windows NT; doing so will cause the pager functions to be denied access to the port. If the modem does not match the hardware settings, the pager code will retry for a long time and eventually fail.
3. Enter the initialization string, which should define the modem as a data modem with an echo on X level4 and a fixed baud rate defined by the local site. Do not include the AT command. The pager function will put it at the beginning of the initialization string.

4. Enter the outside line prefix. This is the number you dial to get outside of your company.
5. Click **OK**.

Change Modem

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

   On the **Change Modem** dialog box you will see a list of fields you can change for the modem definition. Refer to “Add a Modem” on page 99 for explanations of these fields.

2. Click **OK**.

Delete Modem

1. Select a modem name from the **Pager Modem Administration** dialog box and click **Delete** to delete a modem definition file.
2. You will be asked to confirm the deletion. Click **Yes** to confirm.

Pager Notification Logging

The pager notification process uses the firewall log utility to write output logs. All pager messages and errors are written to the general firewall syslog facility. For more information on how to set up and use your firewall log files, see “Chapter 15. Managing Log and Archive Files” on page 103.

Testing Pager Setup

You can verify your pager setup by using the `pager` command. See the *IBM eNetwork Firewall Reference* for details. It is strongly recommended that you use the pager command any time you define or change the setup to be sure your system, modem, carrier, and paging devices all talk with each other correctly and that pages can actually be sent and received.

Execute Commands

You can specify a program that is invoked each time an alert theshold is reached. To specify a program:

1. Click **Log Monitor Administration** and then double-click **NEW**.

   The **Add Log Monitor** dialog box appears.

2. In the **Class Type** drop-down box, select **Execute Comand**. This enables the **Command Filename** field of the panel.
3. In the **Command Filename** field, enter the fully-qualified pathname of the program you want to invoke when an alert threshold is reached.

   The working directory for the commands executed by the log monitor is `\winnt\system32`. Because the command shell is launched from a system process, only system environment variables are set. User environment variables are not set. In general, the launched program should use fully qualified file names instead of relying on path variables.

   The Firewall will pass the full Alert message as the first parameter of the program as follows:

```
Total Authentication Failure Alerts: ICA0001e
Per User Auhentication Failure Alerts: ICA0002e
Per Host Auhentication Failure Alerts: ICA0003e
Message Threshold Alerts: ICA0004e
```

See the *IBM eNetwork Firewall Reference* for a complete description of these messages.

Chapter 15. Managing Log and Archive Files

This chapter describes how to use the log facilities through the configuration client. As users try to access hosts through the various IBM Firewall servers, the IBM Firewall writes entries in the log file maintained by the `IBM Firewall logging service`.

The IBM Firewall can generate large volumes of logging information depending on how you configure your firewall. Log entries can come from a variety of places such as socks and expert filters. Additionally, log files can be written to at a variety of severity levels; for example, *debug, information*, or *error*. This chapter also tells you how to use the log management and log archive management facilities to manage the size of your log and archive files.

Log File Creation 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 `firewall log` facility. Only the primary firewall administrator has access to the `firewall log` facility. Alert messages go to the alert log facility. Administrative audit log information goes to the audit log facility.

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

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

The following priority levels are cumulative with *debug* capturing the most information. *Critical* captures only the most severe firewall events.

- Debug
- Information
- Warning
- Error
- Critical

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

The priority levels do not correspond precisely to the message tag suffix *(i,e,w,s..)*. You might need to experiment to determine how to *shut off* certain messages.

Add Log Facilities

From the configuration client navigation tree, double-click the System Administration file folder icon to expand the view. Double-click the System Logs file folder icon to expand the view. Select Log Facilities. The **Log Facilities** dialog box appears displaying the set of log facilities currently enabled.

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

   The **Add Log Facilities** dialog box appears, as shown in Figure 29.

