A Secure Way to Protect Your Network: IBM SecureWay Firewall for AIX V4.1

Jorge Ferrari, Cristiane Ferreira, Paul Gunther, Tae Beom Lee

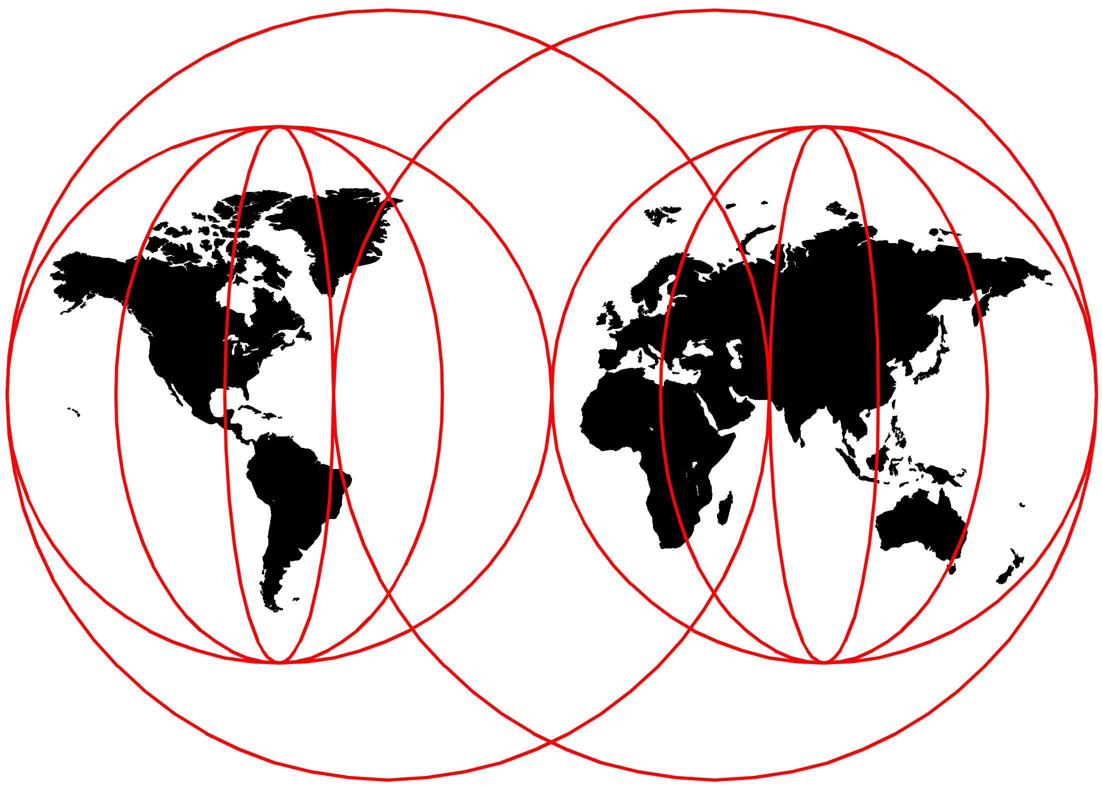

International Technical Support Organization

www.redbooks.ibm.com

SG24-5855-00

International Technical Support Organization

SG24-5855-00

A Secure Way to Protect Your Network: IBM SecureWay Firewall for AIX V4.1

November 1999

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix A, “Special notices” on page 321.

First Edition (November 1999)

This edition applies to the IBM SecureWay Firewall V4.1 for AIX, Program Number 5697-F48 for use with the AIX Operating System.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. HZ8 Building 678
P.O. Box 12195
Research Triangle Park, NC 27709-2195

Contents

Preface

In this redbook we describe the implementation of the IBM SecureWay Firewall V4.1 for AIX for enforcing security at the boundaries of TCP/IP networks. It contains details as well as basic and advanced configuration techniques. It also gives several helpful tips on AIX security and general Internet security.

The chief feature of this redbook is its numerous examples. These examples cover all the major capabilities of the IBM firewall: IP filters; proxy and SOCKS services; logging, monitoring, and alerts; DNS; mail; and remote configuration. It goes a step further by covering configuration examples of the systems around the firewall.

This redbook will help you to plan, install, implement, and manage the IBM SecureWay Firewall V4.1 for AIX. This redbook is aimed at users who have some experience with firewalls. If you are not one of those, we suggest that you read the introductory chapters of *Guarding the Gates Using the IBM eNetwork Firewall V3.3 for Windows NT,* SG24-5209.

The team that wrote this redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization, Raleigh Center.

Jorge Ferrari is a Network Security Specialist in the International Technical Support Organization, Raleigh Center. Before working in the security area, he was a specialist in network design and capacity planning. He holds a degree in Electronic Engineering from the Universidad de Buenos Aires, Buenos Aires, Argentina.

Cristiane Maria Ferreira is a Network Specialist in the RS/6000 Support and Services group at IBM Brazil. She has eight years of experience in UNIX platforms, and she has been working with AIX support for the past four and a half years. Her areas of expertise include TCP/IP networking, firewalls, and networking security.

Paul Gunther is an Advisory RS/6000 Technical Specialist in IBM Australia. He has 13 years of experience in the UNIX field. He holds a degree in Computer Science from the Queensland University of Technology. His areas of expertise include AIX operating system kernel, C development on UNIX, graphics, TCP/IP network security and firewalls. He has written extensively on problem determination techniques, Web enabled applications and network security.

Tae Beom Lee is the president of JOIN Consulting, an IBM Korea business partner in the security and networking area. His specialty is security, auditing and network design. He holds a BS degree in Electronic Engineering from Seoul National University and an MS degree in Computer Science from Sokang University. He also holds Professional Engineer and CISA qualifications.

Thanks to the following people for their invaluable contributions to this project:

Martin Murhammer
International Technical Support Organization, Raleigh Center

Gordon Arnold, Wilfred Jamison, Vach Kompella, William Pranger, Bill Serencsics, Susanne Vergara
IBM Firewall Development

Laurie Bader
IBM WebSphere Development

Comments welcome

Your comments are important to us!

We want our redbooks to be as helpful as possible. Please send us your comments about this or other redbooks in one of the following ways:

- Fax the evaluation form found in "IBM Redbook evaluation" on page 335 to the fax number shown on the form.
- Use the online evaluation form found at `http://www.redbooks.ibm.com/`.
- Send your comments in an Internet note to `redbook@us.ibm.com`.

Chapter 1. New features of Firewall V4.1

IBM SecureWay Firewall V4.1 for AIX introduces many new features, so it can be a key building block to the security services sold through the combination of products that cater to e-business. To enable the products for future growth, IBM SecureWay Firewall V4.1 for AIX takes two primary tactical steps:

1. Reusing and leveraging quality security code that exists around the corporation or may be acquired through contract negotiation. This step is necessary to move more quickly in the offering of features on the Firewall.
2. Creating a *baseline* product from which to launch new initiatives. This step is necessary underpinning for a cogent security solution.

1.1 Firewall V4.1 extensions

The IBM SecureWay Firewall V4.1 for AIX offers numerous extensions:

- Enhanced IPSec and virtual private network (VPN) support
- Multi-processor support
- Filter enhancements
- Secure Mail Proxy enhancements
- Network Address Translation (NAT) enhancements
- An enhanced HTTP Proxy using IBM Web Traffic Express technology
- SOCKS protocol Version 5 support
- Setup Wizard
- Network Security Auditor (NSA) enhancements
- Enhanced logging
- National language support for German speakers

1.1.1 Enhanced IPSec and virtual private network (VPN) support

The IBM SecureWay Firewall V4.1 for AIX includes enhanced IPSec support including triple-DES encryption and support for the new VPN headers. It also supports interoperability with several IBM servers and routers as well as non-IBM VPN devices that support the new VPN headers.

The new VPN headers and encryption details are described in the following RFCs:

- RFC 2401: Security architecture for the Internet Protocol

  This RFC specifies the base architecture for IPSec-compliant systems. The goal of this architecture is to provide various security services for traffic at the IP layer, in both the IP V4 and IP V6 environments.

- RFC 2402: IP Authentication Header

  This RFC specifies the IP Authentication Header (AH) which provides connectionless integrity and data origin authentication for IP datagrams and also provides protection against replays.

- RFC 2403: The Use of HMAC-MD5-96 within ESP and AH

This RFC specifies the use of MD5 (Message Digest 5) combined with HMAC (Hashed Message Authentication Code) as a keyed authentication mechanism within the context of the ESP (Encapsulating Security Payload) and the Authentication Header. The goal of HMAC-MD5 is to ensure that the packet is authentic and cannot be modified in transit.

- RFC 2404: The Use of HMAC-SHA-1-96 within ESP and AH

  This RFC specifies the use of the HMAC algorithm in conjunction with the SHA-1 (Secure Hash Algorithm -1) algorithm as an authentication mechanism within the revised IPSec Encapsulating Security Payload and the revised IPSec Authentication Header. HMAC with SHA-1 provides data origin authentication and integrity protection.

- RFC 2405: The ESP DES-CBC Cipher Algorithm

  This RFC describes the use of the DES cypher algorithm in Cypher Block Chaining (CBC) mode as a confidentiality mechanism within the context of the IPSec Encapsulating Security Payload. DES is a symmetric block cypher algorithm.

- RFC 2406: IP Encapsulating Security Payload (ESP)

  The Encapsulating Security Payload (ESP) header is designed to provide a mix of security services in IPv4 and IPv6. ESP may be applied alone, in combination with the IP Authentication Header, or in a nested fashion (tunnel mode).

Also, new dynamic filters for IPSec are enhanced to support the following:

- Automatic filter rule generation
- Easy to configure tunnels
- Fixed filter rule definitions

For more information on IPSec and VPN refer to Chapter 10, “Virtual Private Network” on page 241.

1.1.2 Multi-processor (MP) support

The IBM SecureWay Firewall V4.1 for AIX can exploit the multi-processor features of the RS/6000 for scaling and performance improvements. Users can benefit from the symmetric multi-processor (SMP) capability.

1.1.3 Filter enhancements

Filters have been enhanced to provide better performance and more flexibility with configurations.

You can tune the performance of your firewall by choosing where to locate different types of filter rules. Filter position can be defined within the connection category hierarchy as follows:

- **Upper Layer**: Has the priority over other layers.
- **Dynamic Filter Rules**: Dynamic filter rules are implicitly activated when a VPN tunnel that needs them is activated.
- **Real Audio Layer**: It’s for Real Audio filters.
- **Lower Layer**: Has lowest priority.

Once you save a connection you cannot change the position, but you can reorder a connection within its own position type. The upper layer has highest priority.

In addition, a frequency indicator provides the number of times a connection is used, so you can use this count to adjust filter positions.

For more information refer to Chapter 4, “Packet filters” on page 43.

1.1.4 Secure Mail Proxy

IBM Firewall Secure Mail Proxy is an application level proxy that acts as a middleman between the originating mail server and receiver, allowing only selected operations through to the destination Mail Transport Agent (MTA).

Secure Mail Proxy relays incoming mail to the appropriate internal mail server(s) and also relays outgoing mail to appropriate Internet mail server(s).

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

- Anti-spam algorithms, including message blocking from known spamers (an exclusion list), verification checks on the validity of messages (known ways of blocking undesirable messages), configurable limits on the number of recipients per mail messages, configurable limits on the maximum size of a message.
- Anti-spoofing support, including integration with strong authentication mechanisms.
- Overflow servers can be used in the case of mail delivery failure or when other overflow situations occur.

For more information refer to Chapter 8, “Secure Mail Proxy” on page 167.

1.1.5 SOCKS protocol Version 5

AIX Firewall V4.1 supplied function conforms to SOCKs V5 specification (RFC 1928/1929). SOCKS protocol Version 5 offers the following advantages:

- Easy deployment of authentication and encryption methods.

  SOCKS V5 clients can be authenticated by any of the supported firewall authentication schemes. However, clients must support the Challenge Response Authentication Method (CRAM) protocol in order to use strong authentication; otherwise, they can only use the user name/password scheme.
- UDP association, which creates a virtual proxy circuit for traversing UDP-based proxy circuits.
- Ability to chain SOCKS servers.
- SOCKS5 Monitor, which displays real-time SOCKS performance information.
- Supports Archie, Finger, FTP, Gopher, HTTP, Proxy, News, SNMP, Telnet, TFTP, RealAudio, Whois, X-Windows and others.
- Supports both SOCKS V4 and SOCKS V5 clients and allows migration of SOCKS via three modes:

  Permissive

  Permits any user and does not enable outbound authentication. Inbound connections are denied.

Migration

Allows SOCKS V4 users to pass unauthenticated but requires SOCKS V5 clients to authenticate. Inbound SOCKS V5 connections are allowed, while inbound SOCKS V4 connections are denied.

Strict

Requires all users to be authenticated, thus requires SOCKS V5.

For more information refer to Chapter 7, “SOCKS server” on page 137.

1.1.6 Network address translation

Network address translation (NAT) has been enhanced to support many-to-one address mappings. By using port numbers to create unique mappings, this feature allows the firewall to translate many secure addresses into one public address. IP address mapping is done by using the same registered IP address with different port numbers.

Support for MAP, TRANSLATE, and EXCLUDE functions are still provided and the IBM Firewall Version 3.x for AIX NAT configuration can be used. The last reserve is used as a many-to-one address.

Also, logs are improved to include the following items:

- IP addr: <port> allocated
- IP addr: <port> released
- Translation errors

NAT now also supports ICMP packets, which is necessary to support Path MTU discovery.

For more information refer to Chapter 9, “Network Address Translation” on page 221.

1.1.7 Enhanced HTTP proxy

The IBM Firewall V4.1 provides a full-featured HTTP proxy implementation based upon the IBM Web Traffic Express (WTE) product as *non-caching* mode. The HTTP proxy efficiently handles browser requests through the IBM Firewall eliminating the need for a SOCKS server for Web browsing. Users can access useful information on the Internet, without compromising the security of their internal networks and without altering their client environment to implement the HTTP proxy.

This proxy supports HTTP version 1.1. For more information refer to Chapter 6, “Proxy” on page 87.

1.1.8 Setup Wizard

A wizard has been provided to aid the user with the initial configuration of the IBM Firewall. This setup Wizard enables a user, who does not have extensive knowledge of the IBM Firewall, to have a basic firewall configuration up and running quickly after installation of the IBM Firewall.

The setup Wizard guides users through the following fundamental tasks:

- Basic security policies
- System administration tasks having to do with interfaces, DNS, mail, and log setup
- Setup to allow secure users to access non-secure networks through the Web, Telnet, or FTP
- Creating an alert log
- Setting up some basic log monitor thresholds

For more information refer to Chapter 3, “Basic configuration” on page 19.

1.1.9 Network Security Auditor

The Network Security Auditor (NSA) is a tool that checks your network servers and firewall for security holes or configuration errors.

By periodically running the Network Security Auditor, you can ensure that nothing has been changed in a way that creates a security vulnerability, especially after you put the firewall online.

For more information refer to Chapter 12, “Network Security Auditor” on page 315.

1.1.10 Enhanced logging

The IBM SecureWay Firewall V4.1 for AIX monitors the messages sent to the AIX syslog for potential crisis situations, based upon user-defined thresholds. In the event of a threshold violation, fwlogmond delivers an alert, in a manner specified by the firewall administrator. The firewall log facility and alert log facility are subsets of the AIX syslog.

Log facilities are enhanced to keep separate logs for various functions for easier management. Additional logs for mail, WTE, and SOCKS are available.

For more information refer to Chapter 11, “Logging, monitoring, and reporting” on page 285.

1.1.11 National language support for German

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

1.2 How to get more information

More information is available from the following sites:

- Web reference

  `www.software.ibm.com/security/firewall/`

- IBM Firewall fixes

  `www.software.ibm.com/security/firewall/support/`

- NSA

  `http://dr.watson.ibm.com/nsa/`

- Aventail - SOCKS V5

  `www.aventail.com`
- SDI

  `www.SECURID.com`

Chapter 2. Installation

The procedure to install IBM SecureWay Firewall V4.1 for AIX is to:

1. Install and configure the AIX operating system.
2. Install and configure IBM SecureWay Firewall V4.1 for AIX.

When installing and configuring the firewall please keep the following points in mind:

- Keep It Simple and Secure (KISS).
- Physically secure your system in a locked area.
- Make a checklist of changes you make so you can periodically check to make sure those settings are still the same.
- Consider how secure the network structure is between your firewall and the secure network.
- Run the minimum number of services (KISS).
- User IDs should be kept to a minimum and set up using the firewall.
- Remove any compilers, assemblers or any other computer language that allows system calls.
- Use the audit and logging functions to monitor the system.

2.1 Requirements

The hardware requirements include the following:

- RISC System/6000
- 1 GB hard drive
- 64 MB memory (minimum)
- At least two network interfaces
  - One secure and one non-secure connection
  - Statically assigned IP addresses for all interfaces
- Supported interfaces:
  - Token-Ring
  - Ethernet
  - Local Area Network for AIX
  - X.25
  - ATM
  - FDDI
  - SLIP

The software requirements include the following:

- AIX V4.3.2 or later (earlier OS versions not supported)

The Secure Mail Proxy requirements include the following:

- Internal mail server
- Overflow server (recommended)

The DNS requirements include the following:

- Internal domain name server
- External domain name server (recommended)

The pager alert support requirements include the following:

- IBM or Hayes compatible modem
- A supported pager
- The service provider must support the Tele-AlphaNumeric Protocol (TAP).

If you decide to authenticate users using the Security Dynamics SecurID card, the requirements are:

- Model SD200 (standard card without buttons)
- PINPAD (card with buttons)

If you plan to use SNMP for monitoring your firewall, it requires:

- System View Agent for AIX SNMP Mapper

2.2 Install AIX

AIX is a multiuser, multipurpose operating system. It offers a wide variety of services that are not needed when installing a firewall. We recommend that you install AIX from scratch for IBM SecureWay Firewall V4.1 for AIX and aim to install a minimal system. A fresh install of AIX will ensure you do not have any other software installed except the minimum you will need to run firewall.

Take the option to install the Trusted Computing Base, as it can only be installed during AIX installation. This will give you the option of later using it to add a further level of security to your system.

2.2.1 Additional filesets

For the convenience of installation process we installed the additional filesets as shown in Table 1.

Use `lslpp -l fileset.name` to see if each fileset is installed. Use `smit install_latest` to install them.

Use the preview option in smit before every install and patch. It will help you identify problems before you actually begin modifying files.

Table 1. Additional filesets

AIX filesets	Description
bos.net.tcp.server	TCP/IP server
bos.net.ate	Asynchronous Terminal Emulator
bos.net.uucp	UNIX to UNIX Copy Program
bos.acct	Accounting service

AIX filesets	Description
bos.sysmgt.trace	Software trace service aids
bos.dosutil	DOS utilities
bos.perf.diag_tool	Performance diagnostic tool
bos.perf.pmr	Performance PMR data
Java.rte	Java runtime
bos.terminfo.ansi.data	Support for ANSI terminals (if you plan to use Telnet from a Windows machine).

2.2.2 Post-AIX installation

Once AIX was installed, we performed a number of tasks before installing the firewall.

2.2.2.1 Install AIX fixes

It is important to apply the latest fixes/updates to AIX to avoid known problems and potential vulnerabilities as publicly advertised by CERT (Computer Emergency Response Team).

> **Note**: The CERT Coordination Center is the organization that grew from the computer emergency response team formed by the Defense Advanced Research Projects Agency (DARPA) in November 1988. The CERT charter is to work with the Internet community to facilitate its response to computer security events involving Internet hosts, to take proactive steps to raise the community's awareness of computer security issues, and to conduct research targeted at improving the security of existing systems.

Apply these *now* because once the firewall code is installed, no further AIX updates can be applied as this may de-harden the firewall or overwrite firewall components.

The procedure to update AIX once the firewall is already installed is to de-install all firewall software, update AIX, then reinstall the firewall software.

At a *minimum* apply fixes from the following sources:

1. Supplied AIX update CDROM
2. Latest AIX maintenance level

Specific fixes can be obtained from:

1. Call the IBM AIX Support Center quoting fix number and the media of your choice.
2. AIX Support Web site:

   `http://service.software.ibm.com/cgi-bin/support/rs6000.support/downloads`

3. Service FTP site:

   `ftp://service.software.ibm.com/aix/fixes/v4/other`

Store fixes in the directory /usr/sys/inst.images on the firewall for installation. If media need to be copied to the hard disk, use these commands:

```
cd /usr/sys/inst.images
rm .toc
smit bffcreate
```

Use the following command to apply fixes:

```
smit update_all
```

Ensure the default update flags are modified as follows:

```
Commit: NO
Save Replaced Files: YES
```

The following command will view all filesets that are installed with version and update information:

```
lslpp -h | pg
```

Advisories

The CERT and IBM ERS advisories should be reviewed for any relevant security fixes that may be required:

```
http://www.cert.org/advisories/index.html
http://www.ers.ibm.com/tech-info/advisories/sva/index.html
```

Note: ERS (Emergency Response Service) is designed to assist enterprises in establishing the right level of security and reliability to support e-business expansion.

2.2.2.2 AIX customization

After AIX V4.3.2 installation we did some further configuration, as shown in Table 2.

Table 2. AIX customization

Task/Command	Explanation
`lscfg | pg`	Review installed hardware, adapters and disks against your order.
`smitty chtz`	Set the time zone.
`smit date`	Set the date and time.
`/usr/sbin/bootinfo -r`	Review available memory in KB.
`lsps -a`	Review currently configured paging space. Hint: it should be usually double the amount of physical RAM (in AIX V4.3.2, this setting is automatic).
`chfs -asize=+16384 /`	Increase / filesystem by 8 MB.
`chfs -asize=+163840 /var`	Increase /var filesystem by 80 MB.
`chfs -asize=+163840 /tmp`	Increase /tmp filesystem by 80 MB.
`chfs -asize=+163840 /home`	Increase /home filesystem by 80 MB.
`sysdumpdev -e`	Review estimated dump space requirements.
`mklv -t sysdump -y pdumplv rootvg 5`	Create a dump device of size 5 partitions.

Task/Command	Explanation
`sysdumpdev -P -p /dev/pdumplv`	Assign dump device.
`sysdumpdev -P -s /dev/hd6`	Assign backup dump device.
`sysdumpdev -d /var/adm/ras`	Assign directory to copy dump.
`sysdumpdev -K`	Allow dumps for hung systems.
`crfs -v jfs -a bf=true -g rootvg -m /var/log -Ayes -asize=4000000` `(Use 2000000 for 2GB hdisk)`	Create large (2 GB) file-enabled filesystem for logs.
`crfs -v jfs -g rootvg -m /usr/local -Ayes -asize=163840` `(81920 for 2GB disk)`	Create filesystem for tools of size 80 MB.
`crfs -v jfs -g rootvg -m /usr/sys/inst.images -Ayes -asize=327680`	Create 80 MB filesystem/installation directory for PTFS and updates.
`mount all`	Mount all filesystems. Ignore errors about filesystems already mounted.
`/usr/lib/errdemon -s4194304 -B32768`	Increase error log buffers and log file size to provide a larger audit trail.
`chdev -lsys0 -amaxuproc=130000`	Increase maximum processes per user.
`chdev -lsys0 -aautorestart=true`	Autorestart after system crash.
`chdev -lsys0 -afullcore=true`	Enable full-size core dump.

`vi /etc/security/login.cfg`
`Add the line`
`herald = "Access Restricted\n\rlogin:"`
`to the 'default:' stanza in file`
	Change default login prompt to reflect business access policy and remove clues to potential hackers about the operating system the firewall is running.

`vi /etc/security/login.cfg`
`Change the sak_enable stanza in like this:`
`sak_enable = true`
	Enable secure attention key (Ctrl+X/Ctrl+R) on login ports to prevent Trojan Horse attack.

`vi /etc/security/user`
`Locate the 'default:' stanza and do the following 2 changes:`
`Add : registry = files`
`Change : umask = 077`
	Enforce use of local authentication file registry held in /etc/security without the possibility of using remote authentication registries, such as NIS or DCE.
	Change the default file creation mask to prevent file sharing among firewall users.

`vi /etc/motd`	Change the message-of-the-day file to reflect business access policy.
`chuser fsize=4194303 root`	Increase maximum file size that can be written.
`chuser core=4194303 root`	Increase maximum core file size that can be written.
`mkdir /var/log/workspace`	Create temporary directory for processing the logs, using the /var/log file system.
`mkdir /var/log/archive`	Create directory to store archived logs, using the /var/log file system.

Task/Command	Explanation

`vi /etc/netsvc.conf`
`Add a line`
`hosts=local,bind4`
	Control DNS lookup order.
	(By default DNS is not enabled on the firewall; however, if it is and the system is rebooted while the DNS server is unavailable, the system will appear to take a long time to IPL while stuck on LED value 581).

`If using xinit/aixterm`
`vi ~/.Xdefaults`
`Add a line`
`Aixterm * login Shell:<tab>`
`true`
	Makes aixterm behave as a login shell, so it runs the profile files (/etc/profile and $HOME/profile) every time a new aixterm window is opened.

2.2.3 Users

During the installation of the firewall, all users other than root, daemon, bin, adm and nobody will be removed. The root account will be disabled for remote logins.

For all user IDs in the system that are not used for regular logins, you should define a mail alias that transfers the mail to a local administrator. Otherwise, mail could pile up accidentally in a mailbox without anyone ever noticing it.

For performance and tuning informations refer to *AIX Version 3.2 & V4 Performance Monitoring and Tuning Guide,* SC23-2365, which is included on the AIX Version 4.3 Base Documentation CD.

2.3 Firewall installation

Installing IBM SecureWay Firewall V4.1 for AIX is a straightforward AIX installp function.

After the installp process is complete it is necessary to reboot the system. The installp messages tell you to do this, but it is easy to overlook them.

2.3.1 Install the code

Before you begin the installation of the firewall software, move the smit.log and smit.script files aside. The firewall installation will log all its activity to these files. You will be able to use these files to list the changes made to the system by the install process.

Start the installation program at the AIX command line with `smit install_latest`. You need to define the installation media (CD-ROM or local disk) and select the components that you need to install.

The contents of the smit.log file for selected filesets are as follows:

```
Filesets listed in this section passed pre-installation verification and will
be installed.

Selected Filesets
  -----------------
  FW.base 4.1.0.0                              # Base IBM SecureWay Firewall
  FW.cfgcli 4.1.0.0                           # IBM SecureWay Firewall Remot...
  FW.http 4.1.0.0                             # IBM SecureWay Firewall HTTP ...
  FW.ipsec 4.1.0.0                            # IBM SecureWay Firewall IPSEC...
```

```
FW.libraries 4.1.0.0                            # IBM SecureWay Firewall Commo...
FW.report 4.1.0.0                               # IBM SecureWay Firewall Repor...
internet_server.msg.en_US.httpd 2.0.0.0         # IBM Firewall Web Proxy Messa...

Requisites
----------
(being installed automatically;  required by filesets listed above)
internet_server.base.httpd 2.0.0.0              # IBM Firewall Web Proxy
internet_server.proxy.exe 2.0.0.0               # IBM Firewall Web Proxy Files
sva.mapper 1.4.2.0                              # SystemView Agent for AIX SNM...
sva.rte 1.4.2.0                                 # SystemView Agent for AIX

<< End of Success Section >>
```

After a successful preview we changed the preview option to *no* and started the installation.

2.3.2 Firewall hardening

A major part of the IBM SecureWay Firewall V4.1 for AIX installation is a process called hardening. System resources that might be used to compromise security are disabled to secure the system. The firewall hardening process does the following:

- All unnecessary programs are removed from inittab.
- Startup entries are disabled from /etc/rc.tcpip except the following lines:

  ```
  start /usr/sbin/syslogd "$src_running"
  start /usr/sbin/named "$src_running" "-b /etc/fwnamed.boot" #FW#
  start /usr/sbin/inetd "$src_running"
  ```
- Startup entries are added to /etc/rc.tcpip for firewall components.
- All unnecessary functions are disabled from /etc/inetd.conf, and the following lines are added:

  ```
  ftp stream tcp nowait root /usr/sbin/pftpd pftpd -ns
  telnet stream tcp nowait root /usr/sbin/ptelnetd ptelnetd
  ibmfwrcs stream tcp nowait root /usr/sbin/ibmfwrcs ibmfwrcs
  ```

  Note that the AIX telnet and ftp servers are commented out, and these new proxy servers are added in their place.
- If the AIX Common Desktop Environment (CDE) is installed on the system, the installation process disables it.
- All logins for nonessential users are disabled.
- The owners are of unowned files and directories set to root.
- Previous firewall users are migrated to this new version.
- Insecure applications are disabled.
- The file system integrity checker database is generated.

We add logging contents from smit.log for your reference. We included firewall configuration filesets that may be referenced during installation and operation.

```
Installing Firewall configuration files:
        cp -p /usr/lpp/FW/config/fwtdefn.conf /etc/security
        cp -p /usr/lpp/FW/config/fwsocks.cfg /etc/security
        cp -p /usr/lpp/FW/config/secag.cfg /etc/security
```

```
        cp -p /usr/lpp/FW/config/fwsecuremail.cfg /etc/security
        cp -p /usr/lpp/FW/config/fwservices.cfg /etc/security
        cp -p /usr/lpp/FW/config/fwrules.cfg /etc/security
        cp -p /usr/lpp/FW/config/fwobjects.cfg /etc/security
        cp -p /usr/lpp/FW/config/fwaudio.cfg /etc/security
        cp -p /usr/lpp/FW/config/logmgmt.cfg /etc/security
        cp -p /usr/lpp/FW/config/fwpriv.users /etc/security
        cp -p /usr/lpp/FW/config/fwcust.pager /etc/security
        cp -p /usr/lpp/FW/config/fw.carriers /etc/security
        cp -p /usr/lpp/FW/config/fwmodem.config /etc/security
        cp -p /usr/lpp/FW/config/hayes.modem /etc/security
        cp -p /usr/lpp/FW/config/rcsfile.cfg /etc/security
        cp -p /usr/lpp/FW/config/fwtpproxy.cfg /etc/security
        cp -p /usr/lpp/FW/config/usrsportster.modem /etc/security
        cp -p /usr/lpp/FW/config/usrcourier.modem /etc/security
ln -fs /usr/lpp/FW/lib/fwusrdb.a /usr/lib/fwusrdb.a
ln -fs /usr/lpp/FW/lib/gwauth4.a /usr/lib/gwauth4.a
ln -fs /usr/lpp/FW/sbin/ptelnetd /usr/sbin/ptelnetd
ln -fs /usr/lpp/FW/sbin/pftpd /usr/sbin/pftpd
cp -ph /usr/lpp/FW/socks5/config/s5.conf /etc/security/s5.conf
ln -fs /usr/lpp/FW/socks5/config/explode.cfg /etc/security/explode.cfg
ln -fs /usr/lpp/FW/socks5/config/socks5.header.cfg
/etc/security/socks5.header.cfg
ln -fs /usr/sbin/fwMonitor /etc/security/socks/bin/fwMonitor
ln -fs /usr/sbin/fwS5convert /etc/security/socks/bin/fwS5convert
ln -fs /usr/sbin/fwSocks5 /etc/security/socks/bin/fwSocks5
ln -fs /usr/lib/libavconfig.a /etc/security/socks/lib/libavconfig.a
ln -fs /usr/lib/libavconfig.la /etc/security/socks/lib/libavconfig.la
ln -fs /usr/lib/libs5mon.a /etc/security/socks/lib/libs5mon.a
ln -fs /usr/lib/libs5mon.la /etc/security/socks/lib/libs5mon.la
ln -fs /usr/lpp/FW/socks5/man/socks5.1 /etc/security/socks/man/socks5.1
ln -fs /usr/lpp/FW/socks5/man/libsocks5.conf.5
/etc/security/socks/man/libsocks5.conf.5
ln -fs /usr/lpp/FW/socks5/man/s5.conf.5 /etc/security/socks/man/s5.conf.5
ln -fs /usr/lpp/FW/socks5/man/socks5.conf.5
fwmail.cocp -p /usr/lpp/FW/config/fwfschk.db.list /etc/security
        cp -p /usr/lpp/FW/config/fwl.cfg /etc/security
        cp -p /usr/lpp/FW/config/fwhscnt.cfg /etc/security
        cp -p /usr/lpp/FW/config/fwconfig.map /etc/security
        cp -p /usr/etc/fwmib.defs /etc/fwmib.defs

Adding Firewall data to /etc/security/login.cfg
Creating reserved firewall users
Setting SNMP parameters for firewall
Removing extra SVA entries in snmpd.conf snmpd.peers if needed
Removing rmitab piobel applications from inittab
        rmitab qdaemon
        rmitab writesrv
        rmitab uprintfd
        rmitab cfgmceh
        rmitab sva
Adding Firewall data to /etc/rc.tcpip
Adding Firewall data to /etc/services
Adding Firewall data to /etc/inetd.conf
Adding Firewall filter support to /etc/rc.net

Disabling Common Desktop Environment
```

```
Disabling all logins for non-essential users
         uucp
         guest
         lpd
         nuucp

Setting owner of unowned files and directories to root
         /usr/lpp/FW/config/http.dat

Setting Firewall attributes for root user, disabling remote login
Migrating old Firewall users to new version
         fwdpuser
         tblee
         tiger
         lion
         cris
         paulprox

Disabling unsecure applications:
         chmod 0000 /usr/bin/tftp
         chmod 0000 /usr/bin/utftp
         chmod 0000 /usr/sbin/tftpd
         chmod 0000 /usr/bin/uucp
         chmod 0000 /usr/sbin/uucpd
         chmod 0000 /usr/bin/rcp
         chmod 0000 /usr/bin/rlogin
         chmod 0000 /usr/sbin/rlogind
         chmod 0000 /usr/bin/rsh
         chmod 0000 /usr/sbin/rshd
Generating file system integrity checker database
```

When the firewall installation is completed, you see the following summary:

```
Installation Summary
--------------------
Name                        Level           Part        Event       Result
-------------------------------------------------------------------------------
sva.rte                     1.4.2.0         USR         APPLY       SUCCESS
sva.mapper                  1.4.2.0         USR         APPLY       SUCCESS
internet_server.msg.en_US.h 2.0.0.0         USR         APPLY       SUCCESS
FW.libraries                4.1.0.0         USR         APPLY       SUCCESS
FW.report                   4.1.0.0         USR         APPLY       SUCCESS
FW.ipsec                    4.1.0.0         USR         APPLY       SUCCESS
FW.cfgcli                   4.1.0.0         USR         APPLY       SUCCESS
FW.base                     4.1.0.0         USR         APPLY       SUCCESS
FW.http                     4.1.0.0         USR         APPLY       SUCCESS
internet_server.base.httpd  2.0.0.0         USR         APPLY       SUCCESS
internet_server.base.httpd  2.0.0.0         ROOT        APPLY       SUCCESS
internet_server.proxy.exe   2.0.0.0         USR         APPLY       SUCCESS
internet_server.proxy.exe   2.0.0.0         ROOT        APPLY       SUCCESS

installpSoftware changes processed during this session require this system and
any of its diskless/dataless clients to be rebooted in order for the changes to
be made effective.
```

2.3.3 Post firewall installation

After installing IBM SecureWay Firewall V4.1 for AIX we need to complete a few more steps.

2.3.3.1 Install firewall fixes

To ensure your firewall performs properly it's a good idea to regularly check for and install patches. The best practice is to apply all current available PTFs. For firewall code fix information link to the following site:

`http://www.software.ibm.com/security/firewall/support/fixes/fwaix.html`

You can download PTFs from the AIX Fix Distribution Service, either downloading individual fixes or by using the fixdist tool for multiple downloads. If you have an IBM Support Contract you can also order the fixes on CD-ROM or magnetic tape.

You need to know at least one of the following:

- Authorized Program Analysis Report (APAR) number
- Program Temporary Fix (PTF) number (firewall PTFs start with the prefix URxxxxxx)
- Fileset name

You can also directly download from the FTP site by fileset name:

`ftp://service.software.ibm.com/aix/fixes/v4/other/`

2.3.3.2 IP forwarding

IBM SecureWay Firewall V4.1 for AIX adds some lines to /etc/rc.net that enables log and *ipforwarding*. This and other *no* options will be honored by the firewall unless contradicted by a rule. IBM SecureWay Firewall V4.1 for AIX already turns off IP source routing so you do not need to turn it off by changing this line. The contents added by the firewall installation is as follows:

```
# Start - IBM FW Additions ---------------------------------------- #FW#
{                                                                  #FW#
/usr/lpp/FW/fwext/fwkernel.config                                  #FW#
} >> $LOGFILE 2>&1                                                 #FW#
# End - IBM FW Additions ------------------------------------------ #FW#

# Start - IBM FW Additions ------------------- #FW#
/usr/sbin/no -o ipforwarding=1 >>$LOGFILE 2>&1 #FW#
/usr/sbin/no -o somaxconn=1024 >>$LOGFILE 2>&1 #FW#
# End - IBM FW Additions --------------------- #FW#
```

2.3.3.3 Invoke the configuration client in local mode

In a command window, logged on as root, do the following steps:

1. Run `xinit`. After AIX Window is started, open an aixterm.
2. Run `fwconfig`. The window in Figure 1 is displayed.

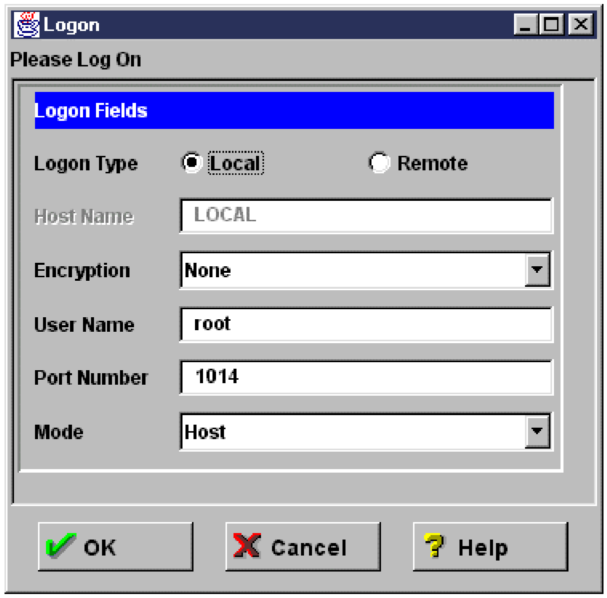

Figure 1. Configuration Client login

3. Select **Local** in the Logon Type field.
4. Enter root for User, select **OK**.
5. Type your root password and click **Submit**.

The Setup Wizard window is automatically opened after you log in for the first time and after that it is available from the help menu. You can do basic configuration settings using it.

For more information refer to Chapter 3, “Basic configuration” on page 19.

2.3.3.4 Enable the Configuration Client in remote mode

The configuration server processes requests from the configuration client. It runs on the firewall machine and it can handle requests from configuration clients that are either on local or remote machines.

The configuration server is initially set up to accept requests only from configuration clients on the local machine. Initial requests are not encrypted. To allow the remote configuration client to connect to the firewall, do the following steps:

1. Check the current parameters of the configuration server by running the command:

```
# fwcfgsrv cmd=list
localonly = yes
encryption = none
sslfile = /etc/security/fwkey.kyr
```

To change these options, use the command fwcfgsrv from the command line. The syntax for this command is:

```
fwcfgsrv cmd=change
                [localonly=yes|no]
                [encryption=none|ssl]
                [sslfile=ssl_file_name]
```

Using `localonly=yes` indicates if the firewall can only be administered from a local machine. Using `localonly=no` indicates that the configuration can occur only on the local machine; this is the default.

The configuration can occur from any machine that is allowed to access that firewall system as a remote configuration client. This permission is issued by creating filters allowing this connection (see Chapter 4, “Packet filters” on page 43).

The field `encryption` indicates whether the configuration server expects incoming data to be encrypted through secure sockets layer (SSL) or not.

Using `encryption=none` means that no encryption will occur; this is the default. Using `encryption=ssl` means that SSL encryption will occur.

If you use SSL encryption, you need to specify the path of the SSL keyfile on the field `sslfile`. The default is /etc/security/fwkey.kyr.

For information on how to create the keyfile, refer to the *IBM SecureWay Firewall Reference Guide for AIX Version 4*, SC31-8418.

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

If a configuration client cannot connect to the firewall machine, and it is installed on a remote machine, use `fwcfgsrv cmd=list` to check that `localonly=no` is set. Also, the language used by the client and the server must match.

If you want to change the `localonly` to `yes`, run the following command:

```
fwcfgsrv cmd=change localonly=yes
```

Chapter 3. Basic configuration

In this chapter we cover two ways of setting up your firewall for the first time.

The first one is using the setup Wizard, which is a new feature shipped with IBM SecureWay Firewall V4.1 for AIX. We recommend this option for new users who are not familiar with the Configuration Client Graphical User Interface (GUI) and who want a simple configuration to begin with.

The second one, which we call "expert configuration", is the classical step-by-step configuration using the Configuration Client GUI. We recommend it for advanced users, who are already familiar with the firewall.

If you start using the Wizard, you can make changes to your configuration using the Configuration Client GUI later on.

3.1 Setup Wizard

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

The Wizard appears automatically after you log onto the firewall for the first time.

Thereafter, the Wizard is available under the help menu item on the Configuration Client GUI. The Wizard is optional; you are not required to use it to configure the firewall.

The Wizard guides you through the following fundamental tasks:

- Basic security policies
- System administration tasks having to do with interfaces, DNS, mail, and log setup
- Setup to allow secure users to access non-secure networks through the Web, Telnet, or FTP
- Creating an alert log
- Setting up some basic log monitor thresholds

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

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

Figure 2 on page 20 shows the Setup Wizard start screen. You can easily follow the setup sequence.

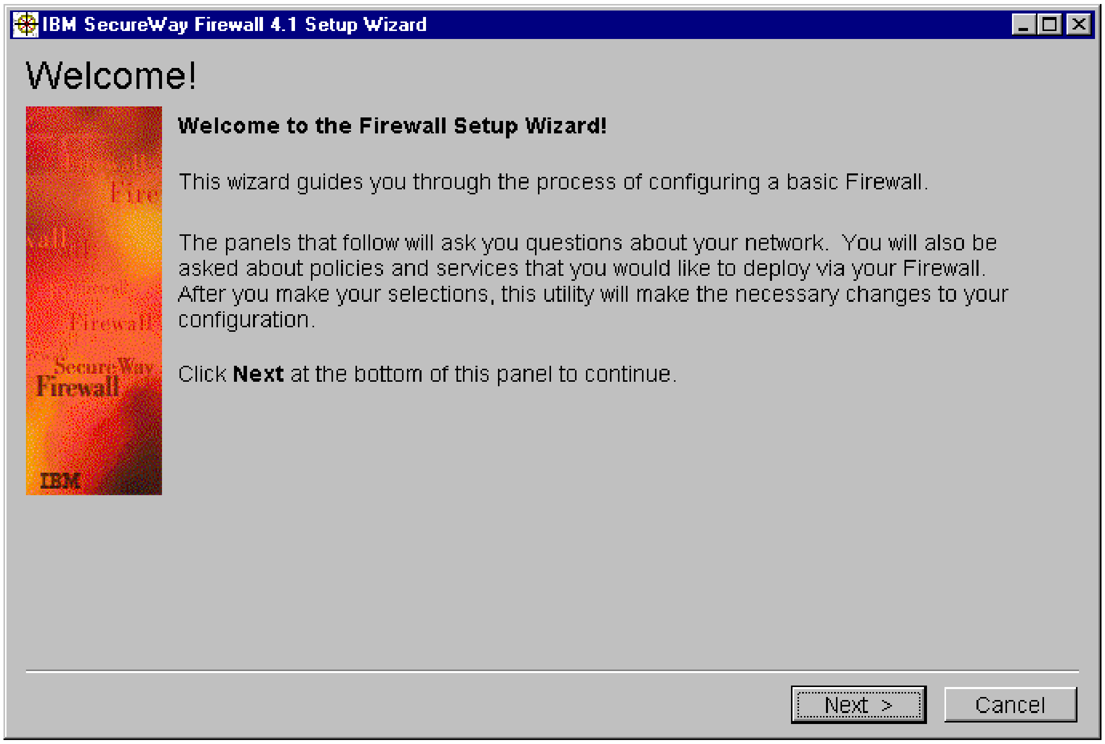

Figure 2. Setup Wizard start screen

3.2 Expert configuration

In this section we cover the basic configuration for users who are using the Configuration Client GUI for an initial setup.

We recommend that you perform all the following steps for your basic configuration right after the installation.

If you are migrating from a previous version of IBM Firewall for AIX, you may follow these steps to check your configuration after the migration. Refer to Chapter 15, “Migration and Backup” on page 381 for more details on migrating from a previous version of IBM Firewall for AIX.

3.2.1 Selecting the secure interface(s)

Log on to the configuration client. The window in Figure 3 is displayed.

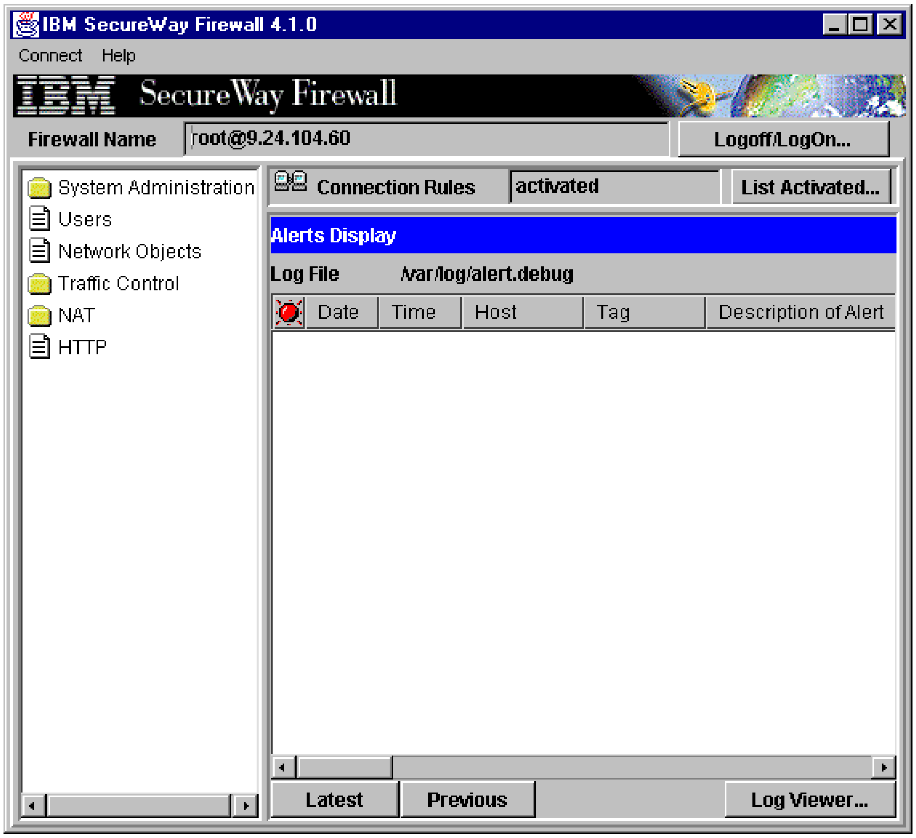

Figure 3. Configuration client main window

Under the System Administration folder, double click **Interfaces**. The window in Figure 4 is displayed.

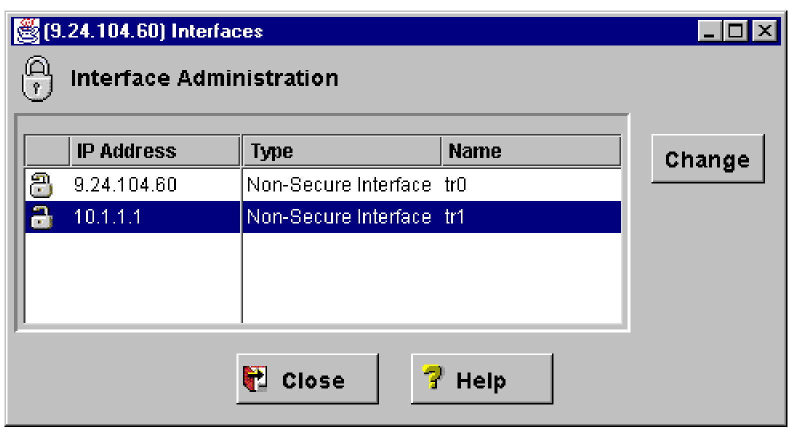

Figure 4. Interfaces window

All TCP/IP interfaces that you previously configured under AIX will be listed as non-secure. The interfaces must be *up* to be shown in the list (you can check the state of the interfaces with the command `ifconfig`).

Select the interface(s) that you defined as secure and click **Change**. The status of the selected interface will be changed to secure.

Repeat these steps until you have changed the status of all your secure interfaces.

3.2.2 Security policy

The purpose of Security Policy is to provide quick and easy configuration of basic filters and to enable or disable Telnet and FTP transparent proxies.

Under the System Administration folder (see Figure 3 on page 21), double-click **Security Policy**. The window shown in Figure 5 is displayed.

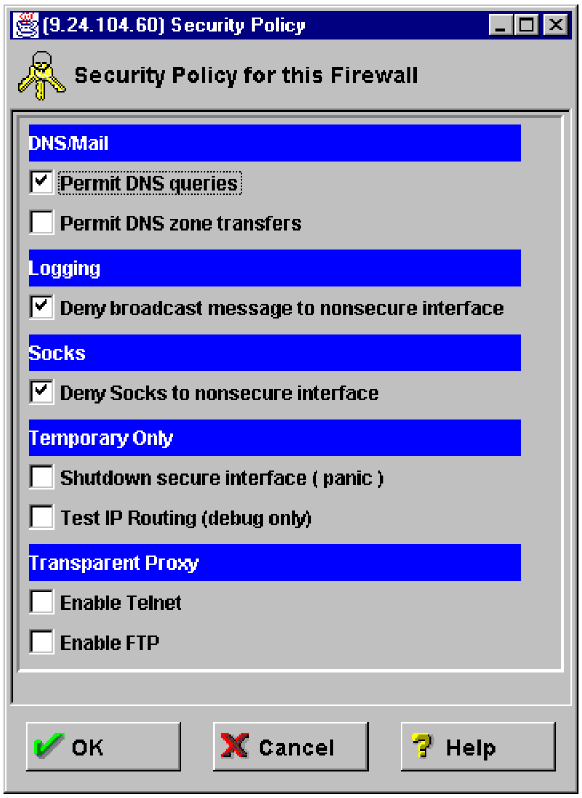

Figure 5. Security Policy window

We checked the following check boxes:

- **Permit DNS queries**: allow all local DNS traffic. It will *not* allow traffic to flow directly from the secure to the non-secure network, and from the non-secure to the secure network (routed traffic). The packets are not logged.
- **Deny broadcast message to non-secure interface**: deny any UDP packet sent to any broadcast address[1] from the non-secure network. The packets are not logged (to keep the log files clean).
- **Deny SOCKS to non-secure interface**: deny any packet sent to port 1080 (SOCKS server) on the non-secure interface. The packets are logged (for audit trail).

These are the basic selections for most environments. If you need to allow DNS zone transfers between your primary and secondary DNS servers, you need to

[1] Broadcast address: any destination IP address where the last octet is equal to 255. In security policy, the filter is created using IP address 0.0.0.255 and mask 0.0.0.255.

choose **Permit DNS zone transfers** also. Remember, you should allow only the traffic that is necessary on your network. If you do not need it, deny it.

The check boxes Shutdown secure interface and Test IP Routing should be used only in certain circumstances. The first will deny all traffic to and from the secure interface. This should be used when you want to interrupt all traffic between secure and non-secure networks. The latter should be used to allow all connections (local and routed) through the firewall. We recommend that you use this selection only to debug problems.

The check boxes Enable Telnet and Enable FTP are used to enable or disable the respective transparent proxies. Refer to Chapter 6, “Proxy” on page 87 for more information on Telnet and FTP transparent proxies.

After you make your selections, press **OK**. In previous versions of IBM Firewall for AIX, a window would be displayed to activate the filters after making changes to the security policy. In IBM SecureWay Firewall V4.1 for AIX, if the connections are already active in your firewall, then the filters for these security policies will be updated and activated immediately. If the connections are not active, you still have to activate them. The transparent proxy status will be updated regardless of the filter’s status.

3.2.3 File system integrity checker

The file system integrity checker is a tool to help you keep track of the changes made to the configuration files. It keeps a list of all configuration files (in the file /etc/security/fwfschk.db.list), and during the installation it generates a database containing the MD5 checksum of each file.

The administrator can use this database to perform checks and verify which files were modified since the last update. The administrator can also update the database with the new checksum of the files after causing modifications to the configuration files.

Using the configuration client, double-click **File System Integrity Checker** inside the System Administration folder (see Figure 3 on page 21). The window shown in Figure 6 is displayed.