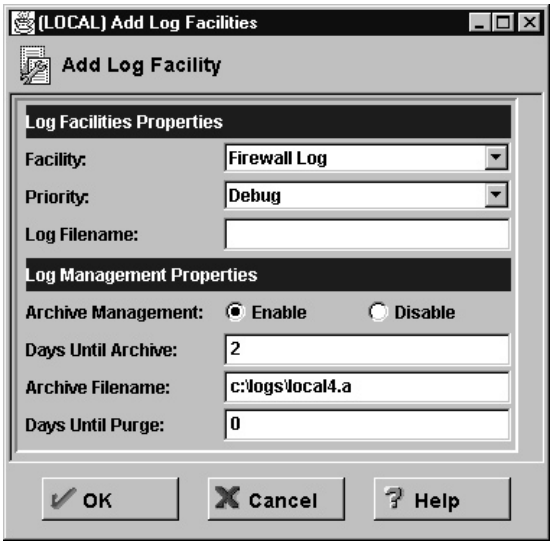

Figure 29. Add Log Facilities

2. Click the **Type** arrow to select type. Type is Filename.
3. The log facility determines the type and source of information that gets logged. Click the **Facility** arrow to select one of the following log facilities:
   - Firewall log - general firewall logs, including filter logging
   - Alert log - log monitor daemon status and threshold violation warnings used to populate the Alerts Display
   - Mail log
4. Click the **Priority** arrow to choose the priority. The logging priorities are listed in order of increasing severity. The priority you select will be the minimum level to be logged.
5. Fill in the log filename. The log filename must have an absolute path (beginning with a drive and a backslash \) and the path to the file must exist.
6. Archive management can be used with a *filename* type log facility only. When enabled, the log file size can be reduced on a periodic basis. Enabling archive management means that you set parameters upon which the `fwlogmgmt` command depends. See “Archiving Logs” on page 105. You can either enable or disable archive management parameters.
7. Select the number of full days until record(s) in an active log should be archived. The value must be zero or greater. Archival will occur when an `fwlogmgmt -l` command finds active log records that qualify under this criteria. Log management does not include the current day when calculating the number of days to keep a log record.
8. Enter an archive filename and full path. The IBM Firewall provides a default archiving function, which uses a directory. However, you can use plug-in archive functions, if you desire.
9. Select the number of full days until an archived log file should be deleted from the archive. The value must be zero or greater. Purge will occur when an

`fwlogmgmt -a` command finds archived file(s) that qualify under this criteria. Log management does not include current day when calculating the number of days to keep an archived file.

10. Click **OK**.

Change Log Facilities

1. Select the firewall logging entry you want to change from the **Log Facilities** dialog box and click **Open**.

   The **Change Log Facilities** dialog box will appear.
2. Change the desired fields. See “Add Log Facilities” on page 103 for an explanation of the fields.
3. Click **OK**.

Delete Log Facilities

1. Select a firewall logging entry from those currently enabled on the **Log Facilities** dialog box and click **Delete**.

   The **Delete Warning** panel appears.
2. Click **OK** if you want to continue with the delete. Click **Cancel** if you change your mind. This does not delete the actual log file.

Archiving Logs

The archival process:

- Removes qualifying records from an active log
- Places them in a separate file
- Compacts the resulting file
- Places the new file into an archive directory

To start a log management program to archive accumulated log records, you have two options:

1. Run the `fwlogmgmt -l` command from the command line from time to time, or
2. Set up the `fwlogmgmt -l` command as an NT Scheduled Service.

Purging the log archives consists of deleting qualifying archived files from the archive directory.

To purge the archived files you have two options:

1. Run the `fwlogmgmt -a` command from the command line from time to time, or
2. Set up the `fwlogmgmt -a` command as an NT Scheduled Service.

Qualifying records and files are determined by the values specified in the log facilities definitions, as described in “Add Log Facilities” on page 103.

The most efficient or convenient means of running the log management process is to set it up as an NT Scheduled Service. Start it by using the Services object on the control panel.

For example, if you want to set up the log management archiving process to run at 3:00 AM every day, type

```
at 3:00 /every:M,T,W,Th,F,S,Su fwlogmgmt -l
```

Plug-in DLL

See the *IBM eNetwork Firewall Reference* for information on the log archiver plug-in DLL that you can use to replace the Firewall default DLL.

Log Management Outputs

The log management facility does some preliminary integrity checks before proceeding with any log management activities. If any problems are found, diagnostics are sent to the firewall log facility when you run the `fwlogmgmt` command from the command line.

Mail or admin audit (local0) log facilities are subject to different archival rules than other facilities. All log facilities require that archival be enabled in order to be archived. However, firewall (local4) and alert (local1) log records are only archived if their dates exceed criteria specified in the facilities definition at the time the archive process is run; whereas the *entire* mail or audit log file will be archived each time. Also, the information in the mail log is considered to be for debug purposes and there is generally little value in archiving it. Other, more generally useful mail information is logged in the firewall (local4) log.