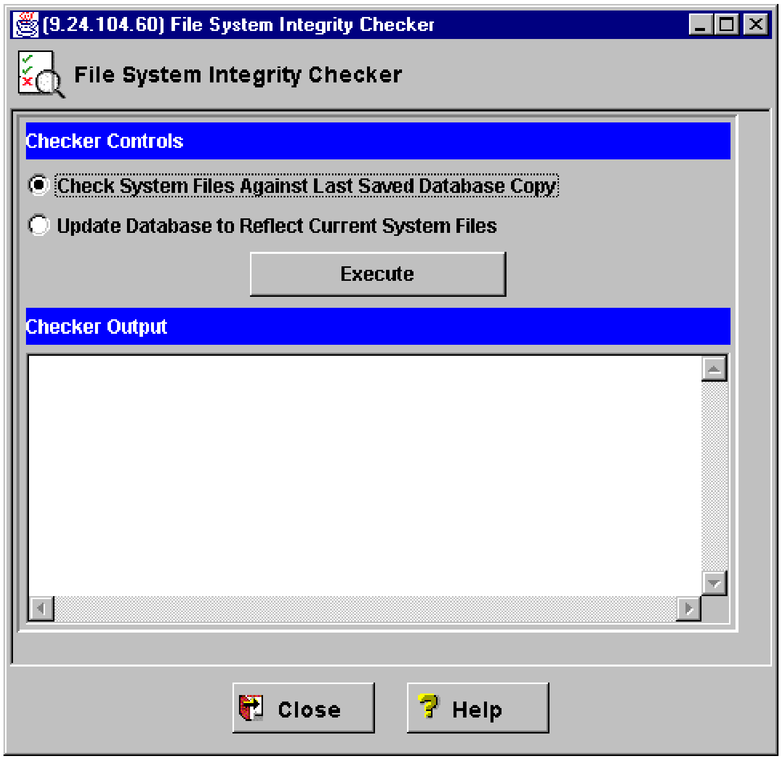

Figure 6. File System Integrity Checker window

To compare the stored checksum database and the current checksum of the configuration files, select **Check System Files Against Last Saved Database Copy** and click **Execute**. The results are shown in the Checker Output pane.

To update the database (in this case, you accept that all changes are valid), select **Update Database to Reflect Current System Files** and click **Execute**.

We recommend that you add the command `fwfschk` to the AIX crontab, so this check is automatically done, and if any inconsistency is found, an alert is logged (refer to Chapter 11, “Logging, monitoring, and reporting” on page 285 for more information about alerts).

The following line will cause the system to run the check every day at 4:00 AM:

```
0 4 * * * /bin/fwfschk -l
```

3.2.4 Users

Once the firewall is installed, you must not use AIX commands or smit to add new users at your firewall. Using the firewall Configuration Client GUI, you are able to use strong authentication for the users.

You can still use AIX commands to change any user’s attributes that are not available in the firewall Configuration Client GUI. For example, if you need to unlock a user’s account that was locked by exceeding the number of failed login attempts, run the following command:

```
# chsec -f /etc/security/lastlog -s <username> -a unsuccessful_login_count=0
```

3.2.4.1 Creating and administering users

Log on to the configuration client (see Figure 3 on page 21), and double-click **Users**. The Users window is displayed (see Figure 7), which lists the users that are authenticated by the firewall.

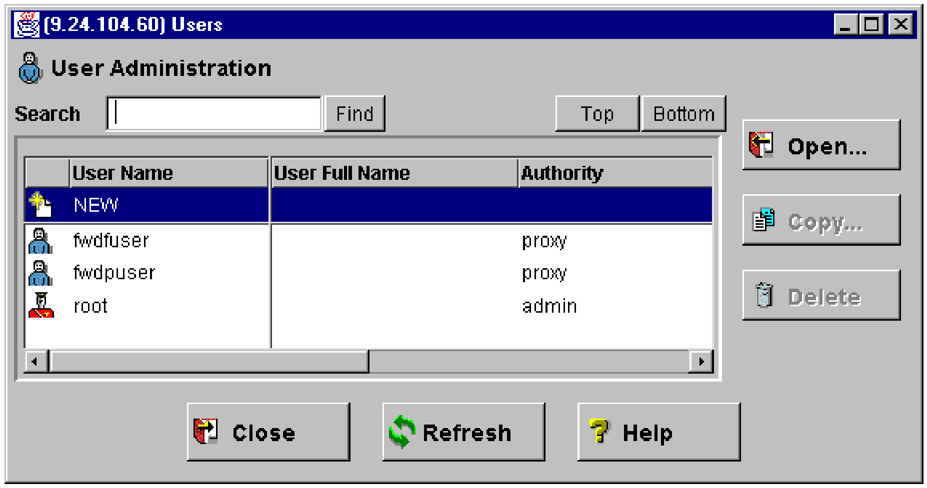

Figure 7. Users main window

When the firewall is installed, it creates an instance for the user root and adds two new users: fwdfuser and fwdpuser.

The user fwdfuser contains the default user authentication. Any user that has not been authorized as a proxy user is authenticated using the information from fwdfuser. Note that all authentication methods for this user are set to deny all. The user fwdpuser contains the default values for a new user. Both fwdfuser and fwdpuser can be changed, but they cannot be deleted.

To add a new user, click **NEW**, then click **Open**. The window shown in Figure 8 is displayed.

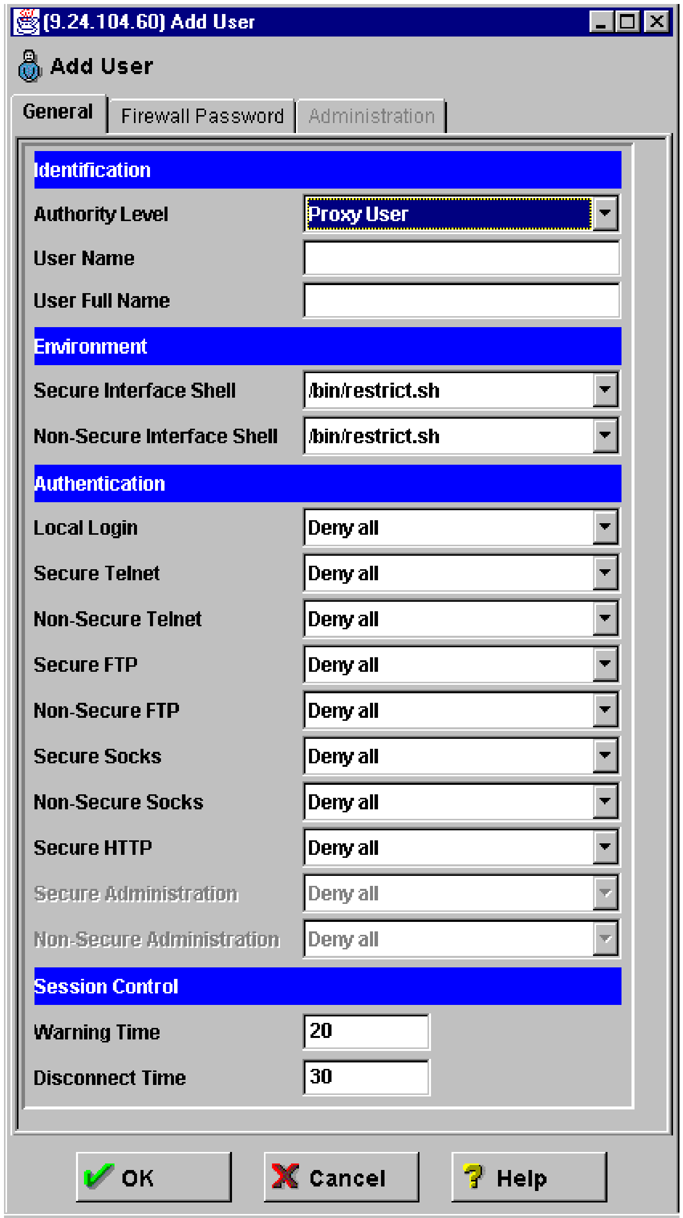

Figure 8. Add User window

The fields in this window are:

- **Authority Level**: the options available are Proxy User (default) or Firewall Administrator.

  By default, proxy users are assigned a restricted shell and they cannot perform administrative tasks. The firewall administrator performs those tasks using the command line interface or the configuration client. When selecting Firewall Administrator, the Administration tab becomes active (we will discuss this later in this section).

- **User Name**: this field defines the user account on the system, and it is limited to 8 characters.
- **User Full Name**: this can be used to store the user's full name or any other convention that you prefer. This field is optional.

- **Secure Interface Shell**: this field determines the shell for this user when the connection is established on the secure interface. The selections available are:

/bin/csh	The C shell
/bin/bsh	The Bourne shell
/bin/ksh	The Korn shell
/bin/restrict.sh	A restricted shell (default selection)
/bin/oneact.sh	A shell that performs a single action

- **Non-Secure Interface Shell**: this field specifies the shell for this user when the connection is established on the secure interface. The selections available are the same as the previous item.
- **Local Login**: this field specifies the method of authentication that is used when the user is logging in locally. The selections available are: Deny all (default), Permit all, Firewall password, SecurID card and User-Supplied. See 3.2.4.2, “Authentication methods” on page 31 for more details on these authentication methods.
- **Secure Telnet**: this field specifies the method of authentication that is used when connecting to the firewall from the secure network using Telnet. The selections available are: Deny all (default), Permit all, Firewall password, SecurID card and User-Supplied.
- **Non-Secure Telnet**: this field specifies the method of authentication that is used when connecting to the firewall from the non-secure network using Telnet. The selections available are: Deny all (default), Permit all, Firewall password, SecurID card and User-Supplied.
- **Secure FTP**: this field specifies the method of authentication that is used when connecting to the firewall from the secure network using FTP. The selections available are: Deny all (default), Permit all, Firewall password, SecurID card and User-Supplied.
- **Non-Secure FTP**: this field specifies the method of authentication that is used when connecting to the firewall from the non-secure network using FTP. The selections available are: Deny all (default), Permit all, Firewall password, SecurID card and User-Supplied.
- **Secure SOCKS**: this field specifies the method of authentication that is used when the user is accessing the SOCKS server from a SOCKS client Version 5 in the secure network. The selections available are: Deny all (default), Permit all, Firewall password and SecurID card.
- **Non-Secure SOCKS**: this field specifies the method of authentication that is used when the user is accessing the SOCKS server from a SOCKS client Version 5 in the non-secure network. The selections available are: Deny all (default), Permit all, Firewall password and SecurID card.
- **Secure HTTP**: this field specifies the method of authentication that is used when the user is accessing the HTTP Proxy from the secure network. The selections available are: Deny all (default), Permit all and Firewall password.
- **Secure Administration**: this field specifies the method of authentication that is used when the firewall administrator is accessing the configuration server from the secure network. The selections available are: Deny all (default),

Permit all, Firewall password, SecurID card and User-Supplied. Note this field is available only when the Authority Level field is set to Firewall Administrator.

- **Non-Secure Administration**: this field specifies the method of authentication that is used when the firewall administrator is accessing the configuration server from the non-secure network. The selections available are: Deny all (default), Permit all, Firewall password, SecurID card and User-Supplied. Note this field is available only when the Authority Level field is set to Firewall Administrator.
- **Warning Time**: this field specifies the maximum time in minutes that the user can remain idle before a warning message is issued to disconnect the user. The default value is 20 minutes.
- **Disconnect Time**: this field specifies the maximum time in minutes that the user can remain idle before being disconnected. The disconnect time must be greater than the warning time. The default value is 30 minutes.

After choosing the selections on the General tab, go to the Firewall Password tab for more fields regarding the password of this user (see Figure 9).

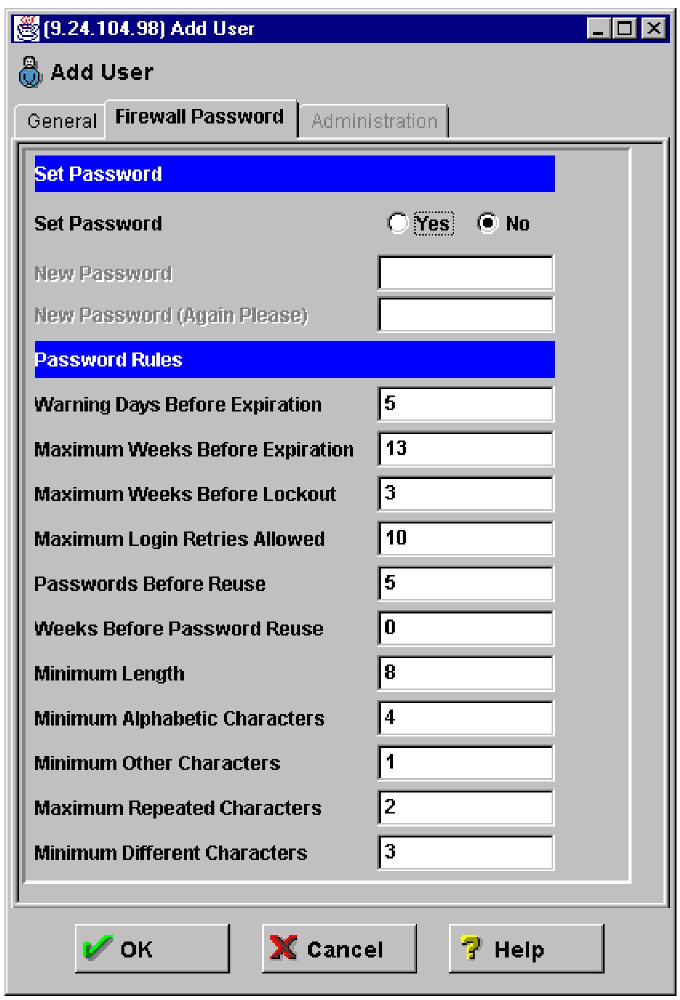

Figure 9. Firewall Password tab

The fields in this window are:

- **Set Password**: you choose whether you want to set a password for this user or not. If this field is set to Yes (the default is No), the fields New Password and New Password (Again Please) become available, so you can type in (once in each field) the new password for the user. The next time this user logs in, this password will be expired (so the user has to change it).
- **Warning Days Before Expiration**: this field specifies how many days before the expiration of the password the user is warned to change it. The default is 5 days.
- **Maximum Weeks Before Expiration**: this field specifies the maximum number of weeks that the user can remain with the same password. The default is 13 weeks.
- **Maximum Weeks Before Lockout**: this field specifies the number of weeks a password can remain unused before it is locked. The default is 3 weeks.
- **Maximum Login Retries Allowed**: this field specifies the maximum number of failed login attempts before the password is locked. The default is 10.
- **Passwords Before Reuse**: this field specifies the 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. The default is 5.
- **Weeks Before Password Reuse**: this field specifies the number of weeks that the passwords are kept in the password history list. The password cannot be changed to any password that is currently in the history list. When this field is set to zero, the field Passwords Before Reuse controls the reuse of passwords. The default is zero.
- **Minimum Length**: this field specifies the minimum number of characters in a password. The default is 8.
- **Minimum Alphabetic Characters**: this field specifies the minimum number of alphabetic characters in a password. The default is 4.
- **Minimum Other Characters**: this field specifies the minimum number of non-alphabetic characters in a password. The default is 1.
- **Maximum Repeated Characters**: this field specifies the maximum number of times any single character can be repeated in the password. The default is 2.
- **Minimum Different Characters**: this field specifies the minimum number of different characters in the password. The default is 3.

If you selected Firewall Administrator in the Authority Level field, you have one more tab available in this window, which is the Administration tab (see Figure 10).

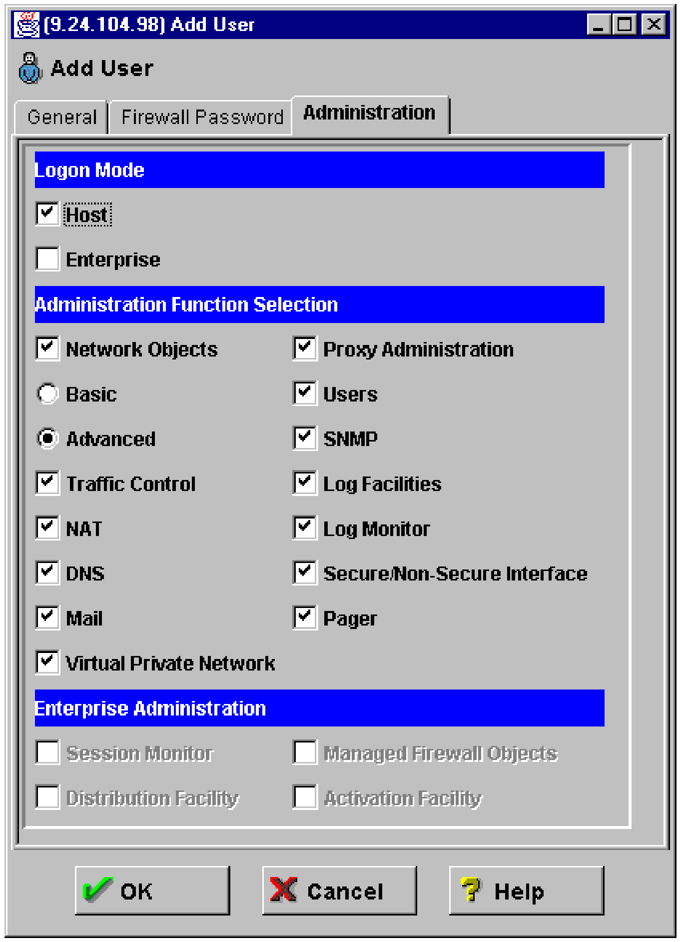

Figure 10. Administration tab

In this window you can specify which functions of the configuration client your administrator is able to configure.

If you select the **Host** check box, you are able to select the check boxes on the Administration Function Selection panel. If you select **Enterprise** you are also able to select the check boxes for the Enterprise Firewall Management (EFM)[2], located in the Enterprise Administration panel.

- **Network Objects**: select this check box to allow the administrator to configure the network objects. You can choose between the Basic access (this administrator can copy or add new objects but cannot modify or delete an existing object) or Advanced (the administrator can use all network objects functions).
- **Traffic Control**: select this check box to allow the administrator to configure the traffic control functions.
- **NAT**: select this check box to allow the administrator to configure the network address translation (NAT) functions.
- **DNS**: select this check box to allow the administrator to configure the Domain Name System (DNS) services.

[2] EFM is a configuration function that allows the firewall administrator to control and update several firewalls from one central server. For more information, refer to *IBM SecureWay Firewall for AIX User's Guide Version 4, GC31-8419.*

- **Mail**: select this check box to allow the administrator to configure the Secure Mail Proxy functions.
- **Virtual Private Network**: select this check box to allow the administrator to configure the virtual private network (VPN) functions.
- **Proxy Administration**: select this check box to allow the administrator to configure the proxy services.
- **Users**: select this check box to allow the administrator to configure the user's functions. A generic firewall administrator is not allowed to create or modify any other administrator users (this can be done only by the root user).
- **SNMP**: select this check box to allow the administrator to configure the SNMP agent functions. For more information on the firewall SNMP agent refer to Chapter 13, "Firewall Management" on page 367.
- **Log Facilities**: select this check box to allow the administrator to configure the log facilities functions.
- **Log Monitor**: select this check box to allow the administrator to configure the log monitor functions.
- **Secure/Non-secure Interface**: select this check box to allow the administrator to configure the status of the interfaces of the firewall.
- **Pager**: select this check box to allow the administrator to configure the pager functions. For more information on the pager functions refer to Chapter 11, "Logging, monitoring, and reporting" on page 285.

The following check boxes are available in the Enterprise Administration panel. For more information on these functions, refer to *IBM SecureWay Firewall for AIX User's Guide Version 4,* GC31-8419.

- **Session Monitor**: select this check box to allow the administrator to configure the session monitor.
- **Distribution Facility**: select this check box to allow the administrator to configure the distribution facility.
- **Managed Firewall Objects**: select this check box to allow the administrator to configure the managed firewall objects.
- **Activation Facility**: select this check box to allow the administrator to configure the activation facility.

3.2.4.2 Authentication methods

An authentication method is used by the firewall to validate the user before it has access to a certain service (Telnet, FTP, SOCKS, and so forth).

The following are the methods provided by IBM SecureWay Firewall V4.1 for AIX (some options on this list may not be available to the service you are configuring):

- **Deny all**: this is the default selection for all services. By using this method, the user has no access to the service (Telnet, FTP, configuration client, and so forth).
- **Permit all**: the user is not authenticated but has access granted to the service.

- **Firewall password**: the user is prompted for a valid AIX password. You can change the password characteristics using the Configuration Client GUI (refer to the previous section for more information).

  When the password is set by the administrator, the user must change it. This new password must comply with the characteristics established by the administrator.

- **SecurID card**: the authentication is done using a Security Dynamics SecurID card or pinpad card. The PIN must be set before using this authentication method. For more information on using this authentication method, refer to *IBM SecureWay Firewall for AIX User's Guide Version 4,* GC31-8419 and *Guarding the Gates Using the IBM eNetwork Firewall V3.3 for Windows NT,* SG24-5209.

- **User-supplied**: the administrator creates and compiles a subroutine that is used as authentication method. You can have one user-supplied method at any given time. Refer to the *IBM SecureWay Firewall Reference Version 4 Release 1,* SC31-8418 for more information on this authentication method.

3.2.4.3 Proxy user and firewall administrator

The main usage for a proxy user is to access the Telnet and FTP proxy. If you specify a nonrestricted command shell, for example /bin/ksh, this user also has access to AIX commands as a common AIX user. This means that this user has no administration privileges - neither AIX administration nor firewall administration.

A firewall administrator can also be configured to use the restricted shell. In this case, the administrator has to perform all the firewall administrator functions using the configuration client. If this administrator uses a nonrestricted shell, such as /bin/ksh, this user can also run the firewall administration commands (see also 3.2.6, "Basic configuration using the command line interface" on page 37).

Next, we give an example of a firewall administrator user. We will discuss the proxy user in more detail in Chapter 6, "Proxy" on page 87.

We created a firewall administrator called "fwadm", using the following configuration (the fields we do not mention here remained with their default values):

- Authority level: firewall administrator
- User name: fwadm
- Secure interface shell: /bin/ksh
- Local login: firewall password
- Secure administration: firewall password
- Initial password: test

On the Administration tab (on the Administration Function Selection panel), we selected:

- Network objects (basic)
- DNS
- Mail
- SNMP
- Secure/Non-secure Interface

We connected to the firewall using the configuration client. On our first login, we were prompt for a new password (see Figure 11).

Figure 11. User being prompted for a new password

Then we tried to input a password that would not match the requirements we configured. The message we received is shown in Figure 12.

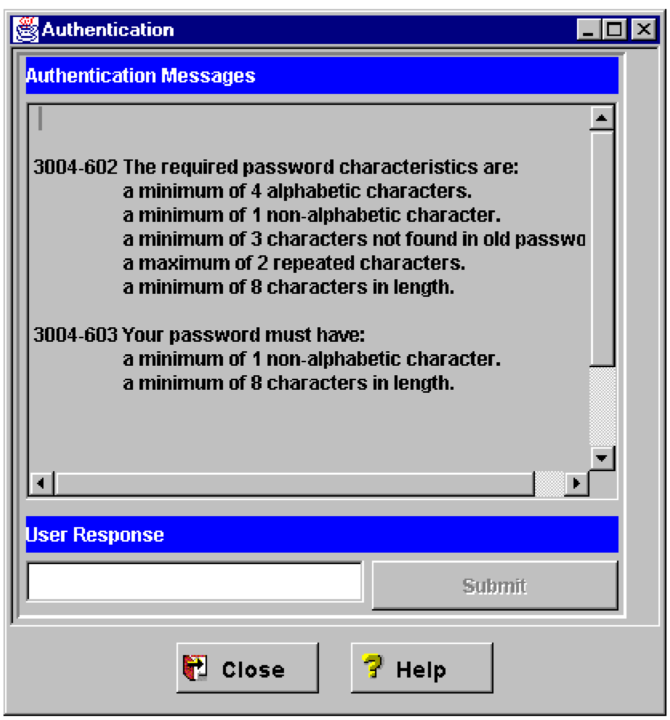

Figure 12. Error when trying to provide an invalid new password

The message 3004-603 shows the user which requirements were not fulfilled by the new password.

We provided a valid password and successfully logged in the configuration client. The main panel contained only the functions we allowed for this administrator (see Figure 13):

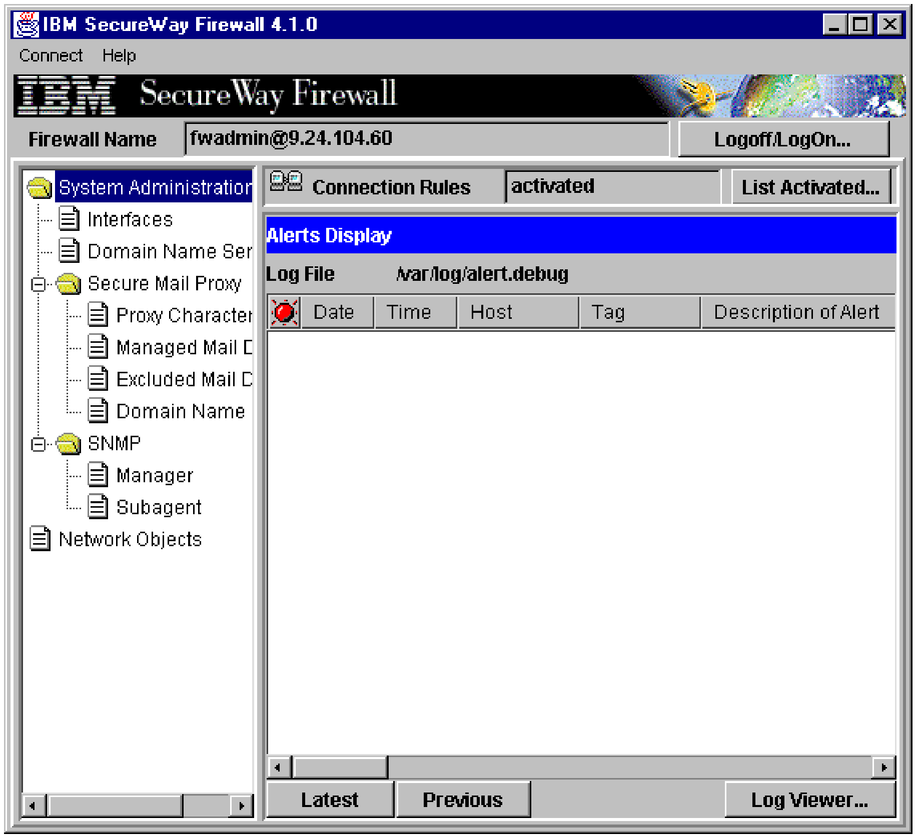

Figure 13. Configuration client window for a restricted firewall administrator

Then we logged in to the firewall console, and we did some tests using some firewall commands (refer to 3.2.6, “Basic configuration using the command line interface” on page 37 for more information).

In the following example, we tried two commands that this administrator is allowed to run:

```
# fwadapter cmd=list

10.2.1.1        Non-Secure Interface
9.24.104.98     Secure Interface

# fwnwobj cmd=add name="Secure Network" type=Network \
addr=9.24.104.0 mask=255.255.255.0
```

The next example shows three commands that this administrator is not allowed to use. The administrator could use them with the list parameter, but no modification on the configuration could be done using them.

```
# fwconns cmd=create name="Test" source="Secure Network" \
destination="The World" servicelist=10

fwadm does not have authorization for the following functional group:
Traffic Control.

# fwuser cmd=change username=fwadm secadmin=permit
fwadm does not have authorization for the following functional group:
Users.

# fwfilter cmd=shutdown
fwadm does not have authorization for the following functional group:
Traffic Control.
```

Note that the error messages show which check box from the user's characteristics should be selected to allow this administrator to do that specific task. For example, the `fwconns` needs to be allowed by selecting the check box **Traffic Control**.

The next example shows one of the commands that the administrator could not use, but this time we ran it with the list parameter, so we can see the output.

```
# fwuser cmd=list

fwadm
fwdfuser
fwdpuser
root
```

3.2.5 Basic logging

We recommend that you set up basic logging before continuing the configuration of the firewall, because it will be very useful for debugging problems later on.

First, log in to the configuration client and double-click **System Administration**, then double-click **System Logs** (see Figure 3 on page 21). To add or change the logging configuration, double-click **Log Facilities**. The window in Figure 14 is displayed.

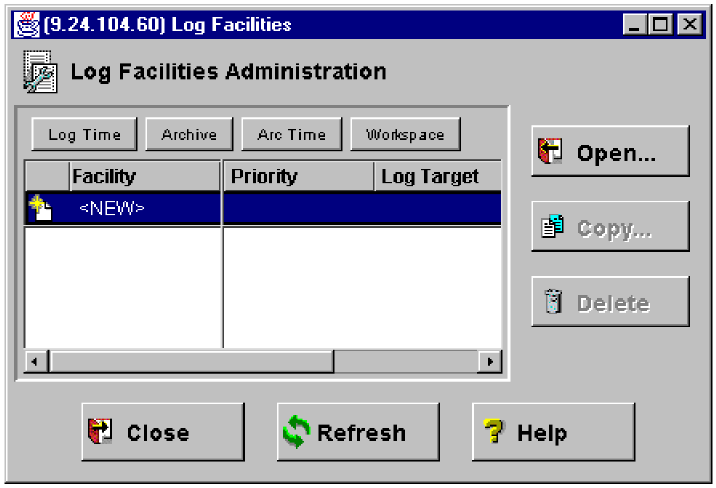

Figure 14. Log Facilities window

Select **NEW**, then click **Open** and the window in Figure 15 is displayed.

Figure 15. Add Log Facilities window

In our basic configuration, we used Filename as the type of logging, so all messages sent to syslog by the firewall are appended to an ASCII file (refer to Chapter 2, “Installation” on page 7 for tips on creating the file systems for logging and archiving).

For more information on the options available on the window in Figure 15 and how to manage the log files, refer to Chapter 11, “Logging, monitoring, and reporting” on page 285.

We created the following facilities:

1. **Firewall Log** (general firewall log, including filter logging):
   - Priority: debug
   - Log Filename: /var/log/firewall.debug
   - Archive Management: Enable
   - Days Until Archive: 7
   - Archive Filename: /var/log/archive/firewall.debug.a
   - Days until purge: 30
   - Workspace: /var/log/workspace
2. **Alert Log** (log monitor threshold violation warnings):
   - Priority: debug
   - Log Filename: /var/log/alert.debug
   - Archive Management: Enable
   - Days Until Archive: 7
   - Archive Filename: /var/log/archive/alert.debug.a
   - Days until purge: 30
   - Workspace: /var/log/workspace

Refer to Chapter 11, "Logging, monitoring, and reporting" on page 285 for more information.

For more information on syslog and other facilities you can use, refer to the syslogd manual page. If you have the AIX online manuals installed on your system, type the following command:

```
# man syslogd
```

3.2.6 Basic configuration using the command line interface

In IBM SecureWay Firewall V4.1 for AIX, the smit panels are no longer available. The administrator can still use the command line interface to manage the firewall.

If you do not have a graphics adapter attached to your firewall box, you can still use the remote configuration client. You need to do a basic configuration on your firewall to allow the connection from the remote client, and set up the configuration server to accept this connection.

We will now demonstrate the use of the command line interface to allow the connection from a remote configuration client.

Perform these steps right after the firewall installation.

1. First, you have to configure the secure adapter. You can list all adapters and identify the secure one(s), then you change the state of these adapters (in this example, 9.24.104.98 is our secure adapter):

```
# fwadapter cmd=list

10.2.1.1        Non-Secure Interface
9.24.104.98     Non-Secure Interface

#fwadapter cmd=change addr=9.24.104.98 state=secure

Command completed successfully.

# fwadapter cmd=list

10.2.1.1        Non-Secure Interface
9.24.104.98     Secure Interface
```

2. Create a network object for your remote configuration client (in this example, the IP address of our client is 9.24.106.97):

```
# fwnwobj cmd=add name="Config Client" type=Host \
addr=9.24.106.97 mask=255.255.255.255

# fwnwobj cmd=list format=long

               id = 501
             type = Host
             name = Config Client
             desc =
             addr = 9.24.106.97
             mask = 255.255.255.255

               id = 1
             type = Network
             name = The World
             desc =
             addr = 0
             mask = 0
```

3. Create the network objects for your firewall interfaces (for more information about network objects, services and connections refer to Chapter 4, “Packet filters” on page 43):

```
# fwnwobj cmd=add name="FW secure" type=firewall \
addr=9.24.104.98 mask=255.255.255.255

# fwnwobj cmd=add name="FW nonsecure" type=firewall \
addr=10.2.1.1 mask=255.255.255.255

# fwnwobj cmd=list

501   Host         Config Client
503   Firewall     FW nonsecure
502   Firewall     FW secure
1     Network      The World
```

4. Identify the predefined services for the remote configuration client:

```
# fwservice cmd=list | grep "Config Client"

31   Config Client non-secure      Permit use of config client from non
30   Config Client secure          Permit use of config client from sec
```

Write down the ID of the service Config Client secure, because you will need it in the next step (in this example, the ID is 30).

5. Add a connection using the service "Config Client secure" (ID=30) from your remote client ("Config Client" object) to the firewall ("FW secure" object):

```
# fwconns cmd=create name="Config Client" desc="permit \
  connection from secure config client" source="Config Client" \
  destination="FW secure" servicelist=30

# fwconns cmd=list

501 Config Client     permit connection from secure config client
```

6. Check the status of the configuration server. Then, change it to accept remote connections:

```
# fwcfgsrv cmd=list

localonly = yes
encryption = none
sslfile = /etc/security/fwkey.kyr

# fwcfgsrv cmd=change localonly=no

Command completed successfully.

# fwcfgsrv cmd=list

localonly = no
encryption = none
sslfile = /etc/security/fwkey.kyr
```

7. Change the attribute of the user root to allow remote administration from the secure network. If you prefer to create a new administrator user instead of using root, go to the next step.

```
# fwuser cmd=change username=root secadmin=password
```

8. If you changed the attribute of the root user, you do not need to perform this step.

   Create a new administrator user adm1 with secure administration allowed. In this example, we used the firewall password as the authentication method and we allowed this administrator to use all available functions at the configuration client. We set the initial password to "test", but it must be changed the first time this user logs in.

```
# fwuser cmd=add username=adm1 password=yes pwdvalue=test \
level=admin secshell=/bin/ksh loclogin=password secauth=password \
secadmin=password fg_all=yes

# fwuser cmd=list username=adm1

username=adm1 fullname= level=admin secshell=/bin/ksh
remshell=/bin/restrict.sh loclogin=password secftp=deny remftp=deny
secauth=password remauth=deny secadmin=password remadmin=deny
warntime=20 disctime=30 loginretries=10 pwdwarntime=5
histsize=5 histexpire=0 maxexpired=3 maxage=13
minlen=8 minalpha=4 minother=1 maxrepeats=2
mindiff=3 modeallowed=host fg_all=yes fg_netobjs1=yes
fg_netobjs2=yes fg_interfaces=yes fg_dns=yes fg_mail=yes
fg_logs=yes fg_logmonitor=yes fg_pagers=yes fg_snmp=yes
fg_sesslfm=yes fg_user=yes fg_proxyserver=yes fg_traffic=yes
fg_vpn=yes fg_addrtrans=yes fg_clone=no fg_dist=no
fg_act=no fg_secag=no secsocks=deny remsocks=deny
sechttp=deny
```

9. If you want to activate the filters now, use the following command:

```
# fwfilter cmd=update
```

Note that you do not have to activate the filters to be able to access the configuration server, because the filters are initially deactivated. We suggest that you add a connection to allow the access to the configuration server so you will not have any problem using the configuration client after activating the filters.

Now you are able to continue your configuration using the remote configuration client.

3.2.6.1 Useful commands

The following is a list of some of the available commands. For more detailed information and documentation on the available commands refer to *IBM SecureWay Firewall Reference Guide for AIX Version 4*, SC31-8418.

- List the state of the interfaces:

  `fwadapter cmd=list`

- List the current status of the configuration server:

  `fwcfgsrv cmd=list`

- Invoke the file system integrity checker:

  `fwfschk -l`

- Update the file system integrity checker database:

  `fwfschk -u`

- List all firewall users:

  `fwuser cmd=list`

- List current active filters:

```
fwfilter cmd=list
```

- Rebuild the configuration and activate the filters:

```
fwfilter cmd=update
```

- Deactivate the filters:

```
fwfilter cmd=shutdown
```

- List the connections:

```
fwconns cmd=list
```

- Refresh the secure mail daemon:

```
fwsecuremail cmd=refresh
```

- List NAT configuration:

```
fwnat cmd=list
```

3.2.7 Further configuration

Once you have this basic configuration ready, the next step is to configure the services you listed in your planning worksheet.

We suggest the following sequence to configure your firewall:

- Filters (connections, services and rules)
- DNS
- Proxy (including adding new proxy users)
- SOCKS
- Mail
- NAT
- VPN

Note that you may not need to configure all those services.

Refer to their respective chapters for more information on how to configure each service.

Chapter 4. Packet filters

Packet filtering is the basis for controlling network traffic through the firewall. By using Wizard or Security Policy, the fundamental filter rules can be set up and activated. By using expert filters, more granular session level traffic control is possible, based on multiple criteria such as protocols, interfaces, IP addresses, directions, time of day, subnet and so on. Filters do not impact the firewall routing tables.

By default the firewall does not allow any traffic flow between the secure and non-secure networks. Connections should be set up to allow specific types of traffic to flow between the secure and non-secure networks.

The control mechanism of the packet filters is based on a set of rules, services, connections and the participating objects.

You can get a good review of the TCP/IP protocol and some common attacks in the redbook *Guarding the Gates Using IBM eNetwork Firewall for Windows NT,* SG24-5209, in Chapter 7, "Coming to Grips with IP Packets."

In this chapter we review the IBM SecureWay Firewall V4.1 for AIX packet filtering mechanisms with typical examples.

4.1 Filter structure

The IBM Firewall filters are created and operated through the following steps:

- Define source and destination *objects* that will be connected for data transmission. An object can be an individual or a group of hosts, networks, firewalls, routers, and interfaces.
- Define deny or permit *rules* that will be applied to control data traffic between the source and destination objects. If appropriate rules are already defined, they can be reused with slight modifications.
- Define *services* that will be applied to a pair of source and destination objects. A service is composed of one or more rules. If there appropriate services are already defined, they can be reused with slight modifications.
- Define *connections* that will allow a pair of source and destination objects to use the designated services.
- *Activate* the connection.

Predefined or standard templates for rules and services are shipped with the firewall. So it is easy to construct your own filters by using them. You can just copy one of them as necessary, give an appropriate new name to it, adjust some parameters according to your environment, save it, and use it.

Figure 16 shows a conceptual filter structure.

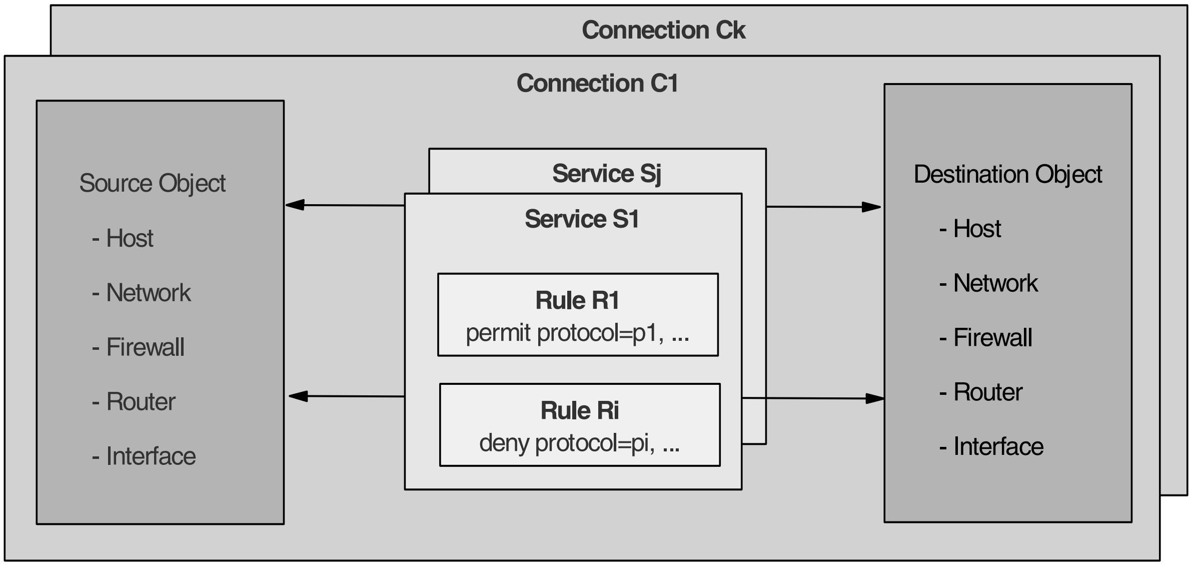

Figure 16. IBM Firewall filter structure

In Figure 16, each shadowed box represents a filter component and its inclusion property. For example, a rule box represents a fundamental or basic filter rule template, which contains no other component. A deny or permit rule is defined within this box. A service box is composed of one or more rules, forming a meaningful service such as FTP, Telnet, HTTP, Mail, and so on. A connection box is composed of one or more services, a source object, and a destination object, completing the filter definition. Connection is the last step in the process of defining a filter.

Note that by using this *box* or *object-oriented* filter structure, you can maximize the benefit of reusability and the ease of definition, change, expansion, maintenance, and so on.

4.1.1 Object and Group

An object is a representation of a network component. It is defined by an IP address and an address mask, so it is possible for one object to represent a whole range of network addresses. A group is a collection of one or more objects. Possible object types are:

- **Host**: A node in your network with mask 255.255.255.255.
- **Network**: A set of IP addresses with a specific mask.
- **Firewall**: The firewall interface with mask 255.255.255.255.
- **Router**: A unique IP address with mask 255.255.255.255.
- **Interface**: A network adapter with mask 255.255.255.255.

When you want to define a new object or group, you must select the option **Network Object**, at the left of the main screen of the GUI. For convenience, you can also invoke the object definition screen from another stage, such as connection screen, where you find that object definitions are missing.

Table 3 on page 45 shows the necessary parameters for network object definition.

Table 3. Network object definition parameters

Category	Field	Subfield	Selection	Description
Single	Identification	Object Type	Host	Node in a network
			Network	Network or subnetwork
			Firewall	Firewall itself
			Router	Router
			Interface	Interface of firewall or other host
		Object Name		Object name
		Description		Object description
	IP Information	IP Address		IP or network address of an object
		Subnet Mask		Subnetwork mask
Group	Identification	Group Name		Name of group
		Description		Group description
	Group Composition	Objects in Group		Select a single object from the object list

Object Name is the name of the object. When you key in the prohibited characters (|, }, {,], [, etc.) in the object name field, the beep will notify you. Try again.

Description describes the object.

IP Address specifies an IP address or the range of IP addresses for this object.

Subnet Mask depends on the type of object you are defining. The subnet mask automatically changes, but you can override it if needed.

The only pre-defined object is *The World*, an object that is matched by any IP address.

A *Group* object is defined by selecting one or more single objects which are already defined. You cannot include a group in another group. A group object has its own name and description.

4.1.2 Rules

We have already discussed how rules are combined within services which are, in turn, embedded within connection definitions. Let us now look at the parameters for rule definitions in more detail.

Table 4 on page 46 shows the necessary parameters for rule definition.

A set of rules are provided with the firewall product. These rules will cover almost all rules necessary to define a service. You can use the predefined rules or can copy them to another name and modify them according to your requirements.

Table 4. Rule definition parameters

Field	Subfield	Selection	Description
Identification	Rule Name		Rule name
	Description		Rule description
	Action	Permit	Permitspecified transmission
		Deny	Deny specified transmission
	Protocol	all	All protocols
		tcp	Transmission Control Protocol
		tcp/ack	TCP with Acknowledgment
		udp	User Datagram Protocol
		icmp	Internet Control Message Protocol
		ospf	Open Shortest Path First protocol
		ipip	IP-in-IP protocol
		esp	Encapsulating Security Protocol
		ah	Authentication Header protocol
	Numeric Protocol		Decimal protocol number(0=any)
Source Port/ ICMP Type	Operation	Any	Any Port# /Type
		Equal to	Equal to Port#/Type
		Not equal to	Not equal to Port#/Type
		Less than	Less than Port#/Type
		Greater than	Greater than Port#/Type
		Less than or equal to	Less than or equal to Port#/Type
		Greater than or equal to	Greater than or equal to Port#/Type
	Port#/ Type		Decimal Port# or ICMP Type
Destination Port/ ICMP Code	Operation	Any	Any Port# /Code
		Equal to	Equal to Port#/Code
		Not equal to	Not equal to Port#/Code
		Less than	Less than Port#/Code
		Greater than	Greater than Port#/Code
		Less than or equal to	Less than or equal to Port#/Code
		Greater than or equal to	Greater than or equal to Port#/Code
	Port#/ Code		Decimal Port# or ICMP Code

Field	Subfield	Selection	Description
Interface Settings	Interface	Both	Both secure and non-secure interface
		Secure	Secure interface only
		Non-secure	Non-secure interface only
		Specific	Specific interface
	Name		Name of specific interface
Direction/ Control	Routing	both	Both local and route
		local	Local only
		route	Route only
	Direction	both	Both inbound and outbound
		inbound	Inbound only
		outbound	Outbound only
	Log Control	yes	Log if this rule matches
		no	Do not log
	Frag. Control	Yes	Matches fragment headers, fragments and non fragments. For fragment, port information is ignored and assumed to match.
		No	Matches only fragments and fragment headers. For fragment headers, port information must match. For fragments, port information is ignored.
		Only	Matches only non fragments. Fragment headers and fragments are excluded by this parameter.
		Header	Matches only non-fragments and fragment headers. Fragments are excluded by this parameter.
Tunnel Information	Tunnel ID		This filter rule will apply to the selected tunnel ID.

Rule Name is the name of the rule. Use a naming convention for keywords used in this name, so rule names will be consistent. This makes it easier to search for them within a list.

Description describes the function of the rule.

Action has the value permit or deny. Any IP packet that matches the other fields in the filter definition will either be passed or blocked, depending on the value of this field. You can also specify a *protocol* by number or name.

The first field of *Source Port/ ICMP Type* specifies the type of operation, and the second the desired port number (for ICMP packets it is the ICMP Type of the message). The port operation field is an arithmetic operator field as shown in the table. The operator is applied to the desired port field, so, for example, if the two fields were greater than 1023, we would match packets only with a source port number of 1024 or higher.

The pair of *Destination Port/ ICMP Code* fields is used in the same way as the source port fields to define which destination port(s) we want the filter to match. For ICMP packets, it refers to the ICMP Code field.

Interfaces Settings defines which interface the packet is flowing through. A specific interface can be defined, for example, when there is more than one secure interface.

In some cases the firewall may act as a router, in which case packets flow through it. In other cases the packets may go to an application on the firewall machine itself (such as a proxy server). This field defines whether the packet has a destination or source of the firewall, or whether the destination and source are both addresses other than the firewall (in which case the firewall is behaving as an IP router). *Local* specifies the traffic that comes to or from the firewall itself. *Route* traffic specifies the traffic that goes through the firewall. *Both* specifies the local and route traffics.

Direction defines whether the packet is coming into or going out of the adapter where the rule is applied. Remember that the rule can apply at any of the firewall adapters, controlled by the *Interface* definition (above).

Log Control defines if the packet should be logged or not. The default log control setting for permitted packets (those that pass the rule) is no and for denied packets is yes. It is important to log extensively on a firewall, because you cannot tell in advance which piece of seemingly unimportant log data will reveal an attack. However, logging every successfully transmitted packet is usually more than you need.

Fragmentation Control controls the fragmented packets, where a packet is divided more that one packets. You can use this field to define the matching rules for fragmented packets.

Tunnel ID identifies the tunnel through which the packet must be sent.

4.1.3 Service

A service defines the type of IP traffic that is permitted or denied between two source and destination objects. For example, you could construct a service to permit Telnet, or a service to deny Ping.

A service is built of one or more rules. IBM Firewall provides you with a large collection of commonly required rules that are predefined by development. When building a service, you usually use these predefined rules. If you don't find the rule that you need, you have to create an extra rule before you define a service.

You also have the ability to move rules up or down in the service, to create a specific order of the rules.

Table 5 shows the necessary parameters for service definition.

Table 5. Service definition parameters

Field	Subfield	Selection	Description
Identification	Service Name		Service name
	Description		Service description
Service Composition	Rule Objects	Flow	You select the necessary rules from rule definitions. The direction is automatically converted to inbound and outbound arrows.
		Name	
		Description	
Service Override Values	Override Log Control	no override	Log control of selected rules will be used
		yes	Override Log Control to yes
		no	Override Log Control to no
	Override Fragmentation Control	no override	Fragmentation control of selected rules will be used
		yes	Override Fragmentation control to yes
		no	Override Fragmentation Control to no
		only	Override Fragmentation Control to only
		headers	Override Fragmentation Control to headers
	Override Tunnel ID		Select a tunnel definition for this service.
Time Control	Control by Time of Day	Begin	Begin time in hh:mm. At this time the time control action is started.
		End	End time in hh:mm. At this time the time control action is stopped.
	Control by Days	Week Days	Specify From-To Week Days: during these week days the time control action is taken.
		Calendar Dates	Specify From-To Calendar Dates: during these calendar dates the time control action is taken.
	Time Control Action	Activate Service During Specified Time	The service is activated during specified time and days.
		Deactivate Service During Specified Time	The service is deactivated during the specified time and days.

Service Name is the name of the service. To simplify searching, you should use a naming convention for keywords used in this name with appropriate descriptions.

When you configure a service you do not specify the objects (that is, network addresses) between which it operates; you define the objects when you place the service in a connection definition. However, you do need to know what type of objects a rule applies to, because you have to define the direction of flow for each rule within the service definition. For example, a service that defines a TCP session from a client to a proxy server on the firewall will only operate as intended if it is included in a connection whose destination object is a firewall IP address.

You must add the rules that you need for this service, and you can move rules up or down to establish the correct order of the rules in the service. Order may be important, because some rules contained in a service may be more restrictive than others. If the less restrictive rule is at the top of the rule list, the packet may never be tested against the more restrictive rule.

The other element of the service composition section is the flow button. This defines whether the rule applies for packets going from the source to the destination object, or to returning packets (those going from destination to source). Very often a service contains an even number of rules, in pairs, with one of the pair controlling the flow in one direction and the other controlling the reverse direction.

Override Log Control has the values yes, no or no override. No override (the default) will let the settings in the rules apply. If you select yes, a log record will be generated for every packet that matches the rule, regardless of the log control setting in the rules that make up the service. This is useful for debugging, but you will not normally want to log so extensively. If you select no, no log records will be generated regardless of the settings in the rules.

Override Frag. Control allows you to override fragmentation settings in the rules.

If you enter a *Tunnel ID* in this field, the session will be passed through the specified secure tunnel.

With the *Time Control* feature, you can activate and deactivate the service during specified times, dates or days of the week.

The flow is indicated by an arrow at the left of the rules. If the rule has a green arrow (arrow points to the right), the filter defined by the rule applies to packets flowing from the source to the destination object. If the rule has a blue arrow (arrow points to the left), the source object and destination object are swapped, so the rule applies to flows from destination to source. You can see these arrows in Figure 27 on page 61.

4.1.4 Connection

A connection defines the IP traffic that is allowed or denied between a pair of network objects. A connection is built on a source and a destination object which are connected by a service component, as shown in Figure 16 on page 44. A service defines the type of IP traffic that is permitted or denied between the source and destination. The source and destination in a connection are each defined by an object or a group of objects.

For example, imagine you have a connection that permits Telnet between a client in the secure network and the proxy server on the firewall. The service in this case is Telnet. To be precise, it is a session from an unprivileged client port to TCP port 23. The source object in this case is any IP address in the secure network, and the destination object is the firewall.

We normally think of a connection definition as something that permits defined services to be used by a pair of network objects. However, they can also be defined to block the defined service. Remember that the firewall only allows the services that are explicitly permitted by the connection definition. All others will be blocked by default.

When you explicitly deny a specific service, it may reduce the matching overhead by dropping off the IP packets that are not allowed before it runs the whole range of rule comparisons that will result in match failure.

Table 6 shows the necessary parameters for connection definition.

Table 6. Connection definition parameters

Field	Subfield	Selection	Description
Identification	Name		Connection name
	Description		Connection description
	Source		A source object selected from the list
	Destination		A destination object selected from list
	Position	Upper Layer	Locate before dynamic filter layer
		RealAudio Layer	Locate at RealAudio layer
		Lower Layer	Locate after dynamic filter layer
Connection Services	Name		Name of the selected service
	Description		Description of selected service
SOCKS	Name		Name of selected SOCKS service
	Description		Description of selected SOCKS service

Name is the name of the connection. This name indicates the complete connection for a pair of objects, which includes service and rules.

Description describes the connection.

Source and Destination form a pair of traffic partners through the firewall. You can select a single or a group object from the list.

Position defines the location of the connection in a connection category hierarchy. You can choose to place the connection either before dynamic filters (Upper Layer) or after dynamic filters (Lower Layer). Once you save a connection you cannot change the position but you can reorder connections within their own position type. The upper layer has rule match comparison priority over the lower layer. The position layers are as follows, in the order of match comparisons:

- **Upper Layer**: Has priority over other layers.
- **Dynamic Filter Rules**: Dynamic filter rules are implicitly activated when a VPN tunnel is activated. Currently it has very coarse granularity.
- **RealAudio Layer**: This is for RealAudio filters.
- **Lower Layer**: Has the least priority.

Connection Services is the name and description of the selected service. You may select one or more service definitions to give a meaningful service to both objects.

Connection can be created by using a SOCKS service. SOCKS are defined by using the SOCKS Administration panel. See Chapter 7, “SOCKS server” on page 137 for more information.

4.1.5 Creation conventions

When you create new objects, rules, services, and connections, consider making your own conventions so that you can properly manage many kinds of related components. Despite the dangers of diversity, we recommend three things:

1. Utilize the predefined rules and services as much as possible. Almost all necessary definitions are provided as standard. You can select one, copy, give it a good name, adjust the necessary parameters, save it, and use it. Do not try to change predefined definitions directly.
2. Make names that can be easily distinguishable, from the predefined template. Make names and descriptions as short as possible, without sacrificing legibility. Even though the input field may be long enough for the name and description, about100 characters each, almost all display fields cannot fit a long name or description.
3. Think of the maintenance of your own rules and services. It is common to reuse rules and services, so when you change a rule or a service it can affect many related components.