Report Utilities

You can use the report utility functions to assist you in generating reports from current or archived log files. Report utilities generate tabulated files of administrative information that are organized and formatted for easy mapping to relational database tables. These tables help the firewall administrator to analyze:

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

Using the utilities and the firewall log, the administrator can create a regular text file of the messages. Additionally, tabulated files can be generated and imported into tables in a relational database system, such as the DB2 family of products. The administrator can then use the Structured Query Language (SQL) to query the data and generate reports.

Report Utilities are installed as part of the Firewall installation. They can also be separately installed and run on a non-firewall host. The configuration client can be used to run them on a firewall. On a non-firewall machine, use the command line.

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

Do not try to use report utilities on any log files prior to the IBM Firewall for AIX V3R1. You can however, use report utilities to process log files from the IBM Firewall for AIX V3R1 or later. You can also use them to process the AIX su log. See the *IBM eNetwork Firewall Reference* for more detailed information on report utilities.

Running Report Utilities Using the Configuration Client

From the configuration client navigation tree, double-click the System Administration file folder icon to expand the view. Double-click the System Logs file folder icon to expand the view. Select **Report Utilities**. The **Report Utilities** dialog box appears, as shown in Figure 30.

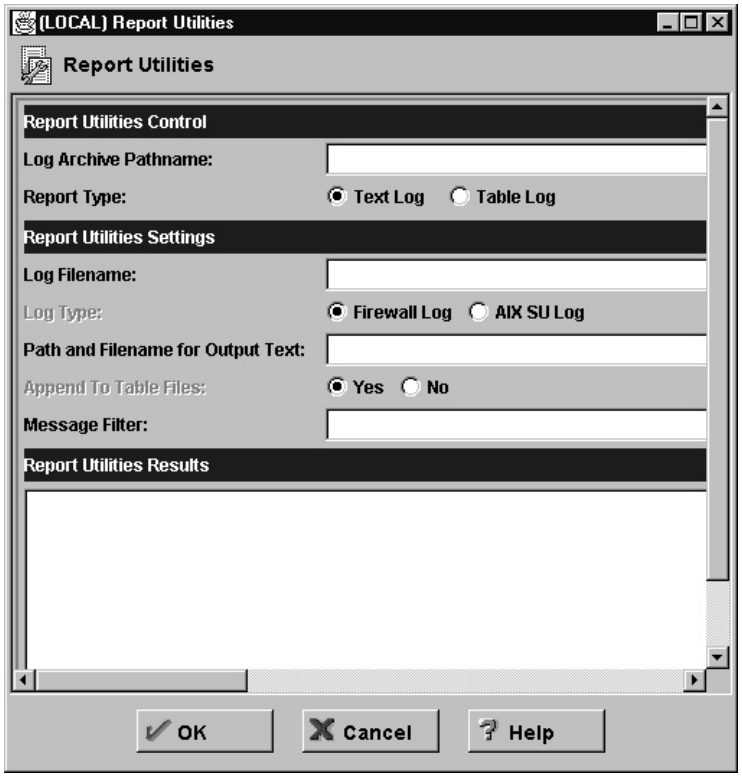

Figure 30. Report Utilities

1. For the default archiver provided with the IBM Firewall, the log archive pathname is the directory that contains compressed log files. In the log archive pathname field, enter the directory that you specified in the archive directory field on the **Log Facilities** dialog box. Enter the absolute path name to the archive directory. If you want to view a log file that is not archived, leave this field blank.
2. Select the **Report Type**. To produce the expanded log message text, select **Text Log**. To create tabulated files for DB2 usage, select **Table Log**. If you import the resulting files into DB2, you can perform SQL queries on the log data. Refer to the *IBM eNetwork Firewall Reference* for more information.
3. The log filename is any one of the compressed archived log files or other valid `firewall log`, or the name of an AIX su log file. If you made an entry in the log archive directory field, you can click the **Log Filename** arrow to choose which log to work with. If you do not enter a log archive in step 1, the log file name you enter here must be the name of a valid, uncompressed firewall log file or an AIX su log file. You must specify a full path.
4. Select the **log type**, either **firewall** or **AIX su**.
5. Enter the **Path and Filename for Output Text**.
6. Select **Yes** to append the results of a table log request to existing tabulated files or **No** to replace the existing files.