In fact, there is no need to significantly change predefined rules or services because they are defined based on protocols that are already standardized and will not change for a long time. So it is possible to define rules and services independent of specific network environments. That’s why a connection is broken into objects, rules and services. In this sense, we can define a connection as a binding process of independent logical services (a set of independent logical rules) to specific real network objects.

4.2 Definition flow

To set up the rule base for your firewall in a structured way, it is important that you have a clear picture of your network infrastructure and the services that you want to provide. In this way it is easy to configure your firewall and maintain a consistent set of objects, rules, services and connections.

When you want to implement your connections there are basically two possible types of connections. The first type, and the easiest to implement, is a standard connection. This connection can be built with predefined services.

The following services are predefined:

- Telnet
- FTP
- HTTP
- SOCKS
- SSL
- SMTP
- RealAudio

- Identd
- SNMP
- Ping
- DNS
- SecureID
- Tunnels
- Remote Logging
- Firewall Configuration

The other type of connections cannot be built from predefined services, so you must define your own services. When defining services, you will probably have to define your own rules as well.

We will first describe how to set up a standard connection (one made up of predefined services and rules) and after that a nonstandard connection.

4.2.1 Standard connections

In this part we explain how to define a connection in the IBM Firewall. The example we use is a Telnet connection between the secure network and the firewall secure interface. This is something you are very likely to need to allow an administrator to log in to the firewall for maintenance purposes.

First, start the configuration GUI using the `fwconfig` command and then select the **Connection Setup** option from the navigator panel on the left of the display. A list of existing connections will appear, as shown in Figure 17. In the list select **<NEW>** and click on **Open** to create a new connection.

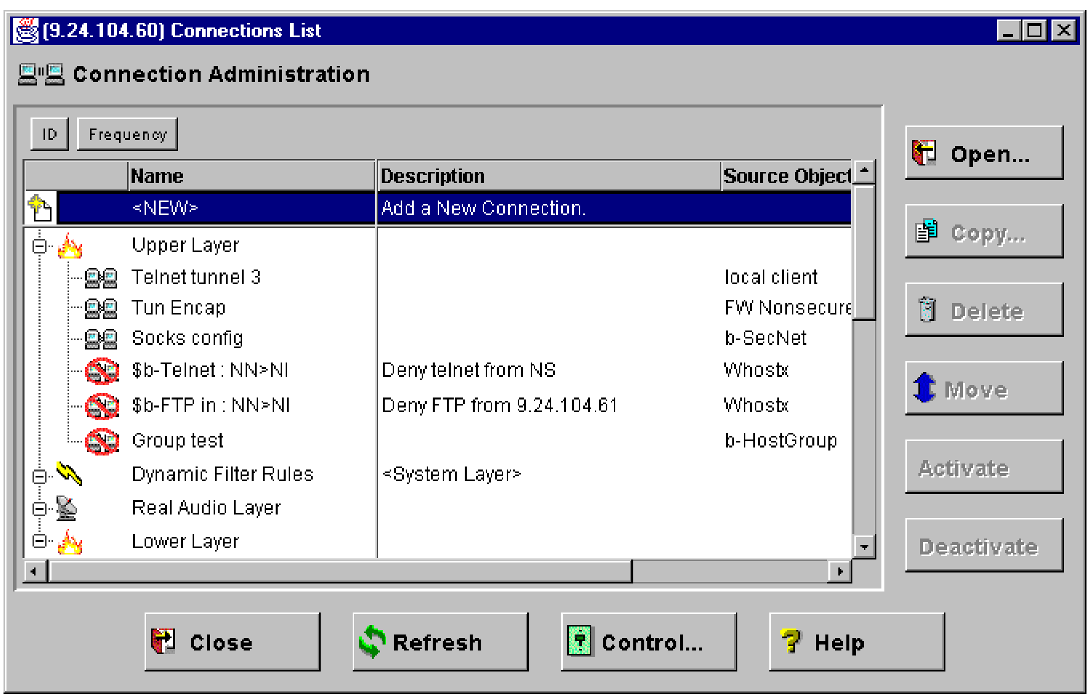

Figure 17. Connection administration window

Figure 18 shows the Add a New Connection window, where you must define all the parameters. First, enter the name and the description of this connection. Remember to use a convention for all the names, as this will make future definitions and modifications easier.

IBM SecureWay Firewall V4.1 for AIX provides a connection match frequency count display. You can see the **Frequency** button in Figure 17., “Connection administration window” on page 53; if you click on this button, a new column will be displayed, where you will see the frequency with which each connection gets a hit (a packet with matches the characteristics described in the connection arrives or leaves the firewall). You can utilize this count values to adjust the order of the connections. Highly used connections should be located above those with low frequency.

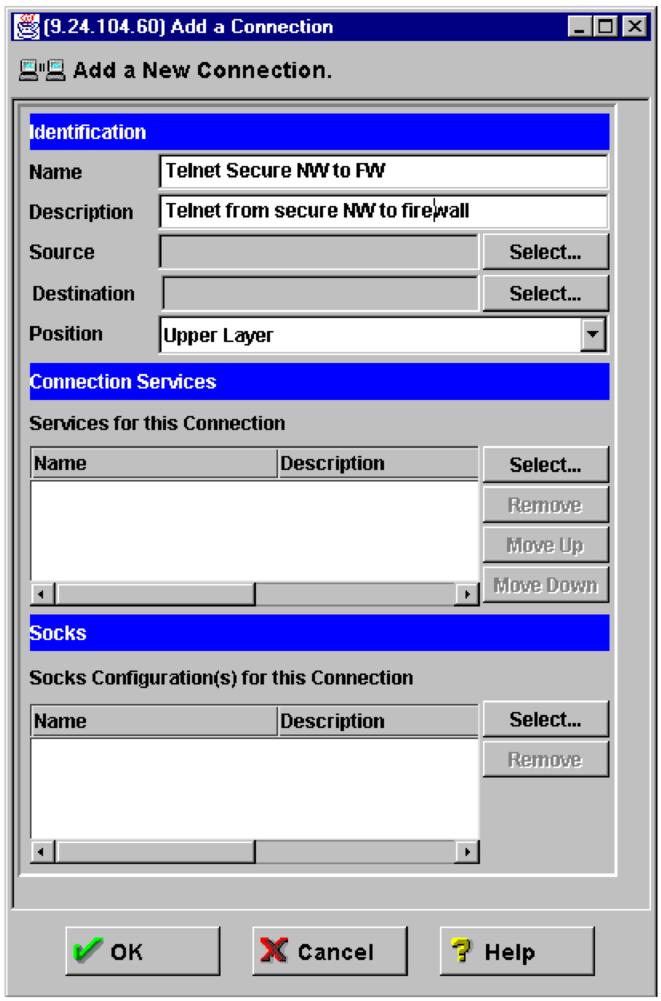

Figure 18. New connection window

Secondly you need to define the source object and destination object. Click **Select** to select each one from the object list. If you have not already defined the object, you can select **New** to define it. The only object predefined is The World, so we will have to define both the source and destination objects to construct our example. Figure 19 shows the definition for our source object, representing any address in the secure network. After you have defined the object click **OK,** then select the new object in the object list and click **OK** to place the object in the source object field. This procedure must be repeated for the destination object, the firewall itself. Figure 20 shows our definition for this object.

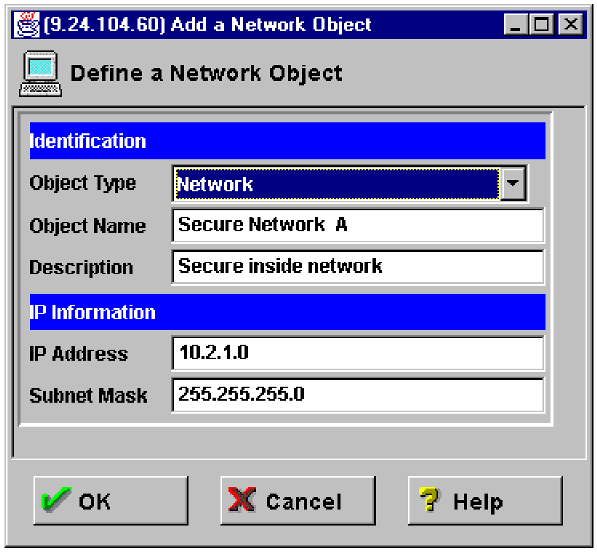

Figure 19. Source object definition window

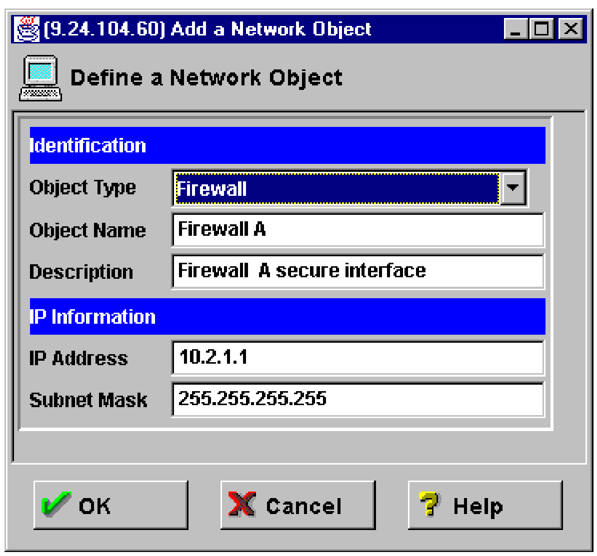

Figure 20. Destination object definition window

Finally you need to select the service between these objects. Click **Select** and a list of all the defined services will appear. In this case we are using a standard service, so select **Permit Proxy Telnet Outbound** and click **OK**. Figure 21 shows the service list.

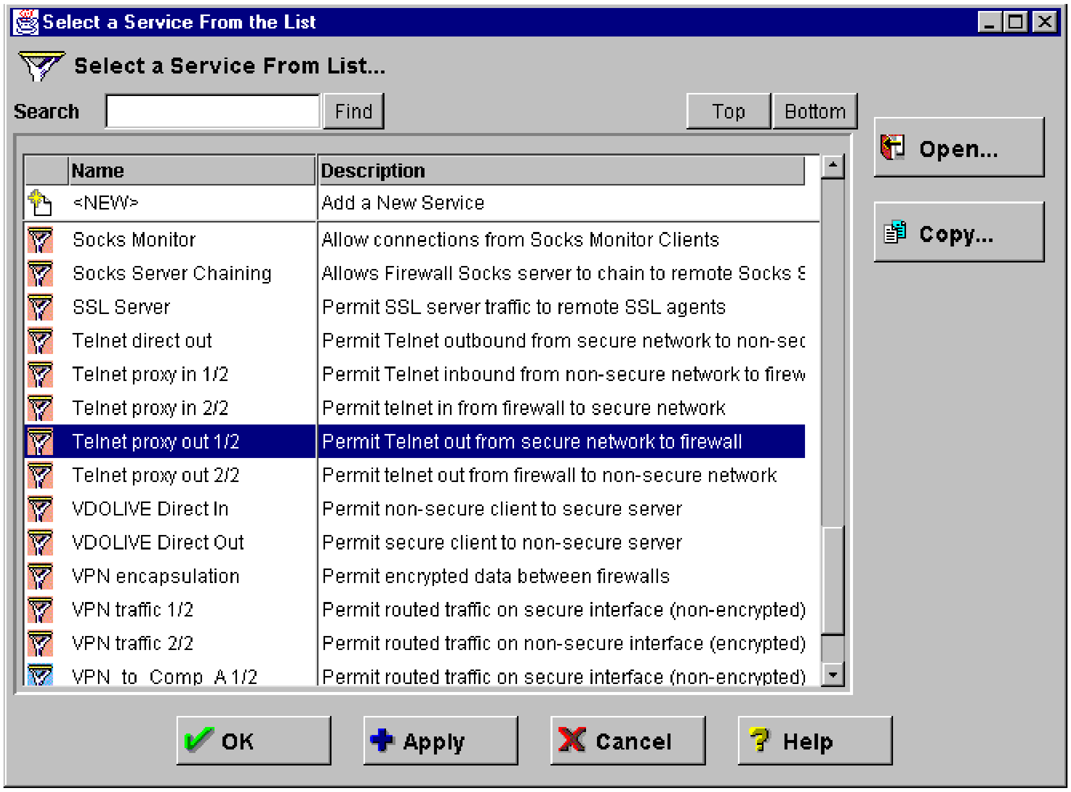

Figure 21. Service selection window

Figure 22 shows the final result. Click on **OK** to save the connection definition.

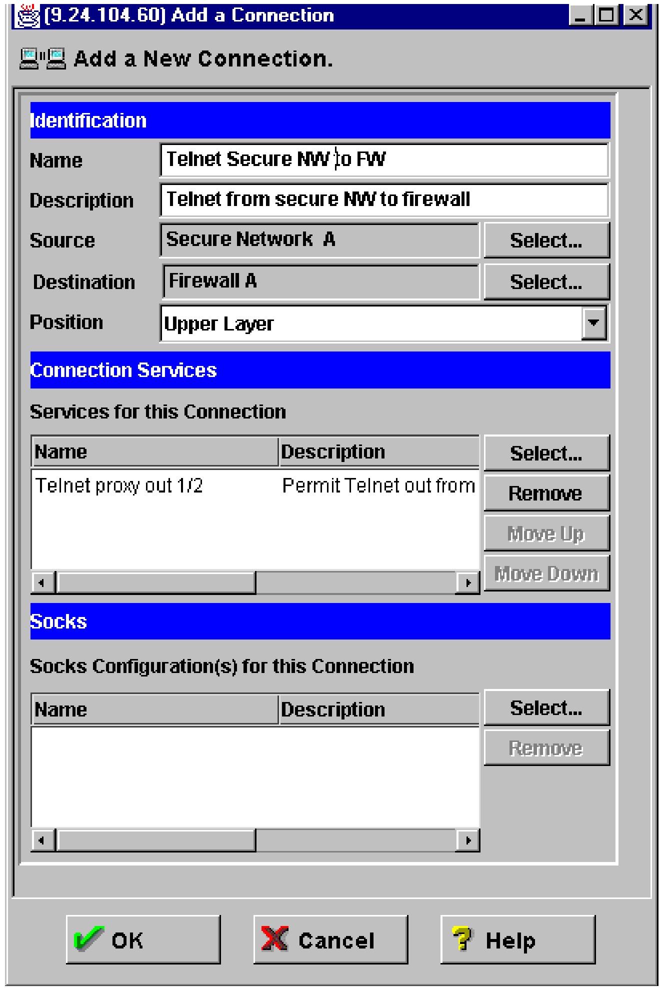

Figure 22. New Connection window, completed

Now you need to activate the rule base and the filter rules will be applied.

4.2.2 Non-standard connections

A non-standard connection is one that cannot be built from predefined services. For example, imagine you have a new application (we call it "CUST") which has a proxy server running on your firewall. It listens on TCP port 400 and you want to be able to access it from the secure network. This is visualized in Figure 23.

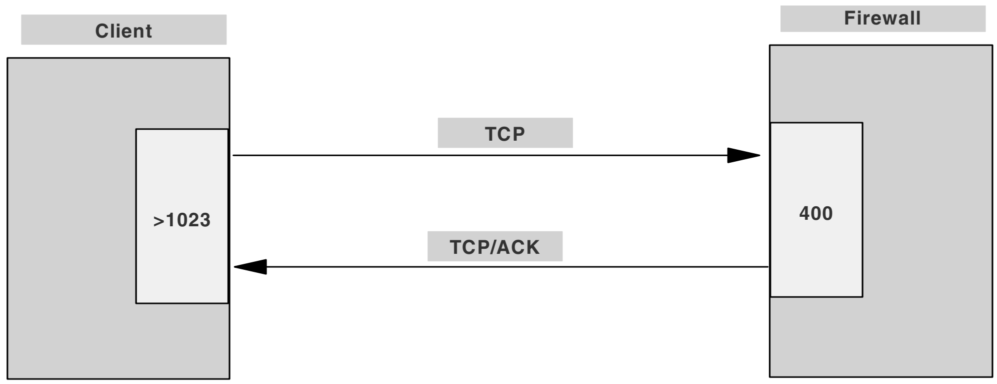

Figure 23. A non-standard connection

To be able to build this connection, we are going to create a service that is called "Permit CUST". First, we have to decide whether we need new rules for this service. Therefore you have to know which rules already exist, by checking the list of rules. Do this by selecting **Traffic Control** then **Connection Templates**, and **Rules** from the initial GUI navigator panel.

In this example we need a rule that permits inbound TCP packets on port 400 of the secure interface. This does not exist, so we must create a new rule.

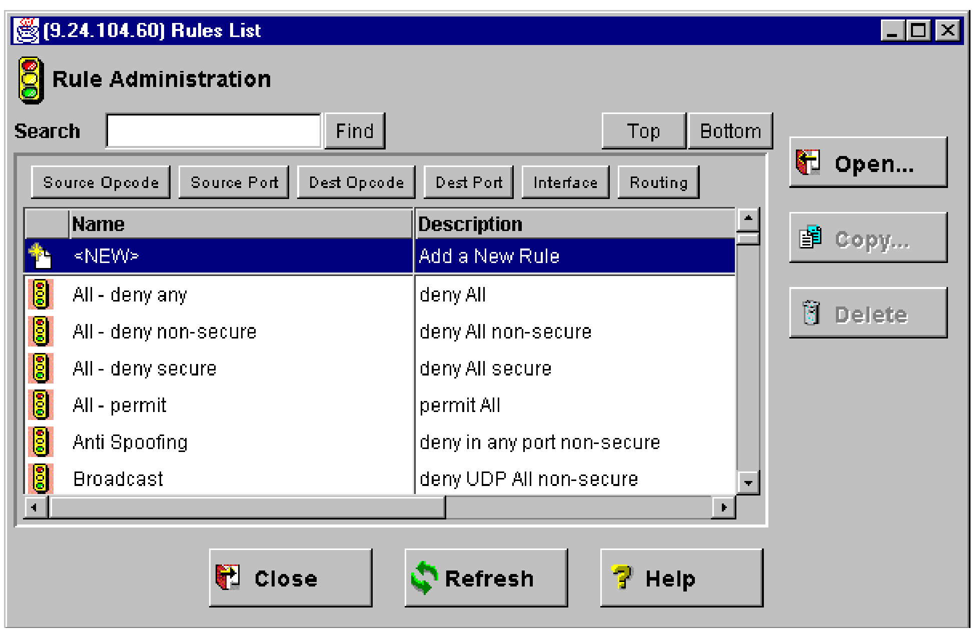

Figure 24. New rule selection window

It is very important to assign a clear name to a rule. For example, do not use the name of a source or destination in the rule name, because they are independent of the rule. A good name may be: "Permit CUST Inbound 1". By giving rules clear names, it is also easier to reuse your rules. In the rule list double-click **<NEW>** as shown in Figure 24. Fill in the parameters for the new rule, as shown in Figure 25. Notice that we have been very specific in defining the rule. It will only allow packets for our CUST application to pass if they appear inbound on the secure side of the firewall.

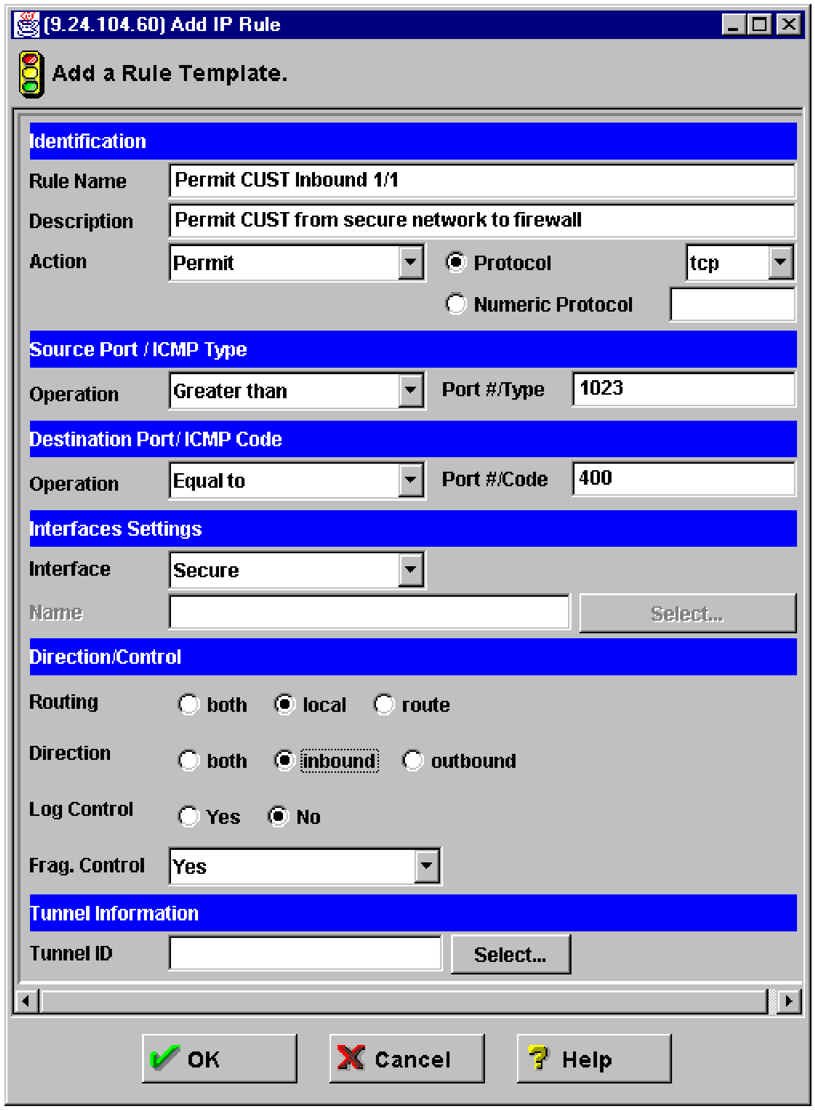

Figure 25. Parameters for "Permit CUST Inbound 1/2"

This rule deals with one direction only, client to server. We also need to create a rule for the response packets from the server to the client. The construction of this rule is visualized in Figure 26 on page 60.

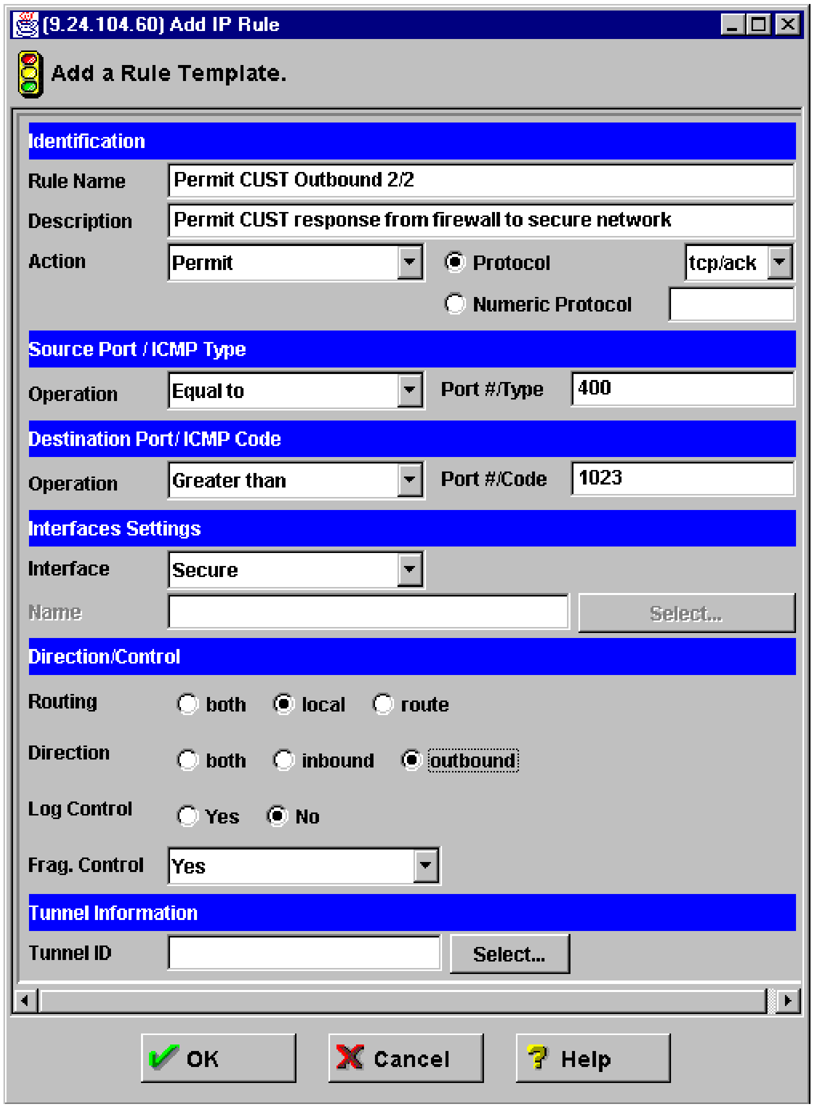

Figure 26. Parameters for "Permit CUST Outbound 2/2"

The differences from the first rule are:

- The protocol is now TCP/ACK.
- Source and destination criteria are swapped.
- The direction is now outbound.

After creating this rule we can build the "Permit CUST" service that invokes the new rules. Select **Traffic Control** then **Connection Templates,** and **Services**. from the initial GUI navigator panel. You will see the list of existing services. In the list double-click **<NEW>** to see the new service dialog window shown in Figure 27.

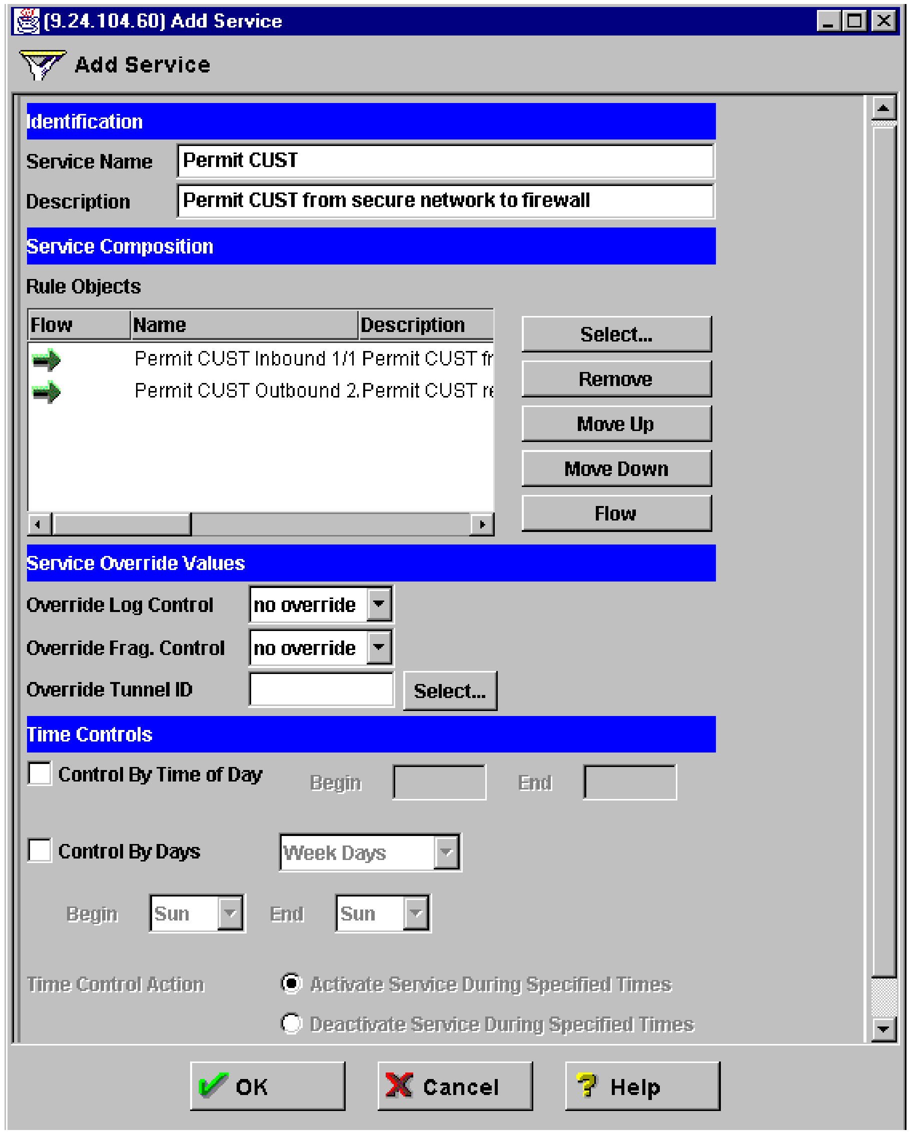

Figure 27. New service "Permit CUST"

Finally, we can configure a connection in the same way as for the previous example with the following content:

- Source object: *Secure Network* (created in the previous example).
- Destination object: *Secure Firewall* (also created in the previous example).
- Service: *Permit CUST*.

After you complete the connection definition, you can see the defined connection on the connection administration window as disabled. You can select the connection you finished and activate it by clicking the **Activate** button. Then this connection becomes immediately effective.

4.2.3 Connection control

After you have defined the connection you have to activate it.

Figure 28 shows the options on the Connection Control window.

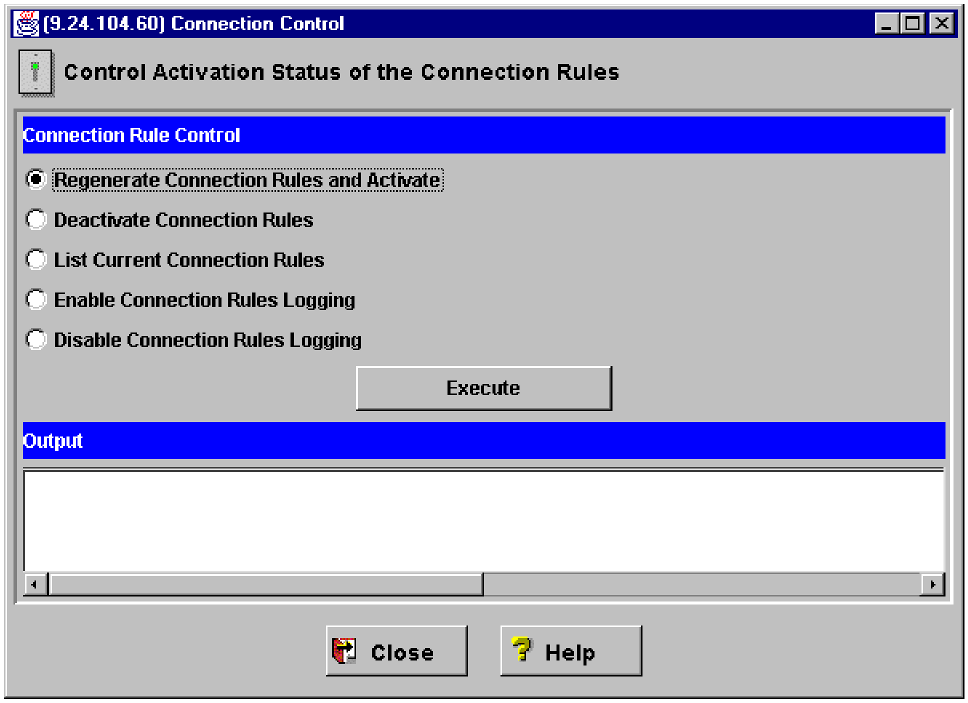

Figure 28. Connection Control window

The functions provided in this screen are as follows;

- Regenerate Connection Rules and Activate: The firewall builds the static filter from the active connection rules. Packet filtering will be done according to the filter rules. To see the static rule interactively the connection should be active.
- Deactivate Connection Rules: If you choose this option, your filter rules will be all deactivated and the firewall is protected only by default rules.
- List Current Connection Rules: If you choose this option you will see the various connections defined hitherto, according to placement hierarchy. Starting from the connection name, you can collapse each icon step by step. By viewing the sequence of connection -> service -> rule -> filter coding, you can verify and debug the steps you have done. Items are treated one by one, interactively. Only the active connections are displayed including setting from the Security Policy window. Double-clicking the **rule symbol (traffic light)** will toggle between *rule coding* and *rule name.*
- **Enable Connection Rules Logging**: Choose this option to enable logging. Firewall logs selected traffic to the log facility.
- **Disable Connection Rules Logging**: Choose this option to disable logging. Firewall stops logging to log facility.

4.3 Filter examples

In this section we suggest some filter definition examples. The network environment and the example filter planning worksheet are shown in Figure 29 and Table 7 respectively.

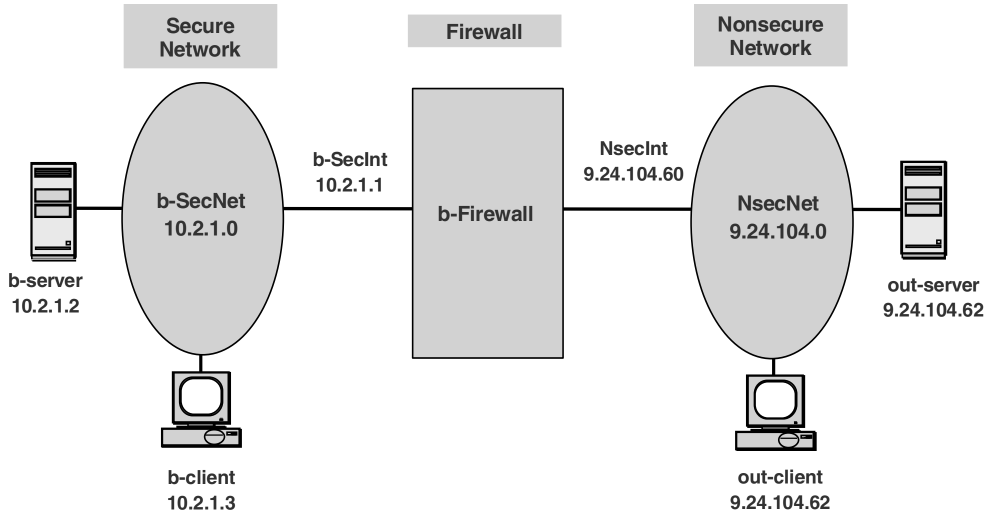

Figure 29. Example network

In the example network, the secure network is 10.2.1.0 and the non-secure network is 9.24.104.0. Actually we used ITSO 9.24.104.0 network as non-secure network. Note that the secure and non-secure definition is relative to the firewall.

We suggest three cases as filter setting examples. They are ICMP, Telnet and FTP.

Table 7. Example filter planning sheet

Protocol	Source		Source IP & Mask	Desti-nation	Dest. IP & Mask	Direction		Access		Use of Log
	Secure	Non-secure				Out	In	D	P	
ICMP	b-SecNet		10.2.1.0 /24	NsecNet	9.24.104.0 /24	o	o	o		no
Telnet	b-client		10.2.1.3 /24	World	any	o			o	yes
		NsecNet	9.24.104.0 /24	b-SecNet	10.2.1.0 /24		o	o		yes
FTP	b-SecNet		10.2.1.0 /24	World	any	o			o	yes
		out-client	9.24.104.62 /24	b-server	10.2.1.2 /24		o		o	yes
		World	any	b-SecNet	10.2.1.0 /24		o	o		yes

4.3.1 ICMP example

The ICMP planning item is in the first row of Table 7. It says that all ICMP traffic is prohibited through the firewall.

The ICMP example protocol model is shown in Figure 30.

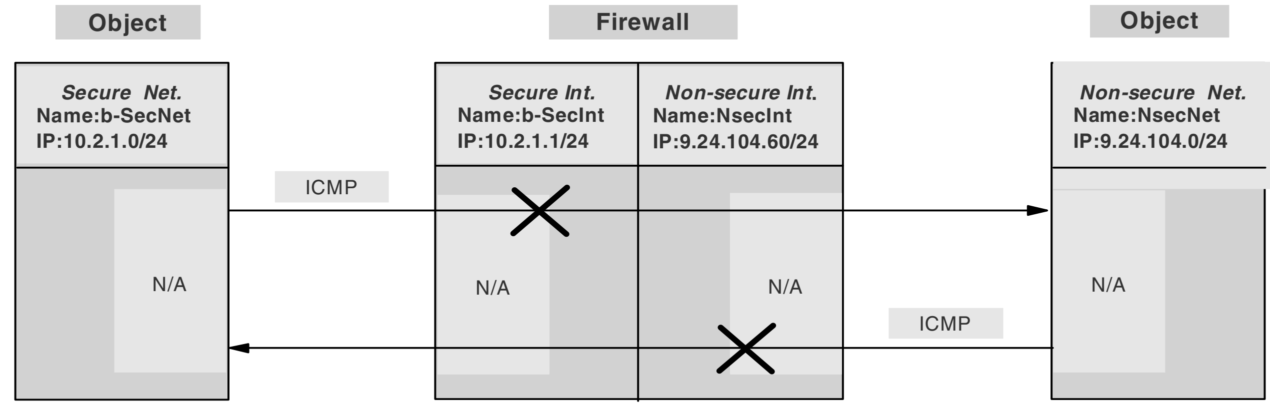

Figure 30. ICMP example protocol model

Now it's time to find out whether predefined services exist that will satisfy this example model. We cannot find predefined rules or services. We need to create new rules and services.

For details of the definition process refer to 4.2, "Definition flow" on page 52.

To manage things better, let's use Table 8 for creating new rules and services.

We used "$b+" or "$b-" as the beginning three characters to gather all the created rules and services in the front part of the list and to clarify that this creation was made for network "b". If you have a single secure network connected to this firewall, you can omit this character. Third position can be "+" or "-", where "+" means permit and "-" means deny. If a service or a connection permits something, you will see "+"; if not you will see "-" in the name string.

You may have your own naming conventions that will fit your real environment best.

Table 8. ICMP example filter definition sheet

Type	Standard Name	New Name	Changes/Remarks
Rule	PING	$b-ICMP deny all	change: *permit PING* to *deny all ICMP*
Service	PING	$b-ICMP deny all	select: $b-ICMP deny all
Connection		$b-ICMP out deny	object: b-SecNet --> b-SecInt
		$b-ICMP in deny	object: NsecNet --> NsecInt

Figure 31 shows the connections of this example. You can verify the full coding in the window.

In this example and others, we put the connections on the Lower Layer to discriminate from other connection definitions. Usually you locate connections on the Upper Layer. Note that connections cannot move across layers.

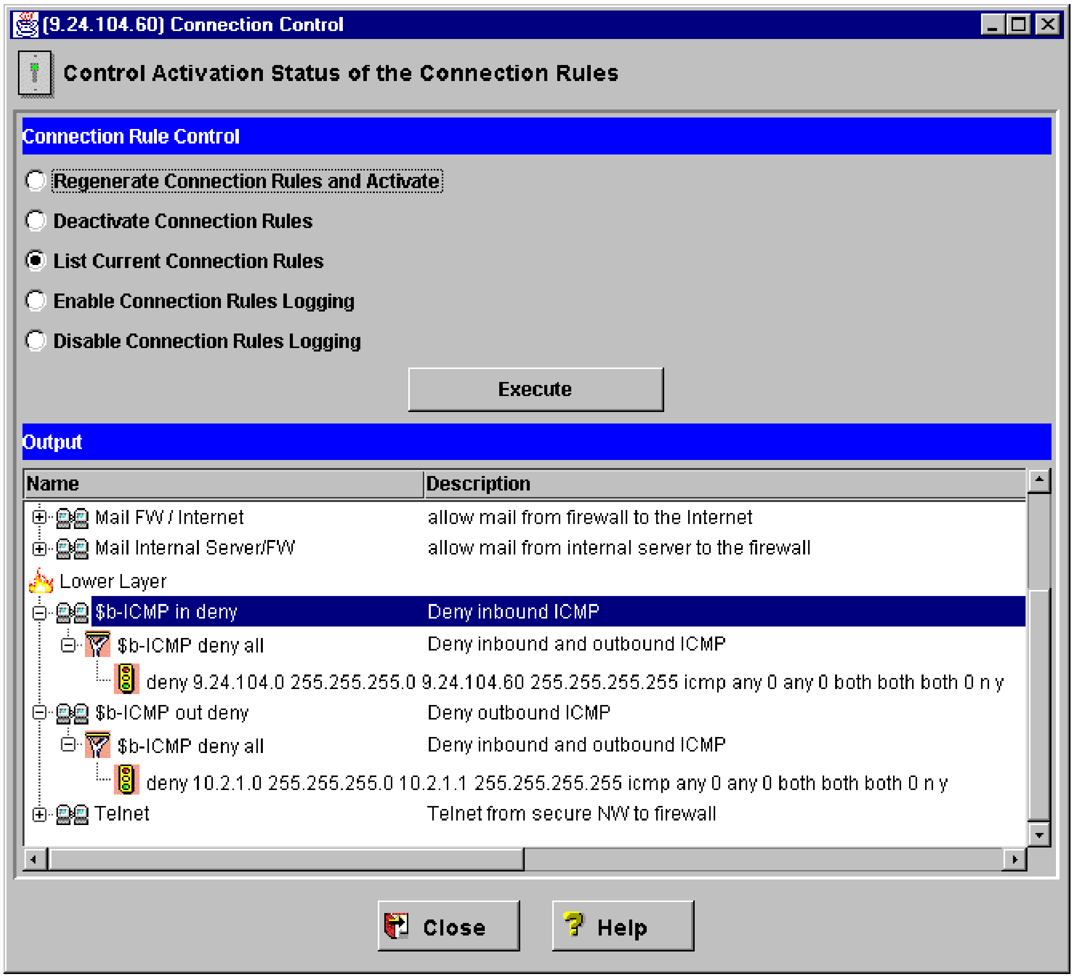

Figure 31. ICMP example definition result

4.3.2 Telnet example

The Telnet planning item is in the second row in Table 7. It says that b-client can telnet any non-secure network, and NsecNet cannot telnet to a secure network. Log control is required in both cases.

The Telnet example protocol model is shown in Figure 32. Note that the inbound Telnet can be prohibited by denying any trial to port 23 of firewall’s non-secure interface from a non-secure network.

Note that we did not change the position of the secure and non-secure objects. Outbound traffic starts from a secure network and inbound traffic starts from a non-secure network.

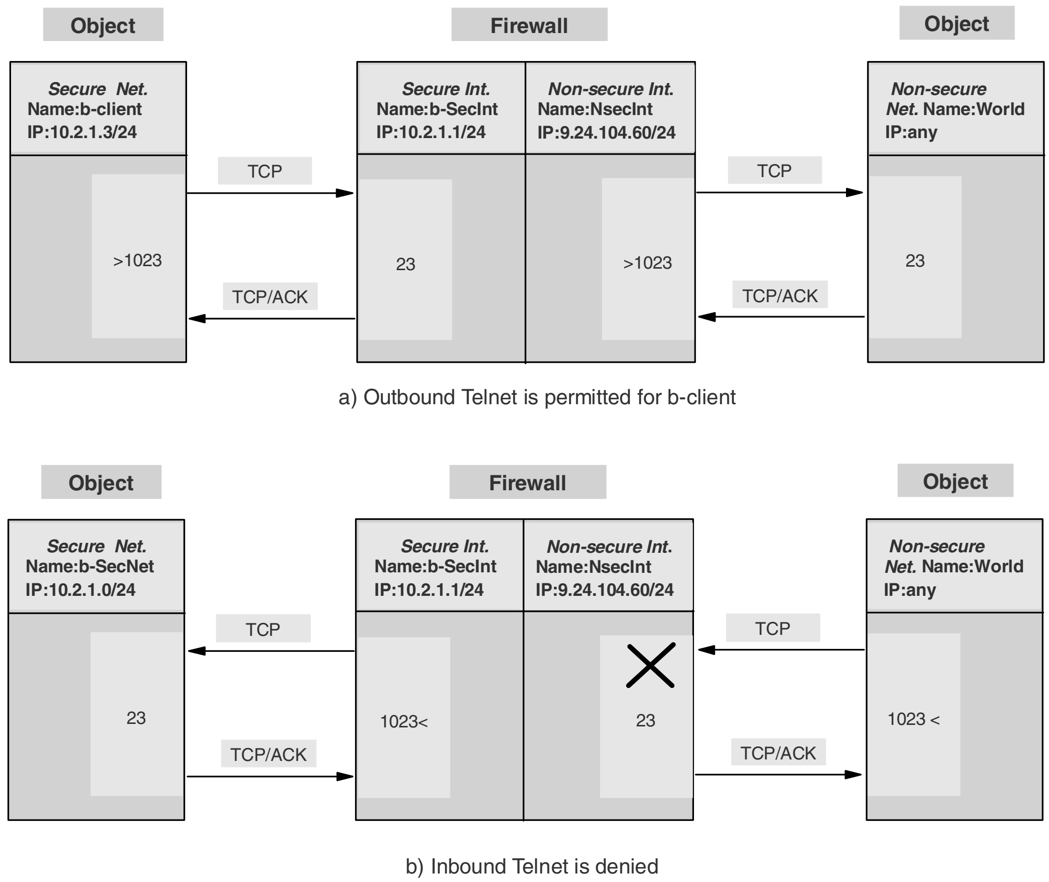

Figure 32. Telnet example protocol model

Now it's time to find out whether predefined services exist that will satisfy this example model. We can find two outbound Telnet permit services from the services list. But we cannot find an inbound Telnet deny service, so we need to create new rules and services for this connection.

For details of this definition process, 4.2, “Definition flow” on page 52.

Let's use Table 9 for creating new rules, services and connections. The naming conventions are the same as in the ICMP example.

Table 9. Telnet example filter definition sheet

Type	Name	New Name	Changes/Remarks
Telnet from b-client to the World			
Rule	Proxy Telnet in secure 1/2		
	Proxy Telnet ACK out secure 2/2		
Service	Telnet proxy out 1/2	$b+Telnet proxy out 1/2	set log=yes
Connection		$b+Telnet: SN>SI	SN=secure network SI=secure interface object: b-client --> b-SecInt
Rule	Proxy Telnet out non-secure 2/2		Note: sequence 2/2 should be 1/2
	Proxy Telnet ACK in non-secure 2/2		
Service	Telnet proxy out 2/2	$b+Telnet proxy out 2/2	set log=yes
Connection		$b+Telnet: NI>NN	NI=non-secure interface NN=non-secure network object: NsecInt --> World
Deny Telnet from the World			
Rule	Proxy Telnet in non-secure 1/2	$b-Proxy Telnet in non-secure 1/1	change: *permit TCP* to *deny TCP*
Service	Telnet proxy in 1/2	$b-Telnet proxy in 1/1	set log=yes
Connection		$b-Telnet: NN>NI	NN=non-secure network NI=non-secure interface object: World --> NsecInt

Figure 33 shows the definitions for this example. You can verify the full coding in the window.

Note that activating the connection for deny (in this example, $`b-Telnet: NN>NI connection`) *while* a client is already in a Telnet session, the client station will hang up. If the connection permits Telnet again (deactivate the deny connection), the session becomes active again. It means that the result of activation and deactivation of a connection is, in some respects, related to current session status.

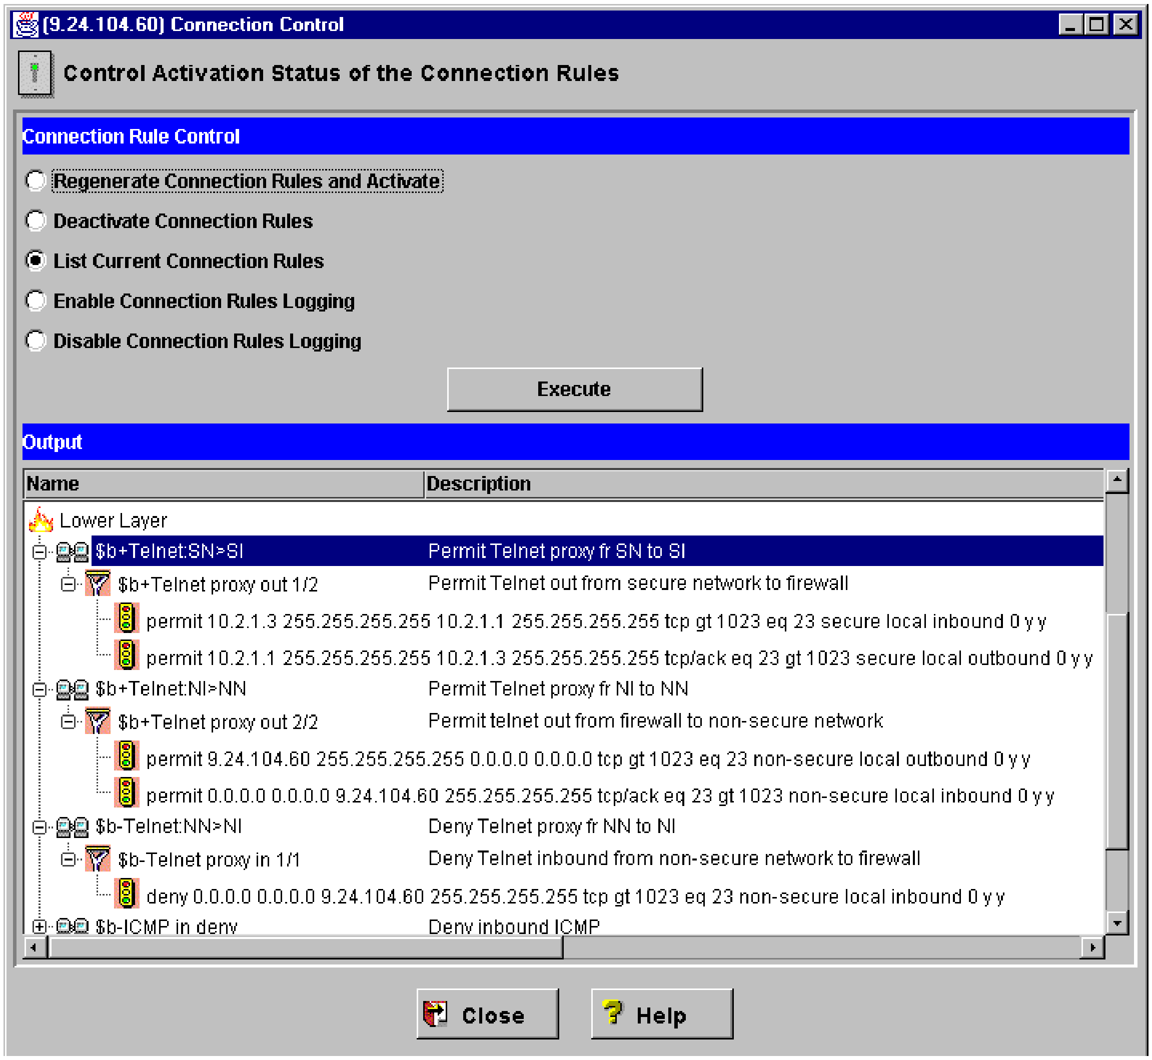

Figure 33. Telnet example definition result

4.3.3 FTP example

The FTP planning item is in the last row in Table 7 on page 63. It says that b-SecNet can ftp to the World and out-client is permitted to ftp to the b-server of a secure network. The rest of the World cannot ftp to SecNet. Log control is required for all cases.

The FTP example protocol model is shown in Figure 34. Note that the inbound FTP can be prohibited by denying any trial to port 21 of the firewall's non-secure interface from a non-secure network. We assume the normal mode FTP.

Note that we did not change the position of secure and non-secure objects. Outbound traffic starts from a secure network and inbound traffic starts from a non-secure network.

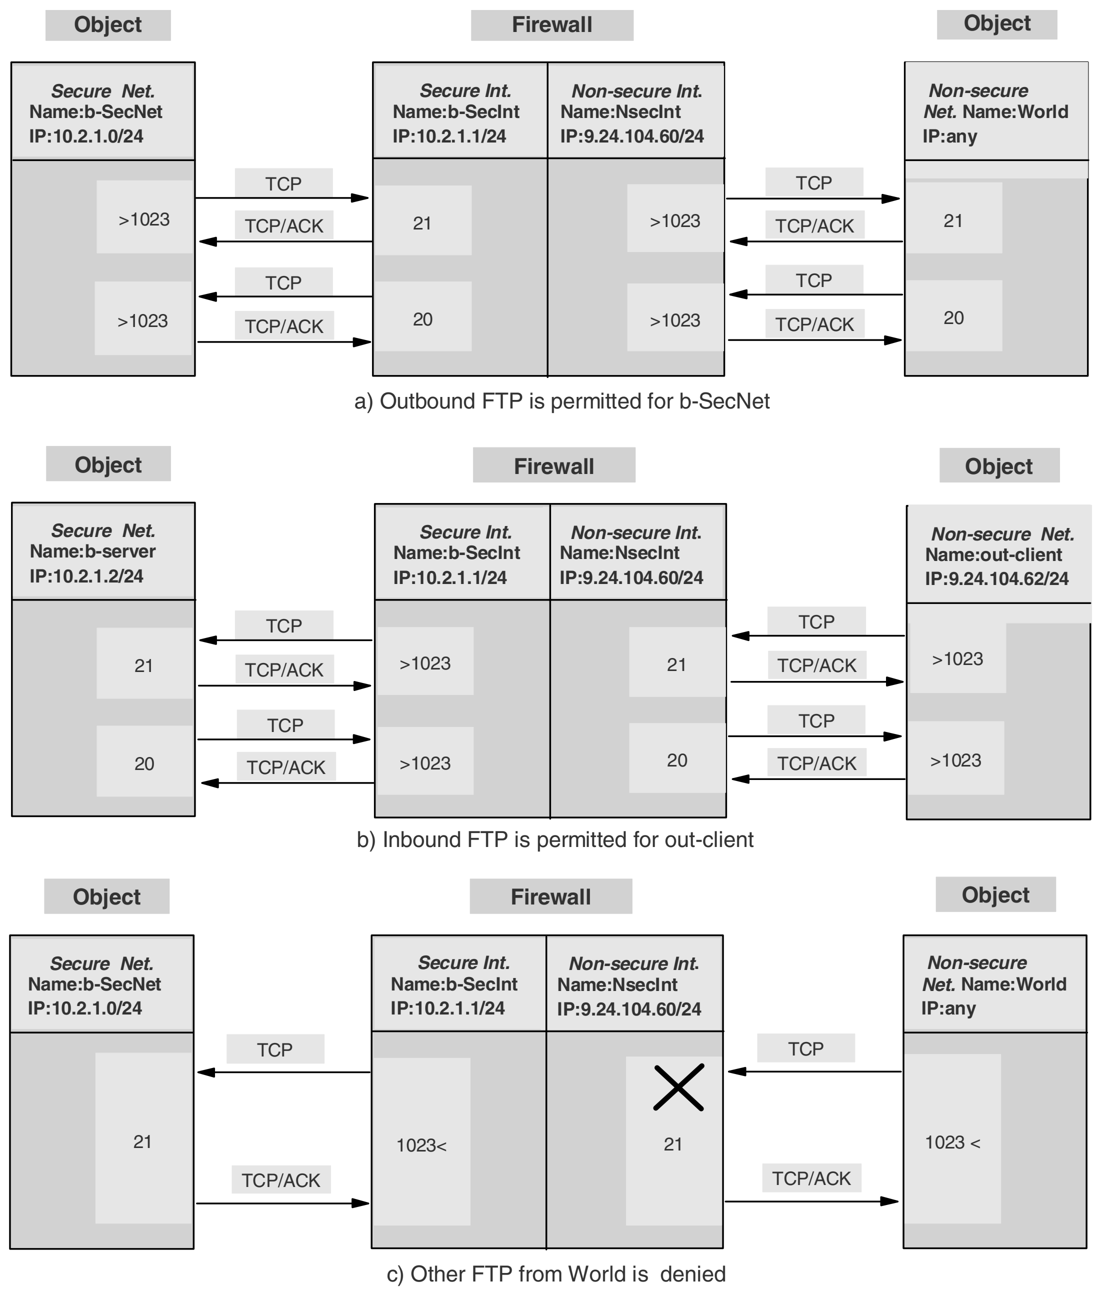

Figure 34. FTP example protocol model

Now it's time to find out whether predefined services exist that will satisfy this example model. We can find two outbound FTP proxy permit services and two inbound FTP proxy permit services from the service list. But we can not find inbound FTP deny service so we need to create new rules and services for this connection.