7. This field allows you to select certain types of messages to be placed in the output text file. The contents of this field are treated as parameters that are placed into a standard Windows NT `Find` command. For example, if you type ″ICA0″ into the field (you must include the quotes), it is as though you are running the following command:

   ```
   fwlogtxt < my.log | find "ICA0"
   ```

   Here are some example entries that you can place into this field and the results:

   ```
   FILTER                          RESULT
   "ICA0"          Lists log monitor threshold alert messages
   "ICA3"          Lists Socks-related messages (#ICA3000 - 3999)
   "ICA2010"       Only lists occurrences of the ICA2010 message
   /V "ICA3"       Lists all messages except Socks messages
   /C "ICA001"     Counts the number of ICA0001 messages
   ```

8. Clicking **OK** produces the requested file(s) in the specified output directory on the firewall machine.
9. The Report Utilities Results area shows any error message from the report utility that was run. To view the log text resulting from a Text Log report type, click **Log Viewer** on the main Firewall configuration client panel, and enter the fully- qualified output file name. The .tbl files resulting from a Table Log report type can be loaded into a database as described in the *IBM eNetwork Firewall Reference.*

Chapter 16. Translating Network Addresses

With the explosive growth of the Internet, IP address depletion problem has become significant. Network address translation (NAT) provides a solution to the IP address depletion problem based upon address reuse.

Addresses within a private network can be assigned out of a very large address space (typically a 10.0.0.0 class A address space). These addresses are private and are not exposed to the Internet. Therefore these addresses can be reused by another IP network. A single registered IP address is used to hide many private network addresses. NAT converts the unregistered addresses and port numbers into a valid registered Internet address and port numbers. In the inbound direction NAT converts the registered Internet address and port numbers back to the unregistered addresses and port numbers. The advantage of NAT is that it transparently allows a network that uses private or illegal addresses to communicate with hosts on the Internet, effectively allowing the private network to have a large address space. Furthermore, by using NAT, addresses in the private network are hidden from the external world providing an additional level of security.

Figure 31 illustrates basic NAT operation in an IBM Firewall environment.

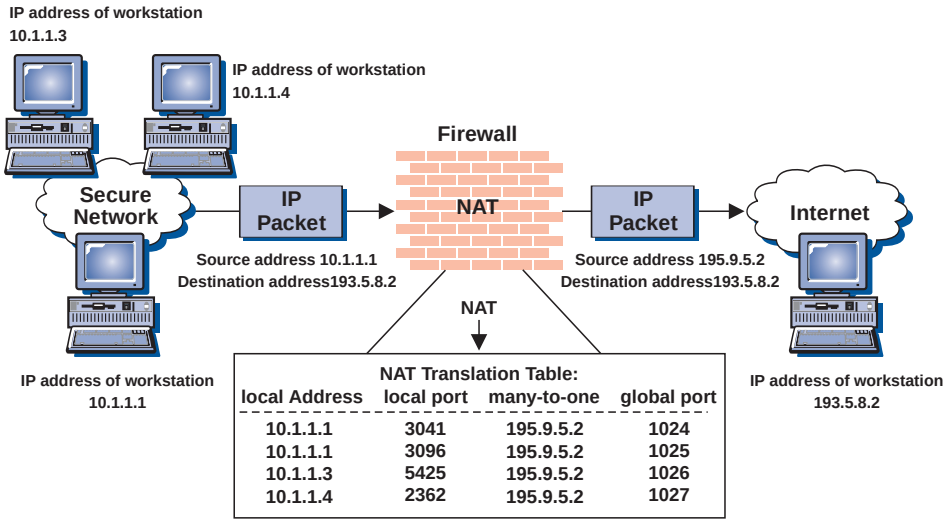

Figure 31. Network Address Translation

TCP/UDP packets generated by secure hosts have their source address replaced with a registered Internet address. The secure host's port will be translated to a unique port number. All outbound packets will have the same source address but a unique port number. This form of translation, referred to as many-to-one translation, allows many secure hosts to hide behind a single address. The packet's checksums in the IP header and TCP/UDP pseudo header are updated. The maximum number of concurrent connections is limited to 64536, (actually 64512) because ports 0–1023 are reserved. Inbound connections are supported with static (rather than dynamic) translation table entries. For example, host 193.5.8.2 can initiate a TCP connection with host 10.1.1.1 (using the global address 195.9.5.2) only if a static entry exists in the NAT translation table that maps 195.9.5.2 to 10.1.1.1.

All packets generated by TCP/UDP applications can be translated. Difficulties arise if the application data contained in the IP packet contains an IP address. One

particularly troublesome application for address translation is FTP. The FTP control connection issues ″PORT″ commands or ″PASV″ replies that contain an ascii encoded IP address in the message. In this case NAT must modify not only the addresses in the IP header but also the ascii address and the port number in the payload.

The IBM eNetwork Firewall NAT Implementation

The IBM Firewall NAT implementation supports basic address translation as described above with the following caveats:

- TCP/UDP applications (except for FTP as described below) that contain IP address information in the payload will only have the packet header fields translated as described above. This implies that UDP applications such as DNS or SNMP will not have address information contained in the payload translated.
- FTP PORT commands are fully translated. However, the address embedded in a PASV response packet is not translated.
- NAT does not detect a TCP disconnect but rather relies on a configurable idle time-out before removing a dynamic translation table entry and inserting the registered IP address back in the pool of available addresses.

Example Interaction Between NAT, Filters, and Tunnels

Figure 32 illustrates an example interaction between NAT, filters, and tunnels.

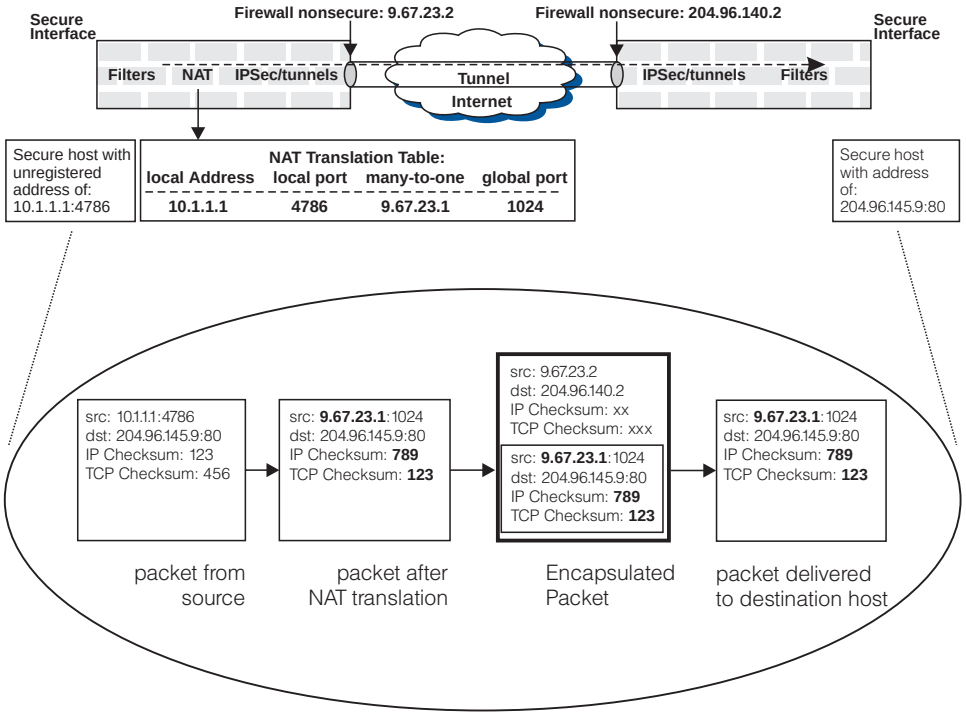

Figure 32. Example Interaction between NAT, Filters and Tunnels

Assume an IPSec ESP tunnel is manually established between firewalls 9.67.23.2 and 204.96.140.2. NAT is active only at the 9.67.23.2 firewall because this secure network uses private addresses. The secure network at the other end of the tunnel

is not using NAT. In addition to illustrating basic NAT translation (the bold fields in the second packet from the left illustrate the fields in the packet that are modified during outbound address translation), Figure 32 on page 110 also illustrates that the translated packet from the host is encapsulated in an IP packet that is *not* translated.

In general, filtering is applied to outbound packets prior to NAT and to inbound packets after NAT translation. Therefore the filter rules are based on untranslated addresses. When NAT and tunnels are involved, the filter rules at the firewall that has NAT active are also based on untranslated addresses. At the other end of the tunnel (assuming NAT is not active at this firewall), the filter rules for inbound packets are based on translated source and destination addresses (for the inbound and outbound cases respectively). If NAT is active at both ends of the tunnel the discussion above applies in both directions.