For details of this definition process, refer to 4.2, "Definition flow" on page 52.

Let's use the Table 10 for creating new rules, services and connections. The naming conventions are same as in the ICMP example.

Table 10. FTP example definition sheet

Type	Standard Name	New Name	Changes/Remarks
FTP from b-SecNet to the World			
Rule	Proxy FTP Control in secure 1/2		
	Proxy FTP Control Ack out secure 2/2		
	Proxy FTP Data out secure 1/2		
	Proxy FTP Data Ack in secure 2/2		
Service	FTP proxy out 1/2	$b+FTP proxy out 1/2	set log=yes, delete 2 FTP passive rules
Connection		$b+FTP out: SN>SI	SN=secure network SI=secure interface object: b_SecNet --> b-SecInt
Rule	Proxy FTP Control out non-secure 1/2		
	Proxy FTP Control Ack in non-secure 2/2		
	Proxy FTP Data in non-secure 1/2		
	Proxy FTP Data Ack out non-secure 2/2		
Service	FTP proxy out 2/2	$b+FTP proxy out 2/2	set log=yes, delete 2 FTP passive rules
Connection		$b+FTP out: NI>NN	NI=non-secure interface NN=non-secure network object: NsecInt --> World
FTP from out-client to b-server			
Rule	Proxy FTP Control in non-secure 1/2		
	Proxy FTP Control Ack out non-secure 2/2		
	Proxy FTP Data out non-secure 1/2		
	Proxy FTP Data Ack in non-secure 2/2		
Service	FTP proxy in 1/2	$b+FTP proxy in 1/2	set log=yes, delete 2 FTP passive rules
Connection		$b+FTP in: NN>NI	NN=non-secure network NI=non-secure interface object: out-client --> NsecInt

Type	Standard Name	New Name	Changes/Remarks
Rule	Proxy FTP Control out secure 1/2		
	Proxy FTP Control Ack in secure 2/2		
	Proxy FTP Data in secure 1/2		
	Proxy FTP Data Ack out secure 2/2		
Service	FTP proxy in 2/2	$b+FTP proxy in 2/2	set log=yes, delete 2 FTP passive rules
Connection		$b+FTP in: SI>SN	SN=secure network SI=secure interface object: b-SecInt --> b-server
Deny FTP from the World			
Rule	Proxy FTP Control in non-secure 1/2	$b-Proxy FTP Control in non-secure 1/1	change: *permit TCP* to *deny TCP*
Service	FTP proxy in 1/2	$b-FTP proxy in 1/1	set log=yes select $b-Proxy FTP Control in non-secure 1/1
Connection		$b-FTP in: NN>NI	NN=non-secure network NI=non-secure interface object: World --> NsecInt

Figure 35 and Figure 36 show the definitions for this example. You can verify the full coding in the window.

Note that the order of permissions for out-client and deny from the World. The permit connection must be located before the deny connection. If the order is changed *out-client* will not be able to ftp to *b-server*.

Note that activating the connection for deny (in this example, $b-FTP in: NN>NI connection) *while* a client is already in FTP session, the client station will hang up. If the connection permits FTP again (deactivate the deny connection), the session becomes active again. It means that the result of activation and deactivation of a connection is, in some respects, related to current session status.

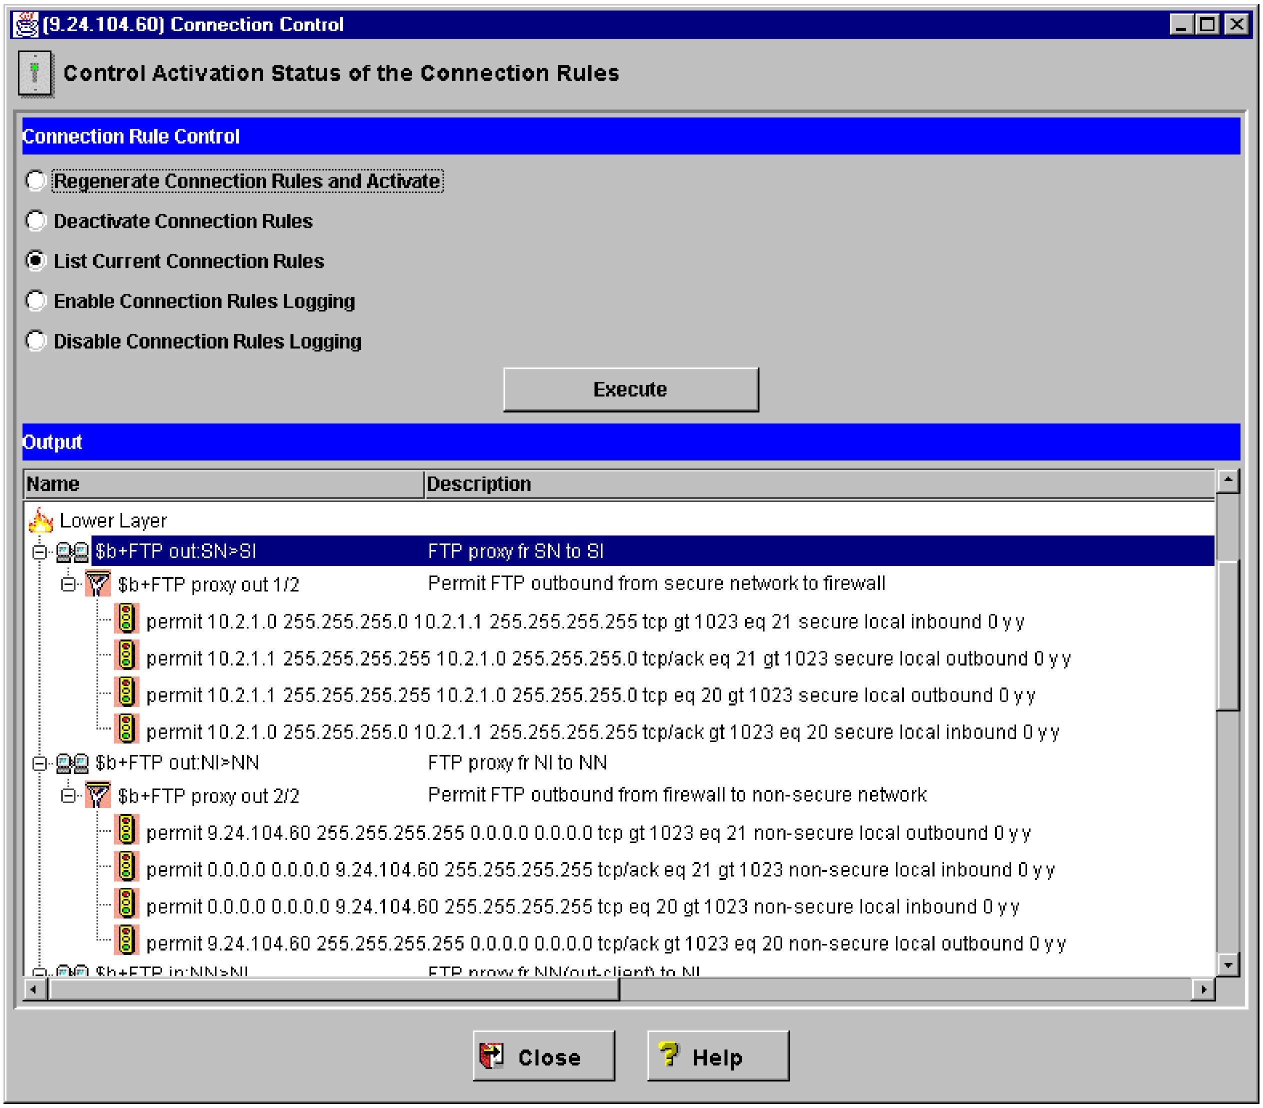

Figure 35. FTP example definition result: outbound

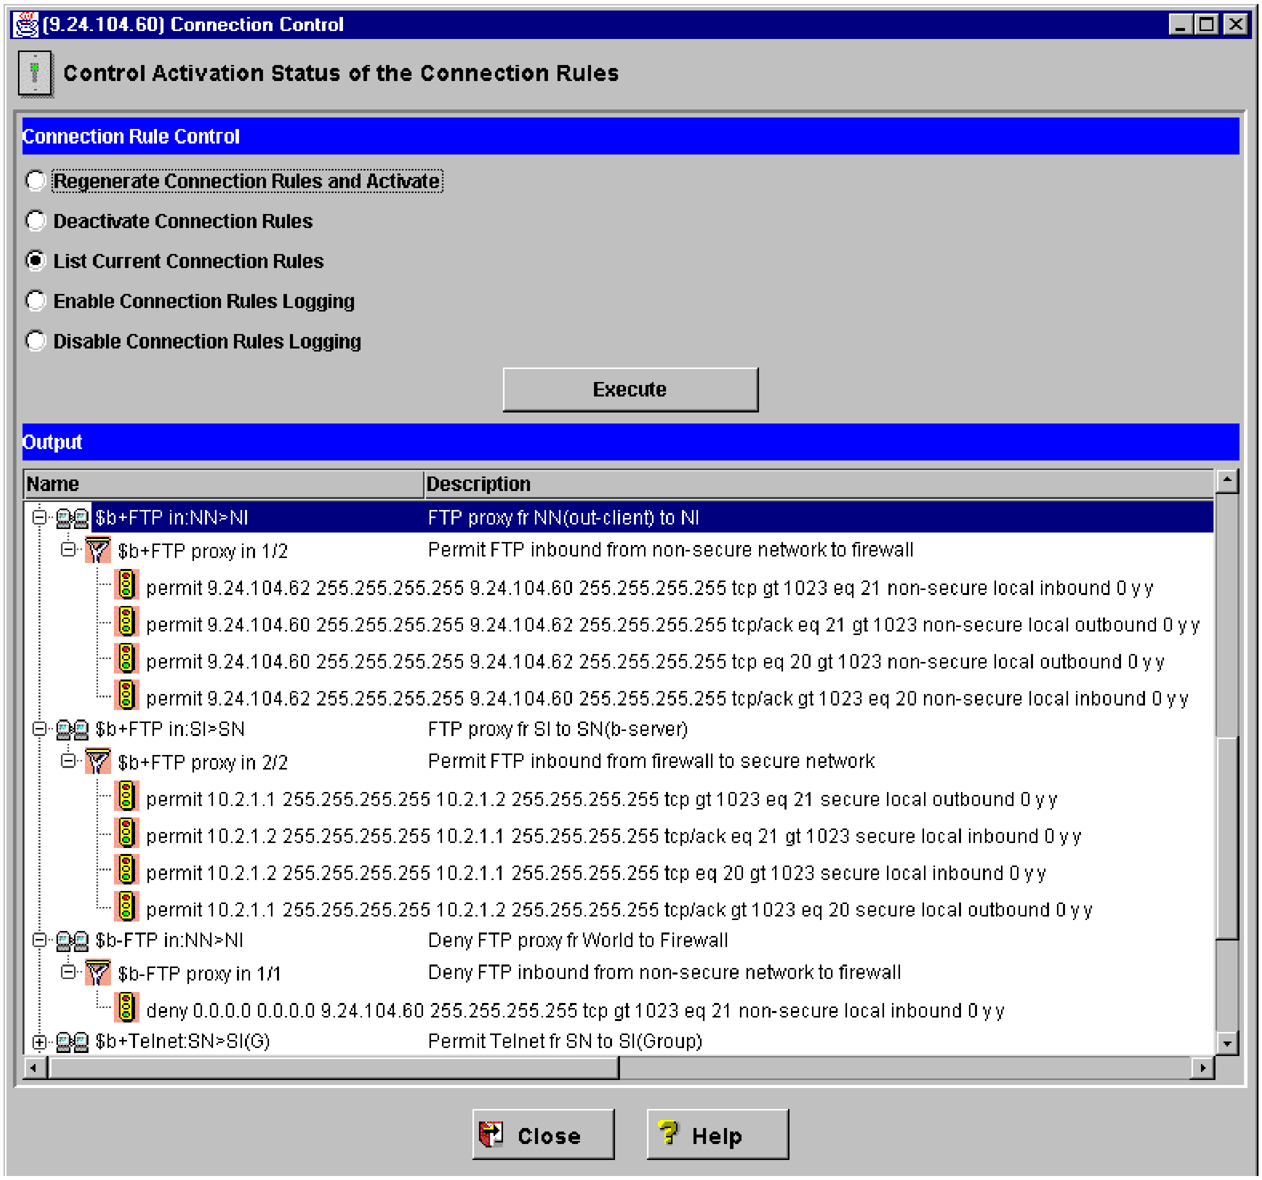

Figure 36. FTP example definition result: inbound

4.4 Filter rules samples

You will find many samples of how to configure the filter rules in the redbook *Guarding the Gates Using the IBM eNetwork Firewall V3.3 for Windows NT,* SG24-5209, in Chapter 9, "Examples of rules for specific services."

4.5 ICMP traffic and MTUs

Many firewall installations do not allow ICMP traffic through the firewall. This is a safe approach, but you should be aware that you may encounter problems if the packets sent to the Internet from the firewall get to a host that cannot handle the size of those packets (they are too large) and cannot be fragmented.

The Maximum Transmission Unit (MTU) for an interface is the maximum datagram size it can handle. When one host sends data to another host it is preferable that the datagrams have the largest size that does not require fragmentation anywhere along the path from the source to the destination. This datagram size is referred to as the Path MTU (PMTU).

At the startup of a TCP/IP connection only the MTU values between endpoints are usually considered. So, it is possible that a packet arrives at an intermediate host that has a smaller MTU. This is handled by fragmenting or sending ICMP type 3 code 4 packets when fragmentation is not allowed.

However, if you do not want to allow ICMP traffic, the firewall must have the minimum MTU on the path. Also, all hosts and routers directly connected to the non-secure interface must have this MTU.

The default MTU for the IBM SecureWay Firewall V4.1 for AIX is based on the network topology. The following table shows the maximum MTU sizes for different media:

Table 11. MTU sizes

Network	MTU (Bytes)
16 Mbps Token-Ring	17914
4 Mbps Token-Ring	4464
FDDI	4352
Ethernet	1500
IEEE 802.3/802.2	1492
X.25	576

You cannot use the above values if you are connected to the Internet. For a connection to the Internet we recommend 1440 bytes, which is IEEE 802.3 minus 52 bytes for the packet headers. This should give you a reasonable throughput without having problems passing these packets to most of the other media that use a larger MTU. In AIX, follow these steps to set the MTU value for an interface:

1. Type the following command:

```
# smit chif
```

2. Select the interface you want to change.
3. In the next screen, the cursor is positioned on the field *Maximum IP PACKET SIZE for THIS DEVICE*, and it shows the current value. Type the new value and press Enter. See the following screen:

```
 Network Interface Drivers

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                        [Entry Fields]
  Network Interface                                   tr0
  Maximum IP PACKET SIZE for THIS DEVICE             [1492]                   +#









Esc+1=Help          Esc+2=Refresh       Esc+3=Cancel        Esc+4=Liste
Esc+5=Reset         Esc+6=Command       Esc+7=Edit          Esc+8=Image
Esc+9=Shell         Esc+0=Exit          Enter=Do
```

Figure 37. Changing the MTU size for a device

Chapter 5. Domain Name System (DNS) Service

The DNS configuration in the firewall provides name resolution services to hosts in secure network while keeping internal hosts information secret to outside networks.

The DNS server on the firewall behaves in three different ways, according to the source of the query (the location of the machine that is requesting a name resolution).

1. From the outside (Internet)
   - The firewall behaves like a *real wall.* Outsiders cannot see through the wall.
   - When an outsider asks some internal name resolution information to the DNS server on the firewall, the reply is *no* answer or a *limited* answer.
2. From the inside (secure network)
   - The firewall behaves like a transparent glass panel. Insiders can see through the glass panel.
   - When an insider asks the outside name resolution information to firewall DNS function, it will fetch the required information and deliver it to the requester.
3. In the firewall itself
   - The firewall itself is something like a broker or gateway. It blocks inbound queries and forwards outbound queries.
   - Sometimes the firewall has its own resolution information and provides services directly to the requester. The information can be *cached* information or normal DNS *database* information, depending on your network design.

5.1 DNS basics

Let's look more closely at the basic firewall DNS functions. You will see three different cases on how the DNS queries are resolved by the firewall DNS server and how to define DNS servers in a firewall environment.

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

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

5.1.1 DNS flow

The firewall is usually configured to act as a gateway between the DNS server in the secure network and the DNS server in the non-secure network. The firewall may *cache* name resolution information for quick reference, thus improving performance. If there is no DNS server in the secure network or in the non-secure network, you can configure the firewall to answer queries for those domains (we strongly recommend you take some time and configure separate internal and external DNS servers). In this case the firewall must endure the burden of DNS queries with the risk of exposing internal addresses.

Figure 38 shows the typical DNS name resolution flow through the firewall.

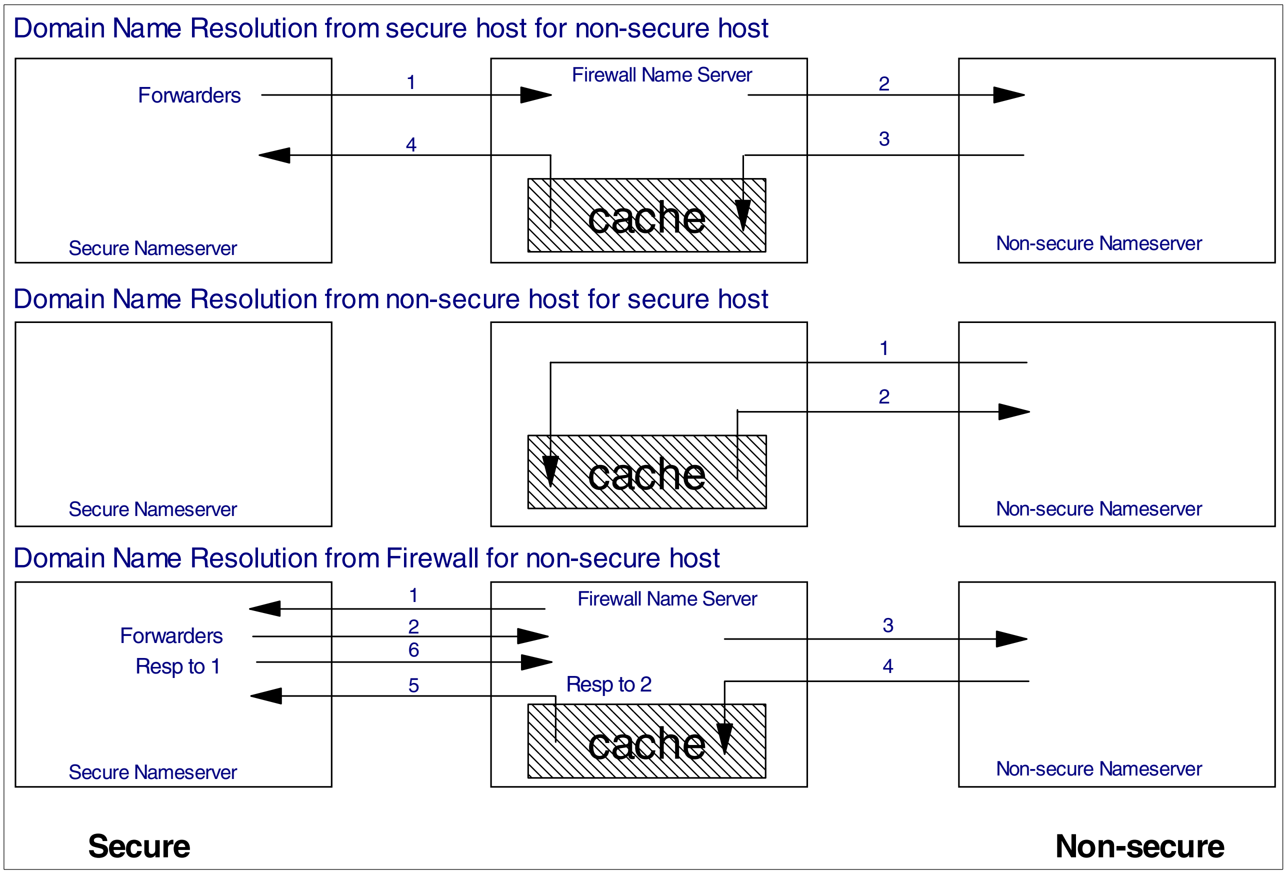

Figure 38. Firewall DNS name resolution flows

Notice that only the names and addresses that we want to reveal (and have defined in the firewall name server) are available to an external host. Notice also that name requests originating on the firewall itself are treated as those from any secure network host, since /etc/resolv.conf refers to the secure network name server.

5.1.2 Configuring the firewall DNS server

First we describe the standard configuration definition for the Domain Name System on the firewall using the GUI. Then we suggest a typical DNS configuration example.

The initial configuration is simple, you just have to open the Configuration Client, and double-click **System Administrator** -> **Domain Name Services**. Figure 39 on page 79 shows the window. After clicking **OK**, the nameserver configuration files are created and the named daemon is started automatically.

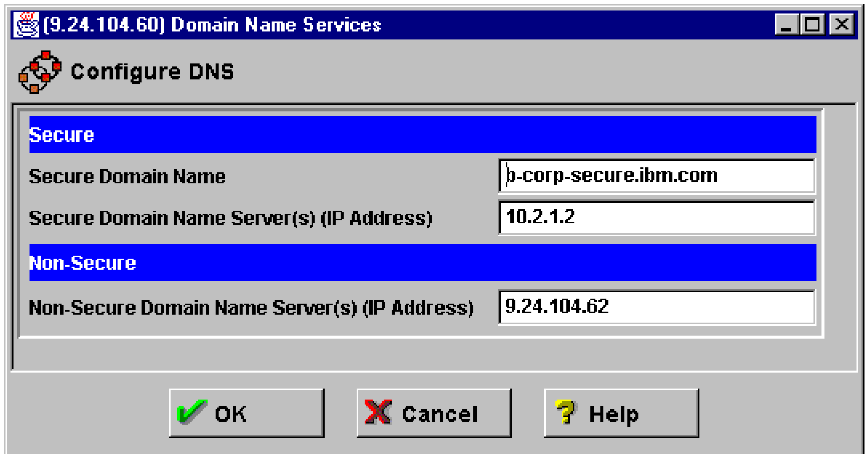

Figure 39. Domain nameserver configuration window

The following nameservers have to be defined:

- The secure domain nameserver is the *SOA* (start of authority) for the internal domain(s), including for the reverse queries (IP address to hostname resolution). It will receive the queries for all machines in the secure network, and if it does not have the answer it will forward the query to the firewall. This server also answers client queries from the firewall itself.
- The non-secure nameserver is the SOA for the non-secure domain. The secure and non-secure domain name may be the same, but this nameserver will have only the external addresses (including the MX records for the mail domain) in its database. This nameserver will also receive the query forwarded by the firewall, and it need to have a cache configuration with all root nameservers[1], or it needs to forward these queries to a higher-level DNS server (for example, a server in the ISP).

Forwarding queries to external servers

In the tests in our lab, we noticed that the firewall has a hardcoded instruction to forward its server queries directly to the root nameservers. This works faster than forwarding them to the non-secure DNS server, but you may experience some problems if your firewall is not directly connected to the Internet (for example, if you are using a lab environment for tests, and the firewall cannot reach the Internet). In this case, you need to add a forwarders instance to the file /etc/fwnamed.boot, pointing to your non-secure DNS server. This will prevent the firewall from forwarding the queries to the root nameservers, and it will force the firewall to forward them to the non-secure nameserver.

You can download an updated list of the root servers from the following URL:

`ftp://ftp.rs.internic.net/domain/named.root`

In the configuration in Figure 39, you also have to provide the domain name of the secure side of the firewall.

[1] Root nameservers are DNS servers that are the SOA for the top-level domain "." (dot). No matter which domain you try to resolve, you can send the query to a root nameserver and it will send you the address of the next nameserver to which you should forward the query.

You do not need to specify the non-secure domain name here, but consider that the non-secure name must be authorized by the Internet Assigned Numbers Authority (IANA) and will follow the national and international conventions for IP domain names. This is the name that you will be known by to the rest of the world. If your firewall configuration includes a DMZ, the servers within the DMZ will be in this domain.

The best practice is to strictly follow the hierarchical domain naming standards, which is the way the DNS was designed to work. DNS will work even if you use a nonhierarchical scheme, but we do not recommend it. Try to use names that will discriminate internal or external domains only by name. That will help you identify whether a resource is inside or outside the firewall and make it easier to create DNS configurations and mail routing rules.