Using the scenario illustrated in Figure 32 on page 110 as an example, and assuming that the goal is to allow secure host 10.1.1.1 to communicate with secure host 204.96.145.9 over a tunnel, the firewall attached to 10.1.1.1 must have a filter rule permitting 10.1.1.1 to communicate with 204.96.145.9 over a tunnel. At the other firewall connected to the destination host, a filter rule is required permitting communication between 9.67.23.1 and 204.96.145.9 through the tunnel.

More About NAT

Use NAT if you want to allow:

- Direct access for a machine behind a firewall to a nonsecure site while protecting the address of that secure machine.
- A number of machines without registered addresses to share a registered address so that they can reach sites on the Internet.
- Others from unsecure locations access to a server behind a firewall.

Obtain the registered addresses for NAT from your ISP. All addresses used for NAT cannot be used for any other purpose.

There are four options for NAT:

Many-to-One
: Involves translating a packet's secure address and port number so that many (up to 65536) internal addresses can share one registered IP address. This one shared registered IP address will hide local addresses but in addition to it, you need another registered Internet address, for the Firewall's nonsecure address. The NAT configuration will identify the registered Internet address that is used for port translation using a many-to-one entry.

Translate
: Used to create a list of secure addresses that will be translated.

Exclude
: Used to create a list of secure addresses that will not be translated.

Map Used to reserve a specific registered address for a specific secure address.

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 33.

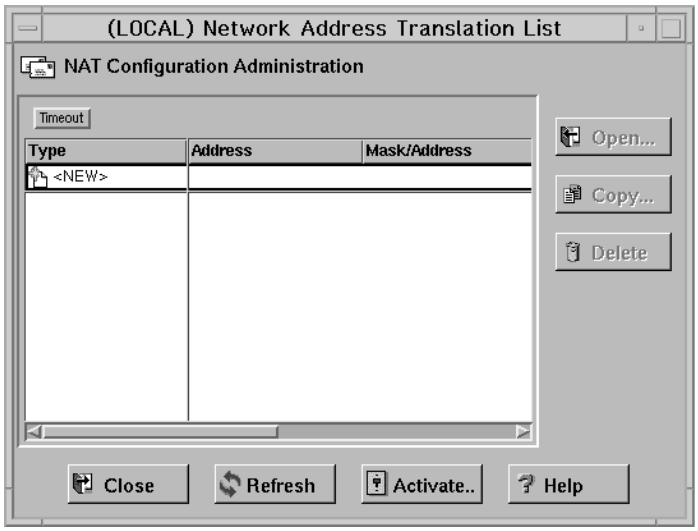

Figure 33. Network Address Translation List

3. Network Address Translation entries contained in the NAT configuration file are displayed on this dialog box. 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 box appears.
2. From the **Add NAT** dialog box, click the arrow in the Type of NAT field and select from the following:
   - Many-to-One Registered Network Address: Adds the IP addresses specified to the reserved IP address. The parameters are reserved IP address and time-out accociated with the translation table.
   - 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.

Many-to-One Registered Network Address

A many-to-one registered address entry translates a packet's secure address and port number so that many (up to 65536) internal addresses can share one

registered IP address. Thus you can hide many local addresses with the one registed IP address. (You will need one additional registered Internet address for the firewall's nonsecure address).

When a secure host sends packets to a nonsecure network, one registered IP address is allocated. This unique registered IP address is used to transport an IP frame between the IBM Firewall and machines outside of the secure network.

If you selected Many-to-One from the Add NAT screen, enter the following values:

Registered IP Address
Obtain this from your ISP. This will be a dotted-decimal IP address that all secure adresses will hide behind.

Choose a network object by clicking **Select** to get the **Select Network Object** dialog box. Select a network object and click **OK**. The network object is added to the Network Object field on the **Add NAT Configuration** dialog box. Or, type a value directly into the field if you have not previously created a network object.

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. The range of values is 5 through 45.

Translate Secured Network Address

A translate secured IP address entry defines a set of secure network addresses that require NAT to perform IP address translation. By default, NAT performs address translation on all secure IP addresses in the translate secured IP address set.

If you selected 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** dialog box. Select a network object and click **OK**. The network object is added to the Network Object field on the **Add NAT Configuration** dialog box. Or, type a value directly into the field if you have not previously created a network object.

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

An exclude secure IP address entry defines a set of secure network addresses that does not require NAT to perform IP address translation. By default, NAT performs address translation on all secure IP addresses in the translate secured IP address set.

If you selected Exclude from the **Add NAT** dialog 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** dialog box. Select a network object and click **OK**. The network object is added to the Network Object field on the **Add NAT Configuration** dialog box. Or, type a value directly into the field if you have not previously created a network object.

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

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.

If you selected Map from the **Add NAT Configuration** dialog box, enter the following values:

Secured IP Address
A dotted-decimal IP address that should be translated into a specified registered IP address.

Choose a network object by clicking **Select** to get the **Select Network Object** dialog box. Select a network object and click **OK**. The network object is added to the Network Object field on the **Add NAT Configuration** dialog box. Or, type a value directly into the field if you have not previously created a network object.

Registered IP Address field
A dotted-decimal IP address into which a specified secure IP address should be translated.

You can choose a network object by clicking **Select** to get the **Select Network Object** dialog box. Select a network object and click **OK**. The network object is added to the Network Object field on the **Add NAT Configuration** dialog box.

Change NAT Entry

Select an existing NAT entry from the **NAT Configuration** dialog box 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** dialog box and click **Delete** to remove a Network Translation entry from the NAT configuration file.

   A confirmation dialog box appears.
2. Select Yes or No.

NAT Activation

1. From the configuration client navigation tree, double-click the 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 dialog box similar to the one shown in Figure 34 appears.

Figure 34. 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

NAT will log a variety of error conditions provided that both NAT logging and filter logging are enabled. NAT logging is enabled through the **NAT Activation** panel or by using the **fwnat** command. Filter logging is enabled through the **Log Facility** panel or by using the **fwlog** command.

The following activities will be logged to the firewall log facility:

- Updates to the NAT tables from the Administrator (for example, static or MAP entries)
- Dynamic updates to the NAT translation table
- Error messages
- Failed translation attempts that result in the packet being discarded
- Each time NAT is activated and deactivated

Create Filter Rules for NAT

After you have completed the NAT configuration, you have to create the filter rules for the connections that are going to use NAT. Review “Chapter 8. Controlling Traffic Through the Firewall” on page 43 and use the predefined services that are for direct connections. Examples of predefined services that are for direct connections are:

- HTTP direct out
- Telnet direct out

See “Building Connections Using Predefined Services” on page 44 for more information.

If you want a service to come directly into your network, you will have to create one. See “Using the Configuration Client to Create Services” on page 64 for information on how to do this.

Appendix. 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:

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.

This product includes software developed by the University of California, Berkeley and its contributors.

Trademarks

The following terms are trademarks of the IBM corporation in the United States or other countries or both:

- Common User Access
- DB2
- eNetwork
- IBM

- OS/2

Microsoft, Windows, Windows NT, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

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

For additional information about security on the Internet, visit the IBM eNetwork Firewall home page at **http://www.software.ibm.com/enetwork/firewall**.

Information in IBM Publications

Other IBM sources of information on firewalls, Internet security, and general security topics are listed here.

Firewall Topics

The following documents are available on the IBM Firewall CD-ROM and the IBM eNetwork Firewall home page.

- *IBM eNetwork Firewall User's Guide*, GC31-8658
- *IBM eNetwork Firewall Reference*, SC31-8659
- *Guarding the Gates Using the IBM eNetwork Firewall for NT 3.2*, SG24-5209

Internet and World Wide Web Topics

- *A Guide to the Internet Connection Servers*, SG24-4805
- *Accessing CICS Business Applications from the World Wide Web*, SG24-4547
- *Accessing OS/390 OpenEdition MVS from the Internet*, SG24-4721
- *Accessing the Internet*, SG24-2597
- *Building the Infrastructure for the Internet*, SG24-4824
- *Cool Title about the AS/400 and Internet*, SG24-4815
- *The Domino Defense: Security in Lotus Notes and the Internet*, SG24-4848
- *Examples of Using MQSeries on WWW*, SG24-4882
- *How to Secure the Internet Connection Server for MVS/ESA*, SG24-4803
- *Lotus Domino Server Release 4.5 on AIX Systems: Installation, Customization, and Administration*, SG24-4694
- *Netscape Proxy Server*, SK2T-7444
- *Running CICS Transactions through the Web: The CICS Internet Gateway to VSE/ESA*, SG24-4799
- *Safe Surfing: How to Build a Secure World Wide Web Connection*, SG24-4564
- *Teach Yourself CGI Programming with PERL in a Week*, SR23-7343
- *Using the Information Super Highway*, GG24-2499
- *World Wide Web Access to DB2*, SG24-4716

General Security Topics

- *The Basics of IP Network Design*, SG24-2580
- *Elements of Security: AIX V4.1*, GG24-4433
- *Enterprise-Wide Security Architecture and Solutions Presentation Guide*, SG24-4579
- *HACMP/6000 Customization Examples*, SG24-4498
- *IBM Global Network (IGN) Security Policy*, GC34-2206
- *IBM Security Architecture: Securing the Open Client/Server Distributed Enterprise*, SC24-8135
- *IBM Systems Monitor: Anatomy of a Smart Agent*, SG24-4398
- *Security Overview of Open Systems Networking*, GG24-3815
- *Systems Monitor for AIX User's Guide*, SC31-8173
- *TCP/IP Tutorial and Technical Overview*, GG24-3376

Information in Industry Publications

These industry publications pertain to TCP/IP and UNIX:

- Albitz, Paul, and Cricket Liu. *DNS and BIND*. Sebastopol, CA: O'Reilly and Associates, 1997. (ISBN: 1-56592-236-0)
- Costales, Brian with Eric Allman. *Sendmail*. O'Reilly and Associates, Inc. (ISBN: 1-56592-222-0)
- Hunt, Craig. *TCP/IP Network Administration*. O'Reilly and Associates, Inc. (ISBN: 0-937175-82-X)

- Nemeth, Snyder, et al. *UNIX System Administration Handbook*. Prentice Hall. (ISBN: 0-13-151051-7

These industry publications pertain to firewalls and security on the Internet:

- Ahuja, Vijay. *Network and Internet Security*. Boston: Academic Press Professional, 1996. (ISBN: 0120455951)
- Ahuja, Vijay. *Secure Commerce on the Internet*. Boston: Academic Press Professional, 1997. (ISBN: 0120455978)
- Anderson, Bart, et al. *The Waite Group's UNIX Communications and the Internet*. Indianapolis, IN: Sams Pub., 1995. (ISBN: 0672305372)
- Atkins, Derek, et al. *Internet Security: Professional Reference*. Indianapolis, IN: New Riders Publishing, 1996. (ISBN: 1562055577)
- Chapman, D. Brent, and Elizabeth D. Zwicky. *Building Internet Firewalls*. Sebastopol, CA: O'Reilly and Associates, 1995. (ISBN: 1565921240)
- Cheswick, Willam R., and Steven M. Bellovin. *Firewalls and Internet Security*. New York: Addison-Wesley, 1994. (ISBN: 0201633574)
- Cooper, Frederic J., et al. *Implementing Internet Security*. Indianapolis, IN: New Riders Publishing, 1995. (ISBN: 1562054716)
- Curry, David. *UNIX System Security: Guide for Users and Systems Administrators*. Sebastopol, CA: O'Reilly and Associates, 1994. (ISBN: 0201563274)
- Garfinkel, Simson, and Gene Spafford. *Practical UNIX Security*. Sebastopol, CA: O'Reilly and Associates, 1991. (ISBN: 0937175722)
- Garfinkel, Simson, and Gene Spafford. *Practical UNIX and Internet Security*. Sebastopol, CA: O'Reilly and Associates, 1996. (ISBN: 1565921488)
- Hare, Chris, and Karanjit Siyan. *Internet Firewalls and Network Security*. Indianapolis, IN: New Riders Publishing, 1996. (ISBN: 1562056328)
- Randall, Neil. *Teach Yourself the Internet in a Week*. Indianapolis, IN: Sams.Net, 1995. (ISBN: 0672307359)
- Stallings, William. *Internet Security Handbook*. Foster City, CA: IDG Books, 1995. (ISBN: 0077092546)
- Stevens, W. Richard. *TCP/IP Illustrated*. Reading, MA: Addison-Wesley, 1994. (ISBN: 0201634953)

Glossary

You can access the IBM Software glossary at:
http://www.networking.ibm.com/nsg/nsgmain.htm.

Index

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

IBM eNetwork Firewall for Windows NT
User's Guide
Version 3 Release 2.1.1

Publication No. GC31-8658-01

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
GC31-8658-01

IBM

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

Cut or Fold
Along Line

Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

GC31-8658-01