5.2 DNS configuration example

The DNS configuration example we suggest here is based on our lab configuration as shown in Figure 40 [2]

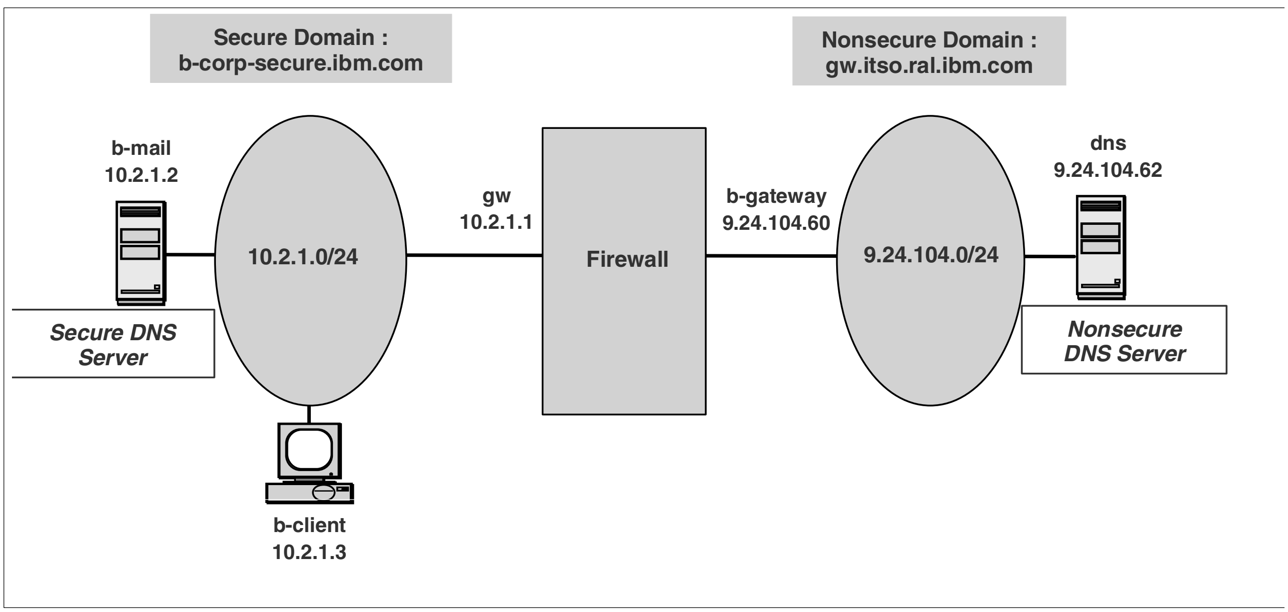

Figure 40. Lab DNS configuration

5.2.1 Firewall DNS configuration

The following files are created by the firewall when you configure DNS as shown in Figure 39.

- /etc/resolv.conf
- /etc/fwnamed.boot
- /etc/fwnamed.loc
- /etc/fwnamed.ca

[2] The convention for subnets used in this figure, 10.2.1.0/24 is the same as 10.2.1.0 mask 255.255.255.0. The number 24 refers to the number of bits ON from the leftmost bit of the mask.

The /etc/resolv.conf file points to the internal nameserver. In our example, the IP address is 10.2.1.2. The contents of /etc/resolv.conf file is as follows:

```
domain     b-corp-secure.ibm.com
nameserver 10.2.1.2
```

This means that when the firewall machine tries to resolve a name, it behaves exactly like a host in the secure network.

The /etc/fwnamed.boot file is the base file from which the DNS configuration is defined. In this case it just specifies the root server hints file /etc/fwnamed.ca, and the loopback/localhost reverse address file. The contents of /etc/fwnamed.boot file is as follows:

```
; Created by IBM Firewall 1999242232
forwarders 9.24.104.62
cache      .                    /etc/fwnamed.ca
primary    0.0.127.in-addr.arpa /etc/fwnamed.loc
```

We used the forwarders directive to send all unresolved DNS requests to the external DNS server of our Lab environment. As we mentioned before, the firewall will try to forward all queries directly to the root nameserver if we do not add this line, and since we are not directly connected to the Internet, we need it.

The /etc/fwnamed.loc file just contains the reverse resolution information for the loopback/localhost address 127.0.0.1. The contents of /etc/fwnamed.loc file is as follows:

```
; Created by IBM Firewall 1999242232
@  IN SOA b-gateway.b-corp-secure.ibm.com. root.b-gateway.b-corp-secure.ibm.com.
(1999242232 3600 600 3600000 86400 )
 IN NS b-gateway.b-corp-secure.ibm.com.
1 IN PTR localhost.
```

The cache hints file /etc/fwnamed.ca specifies the nameserver(s) used to request the list of root nameservers; in this case the external nameserver. The DNS on the firewall will ask that nameserver for the current list of root nameservers, and will cache it in memory. It will repeat this process when the cached list time-to-live expires. In our case this is the DNS system in the non-secure network, 9.24.104.62. The contents of /etc/fwnamed.ca file are as follows:

```
; Created by IBM Firewall 1999242232
.  3600000 IN NS  externaldns.9.24.104.62.
externaldns.9.24.104.62.  3600000 IN A 9.24.104.62
```

Take care if you manually stop and start the firewall DNS server: `stopsrc -s named` will stop the service, but to restart you must specify the configuration file with the command `startsrc -s named -a "-b /etc/fwnamed.boot`. Otherwise it will attempt to use the default file /etc/named.boot.

5.2.2 External DNS server configuration

Unlike the firewall system, name resolution requests on the external DNS server go to the nameserver on the same system. The IP address of this machine is 9.24.104.62, and we are using this same IP in the domain nameserver list (inside TCP/IP configuration).

We used a Windows NT machine as the external nameserver. We used the DNS Manager to configure the database, and this tool generates the files in the directory c:\WINNT\system32\DNS. The DNS Manager window is shown in Figure 41.

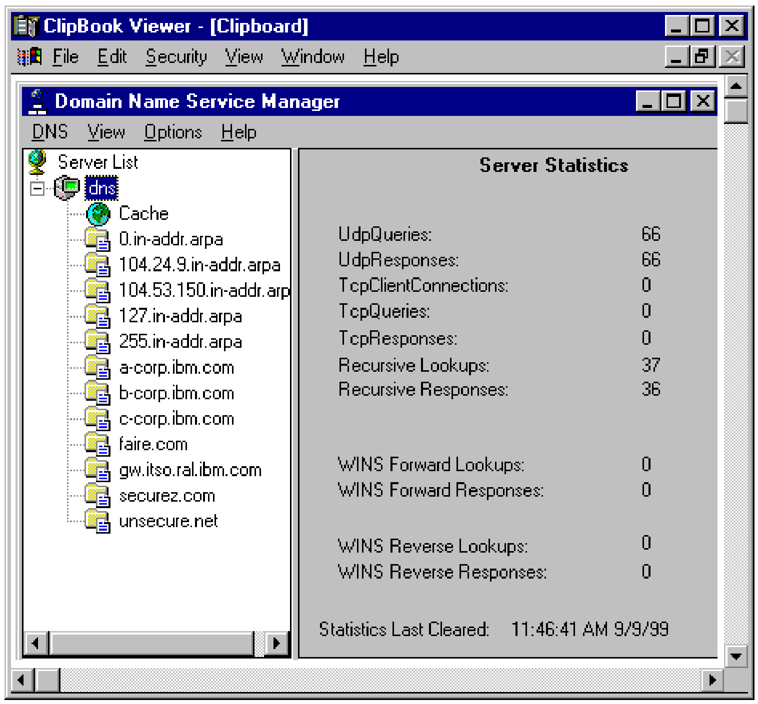

Figure 41. DNS Manager window of Windows NT nameserver

In this DNS server, we added two zones: the main forward zone b-corp.ibm.com (used to resolve hostnames into IP addresses) and the reverse zone 104.24.9.in-addr.arpa (used to resolve IP addresses into hostnames). We also need to add another reverse zone 0.0.127.in-addr.arpa, to allow reverse resolution for the loopback address (127.0.0.1). Finally, we need the cache configuration, to enable this nameserver to redirect all queries that it cannot resolve itself to the root nameservers. In the Windows NT DNS server, the cache configuration is added automatically.

MX record

You must add the MX record in the data file for all mail domains that you have configured in the external DNS server.

We will show some definitions related to our DNS example.

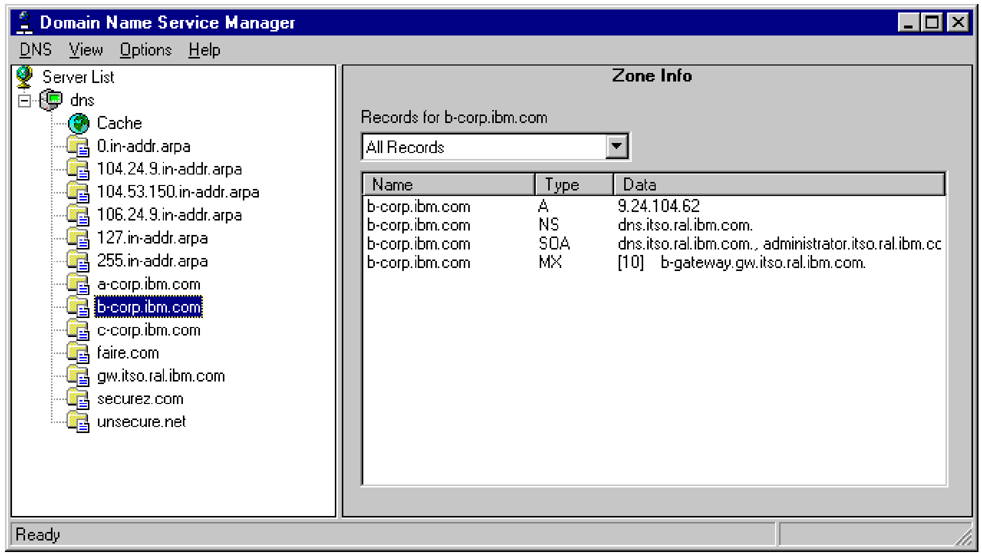

Figure 42. b-corp.ibm.com forward DNS definition

And the reverse DNS definition:

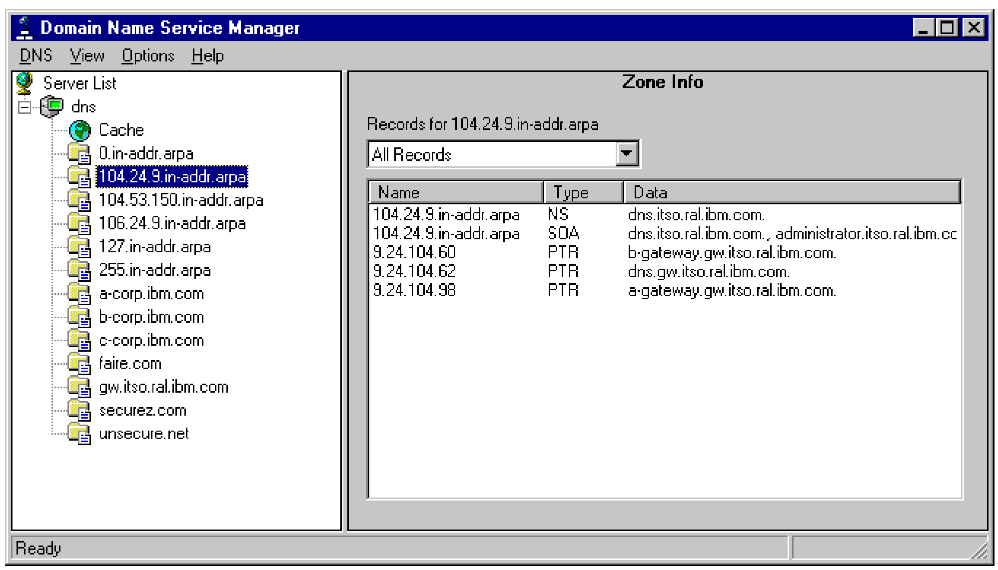

Figure 43. Reverse definition in external DNS

5.2.3 Internal DNS server configuration

We used AIX as the internal DNS server. Since both AIX and NT are using BIND V4 you will notice that the configuration files are similar (AIX also provides BIND V8, but we chose to do the configuration using BIND V4).

The contents of /etc/resolv.conf file is as follows:

```
domain          b-corp-secure.ibm.com
nameserver      10.2.1.2
```

If you let the /etc/resolv.conf file blank it will automatically point to itself as the DNS nameserver, but in this case you need to set the hostname using the fully qualified domain name.

The /etc/named.boot file is also conventional, except that it contains a forwarders record pointing to the firewall. This means that any request for addresses outside its own domain will be forwarded to the DNS server on the firewall. We also used the cache directive to keep the resolved information on the DNS cache for reuse. The contents of /etc/named.boot file is as follows:

```
directory    /etc/dns
forwarders   10.2.1.1
cache        .                         named.ca
primary      b-corp-secure.ibm.com     named.data
primary      b-corp.ibm.com            b-corp.data
primary      c-corp.ibm.com            c-corp.data
primary      1.2.10.in-addr.arpa       named.rev
primary      0.0.127.in-addr.arpa      named.local
```

Note that we also added the domains b-corp.ibm.com and c-corp.ibm.com. We added these domains to be able to get the correct MX record (this configuration is explained in 8.3.5.5, “Overflow server on the firewall” on page 208).

If you want to make this server a *forward-only* server, you can add one line as in the following example:

```
forwarders      10.2.1.1
options         forward-only
```

The forward-only option substatement specifies that the nameserver completely relies on the forwarders servers. Without this statement, the internal DNS server will try to directly contact the root nameservers, if the request from the forwarders servers times out for any reason. These direct requests will always fail, as the firewall filter rules will block them, and in the meantime the internal DNS will appear to *hang*. You may use the *slave* directive instead of forward-only option substatement.

The file named.ca specifies the nameserver used to request the list of root nameservers. It must contain the secure interface address of the firewall itself, since only the nameserver running on the firewall has direct access to the root nameservers and the external DNS, because the filter rules block any other DNS attempts through the firewall. The contents of /etc/dns/named.ca file are as follows:

```
.                          IN NS b-mail.b-corp-secure.ibm.com.
b-mail.b-corp-secure.ibm.com. IN A 10.2.1.1
```

The remaining configuration files for the internal DNS (named.data, named.rev and named.local) are conventional, but we have listed them here for completeness.

We have three zone data files as defined in the boot file. But we will introduce only one data file, named.data. The contents of /etc/dns/named.data file are as follows:

```
@       9999999 INSOA b-mail.b-corp-secure.ibm.com. root.b-mail.b-corp-secure.ibm.c
                           1.1     ; Serial
                           3600    ; Refresh
                           300     ; Retry
                           3600000 ; Expire
                           86400 ) ; Minimum
    9999999  IN NS    b-mail.b-corp-secure.ibm.com.
loopback     9999999    IN A      127.0.0.1; loopback (lo0)name/address
localhost    9999999    IN CNAME  loopback
gw           9999999    IN A      10.2.1.1
b-client     9999999    IN A      10.2.1.3
b-sample     9999999    IN A      10.2.1.4
b-mail       9999999    IN A      10.2.1.2
client2      9999999    IN A      10.2.1.8
exchange     9999999    IN A      10.2.1.9
```

The contents of /etc/dns/named.rev file are as follows:

```
; setting default domain to ... b-corp-secure.ibm.com
@                 9999999 IN  SOA   b-mail.b-corp-secure.ibm.com.
                                       root.b-mail.b-corp-secure.ibm.com. (
                         1.1           ; Serial
                         3600    ; Refresh
                         300     ; Retry
                         3600000 ; Expire
                         86400 ) ; Minimum
                 9999999 IN  NS   b-mail.b-corp-secure.ibm.com.
1     IN PTR gw.b-corp-secure.ibm.com.
3     IN PTR b-client.b-corp-secure.ibm.com.
4     IN PTR b-sample.b-corp-secure.ibm.com.
2     IN PTR b-mail.b-corp-secure.ibm.com.
8     IN PTR client2.b-corp-secure.ibm.com.
9     IN PTR exchange.b-corp-secure.ibm.com.
```

The contents of /etc/dns/named.local file are as follows:

```
@ IN SOA b-mail.b-corp-secure.ibm.com. paulgun.raleigh.ibm.com.
      (
      1999083000 ; serial
      1800    ; secondary chk
      900     ; retry
      172800  ; secondary expires
      1800    ; ttl
      )
      IN      NS      b-mail.b-corp-secure.ibm.com.
1     IN      PTR     localhost.
```

After you have set up the firewall DNS functions, you need to verify whether it works as you intended. You can use *nslookup* tool to track the name resolution process.

Note that BIND Version 8 uses different syntax from previous BIND versions. You may have to convert the Version 4 configuration file by running the Perl script src/bin/named/named-bootconf.pl that is distributed with the BIND source.

Chapter 6. Proxy

A proxy provides client access to network resources to which they do not have direct access. In a firewall environment, clients would typically not have direct access to the Internet.

Clients are configured to access the proxy directly, as though the proxy were providing the service. The proxy makes the same request the client had requested, as though it were the client. The identity of the actual client is hidden by the proxy. The results are then passed back to the client.

The access logs contained within the server will show only the proxy IP address as the source of requests. Using only the IP addresses found in the server's access logs, the server will be unable to determine the actual number of unique clients who are accessing it.

The following diagram illustrates a firewall running a proxy service allowing multiple clients secure access to remote servers on the Internet.

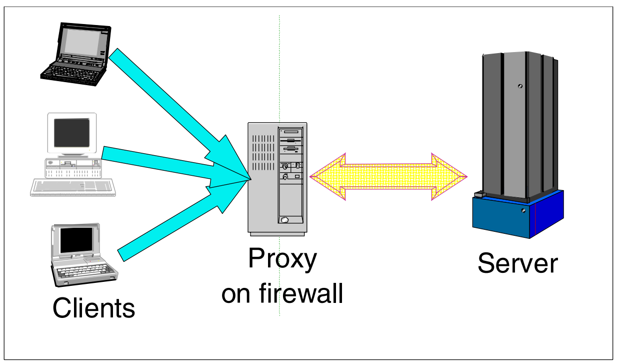

Figure 44. The proxy directly accesses server resources on behalf of the clients

Clients require no modification, apart from accessing the proxy IP address instead of the target server IP address. This makes the proxy ideal for client systems that do not currently support a client SOCKS configuration.

Unlike SOCKS, a different proxy application is provided for each protocol supported.

The IBM SecureWay Firewall V4 supports the following protocols via a specific proxy:

- HTTP
- HTTPS
- FTP via FTP proxy
- Telnet via Telnet proxy
- FTP via HTTP
- WAIS via HTTP
- Gopher via HTTP
- Transparent TELNET

- Transparent FTP
- DNS
- SMTP

Use SOCKS if you need to include protocols that are not in this list, such as NNTP, POP3, IMAP4 or SSH.

This chapter will deal with the HTTP, HTTPS, FTP and TELNET proxies.

6.1 HTTP proxy

The purpose of the Hypertext Transmission Protocol (HTTP) proxy is to relay HTTP requests securely through a firewall.

The HTTP proxy component is a derivative of the IBM Websphere Performance Pack Web Traffic Express product.

This proxy does not provide a caching function, but the customer can purchase WebSphere Cache Manager or WebSphere Performance Pack and install those caching proxy servers on top of the Firewall product.

6.1.1 Scenarios

We are going to describe three basic ways of using the HTTP proxy and then we will detail how to configure this proxy in “Basic configuration” on page 90.

The HTTP proxy can be used in several modes:

- Outbound
- Reverse
- Chained

Each scenario is described in the following sections.

6.1.1.1 Outbound

The most common use of the HTTP proxy will be to permit outbound sessions. Many secure clients will access the proxy service on the firewall. The firewall will then directly access the required resources on the Internet.

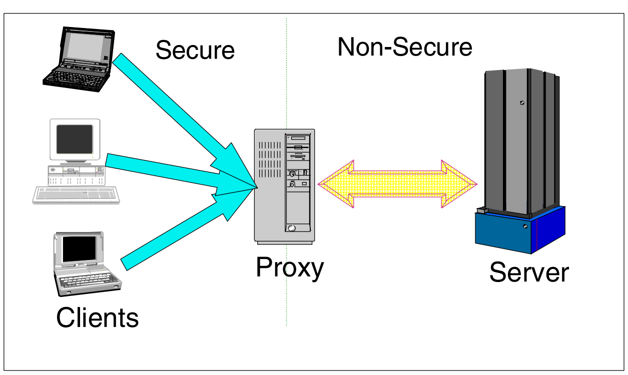

Figure 45. The outbound scenario is most common

6.1.1.2 Reverse HTTP

Reverse connections are possible, allowing the HTTP proxy to relay HTTP requests from many Internet clients to a secure server.

We believe the best practice is to place secure HTTP servers that will be accessed from the Internet in a segregated secure network or DMZ. This can be done by using a third TCP/IP interface on the firewall called the DMZ interface. We do not recommend allowing access to HTTP servers in your secure network from the Internet.

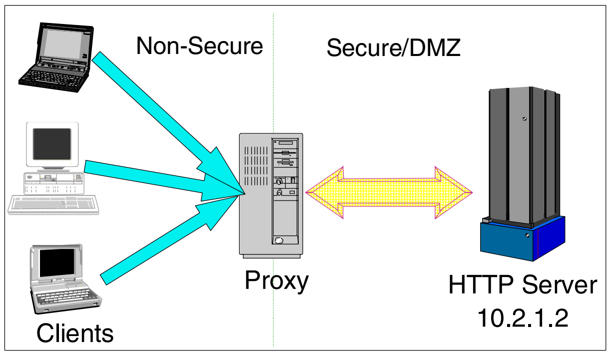

Figure 46. The reverse scenario permits inbound connections

Although we have not yet described the basic configuration of the HTTP proxy, we would like to point out two important proxy settings for this particular scenario. Using the HTTP proxy on the firewall to forward all requests to a specific HTTP server in the DMZ area requires two basic configuration steps:

1. Configure the HTTP port
2. Configure the Proxy directive

Configure HTTP port

Our protected HTTP server, 10.2.1.2, uses TCP port 80. We suggest that you configure the HTTP proxy on the firewall to use the same port, instead of the default 8080. Change the HTTP proxy port as described in “Proxy port” on page 92.

Configure Proxy directive

To configure the HTTP proxy on the firewall to forward all requests to a specific target host, use the "Proxy" directive in the /etc/ibmproxy.conf file.

The following example illustrates the correct entry to use in relation to the diagram shown in Figure 46.

```
Proxy /* http://10.2.1.2/*
```

Reverse proxy authentication

HTTP authentication is supported only on the secure interface, so it cannot be used to authenticate users from the Internet in this case.

NAT versus reverse proxy

Reverse NAT is a technique that can be used to achieve a similar result. See Chapter 9, “Network Address Translation” on page 221. Reverse Proxy HTTP has an important advantage over reverse NAT. This is because the client never has direct access to the target HTTP server in which to exploit HTTP vulnerabilities. The firewall provides the buffer between the Internet and secured HTTP server.

Abuse of HTTP Ports

Allowing inbound sessions from Internet clients requires careful consideration of the actual filter rules that are used.

A URL may contain additional information that is processed by the HTTP proxy. For example a URL may look like the following:

```
http://public.ibm.com:23/
```

This allows the client to override the default HTTP port number of 80. In this example, port 23 or the TELNET port is used.

If the DMZ HTTP server is a proxy server, then it may be possible to make requests to it that were not intended.

Ensuring adequate filter rules will prevent exploitation of this feature.

6.1.1.3 Chaining

Chaining is a technique that permits an HTTP proxy to access another intermediate proxy server before reaching the target server. This may be another HTTP proxy or a SOCKS server as shown in Figure 47.

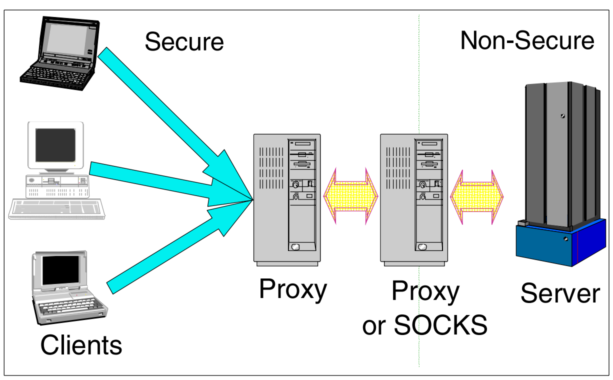

Figure 47. Chained proxies

A scenario that might use chaining is as follows: two small companies, both running independent firewalls, decide to merge. To control access to the Internet they decide to reduce the number of Internet gateways. An existing proxy can be chained to the active Internet gateway without requiring modifications to the client configuration environment.

We describe how to configure the proxy for this scenario in 6.1.2.1.

6.1.2 Basic configuration

The HTTP proxy is enabled by default. Clients can connect through to the proxy on the default TCP port 8080. Configure clients as described in “Client configuration” on page 121.

The following guided tour briefly describes each parameter. We have detailed our experiences that have required modification of the parameters from the default setting.

Changing some of the configuration parameters can increase performance or capacity, while other parameters can be detrimental to performance.

To begin HTTP proxy configuration, select **HTTP** in the main window:

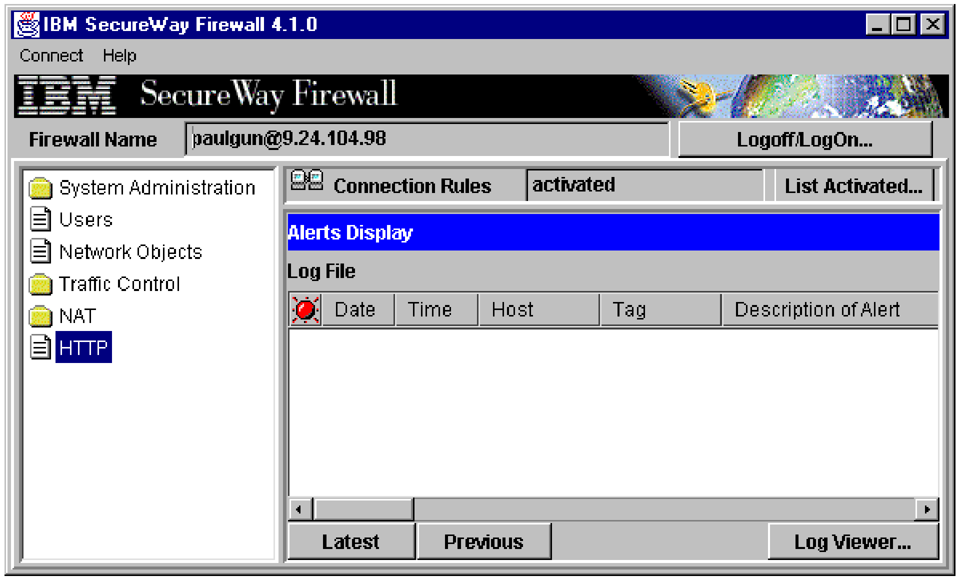

Figure 48. Select HTTP from the main window

The following sections walk through each tab in the configuration window.

6.1.2.1 Proxy settings

The proxy settings tab configures the most common parameters in the HTTP proxy.

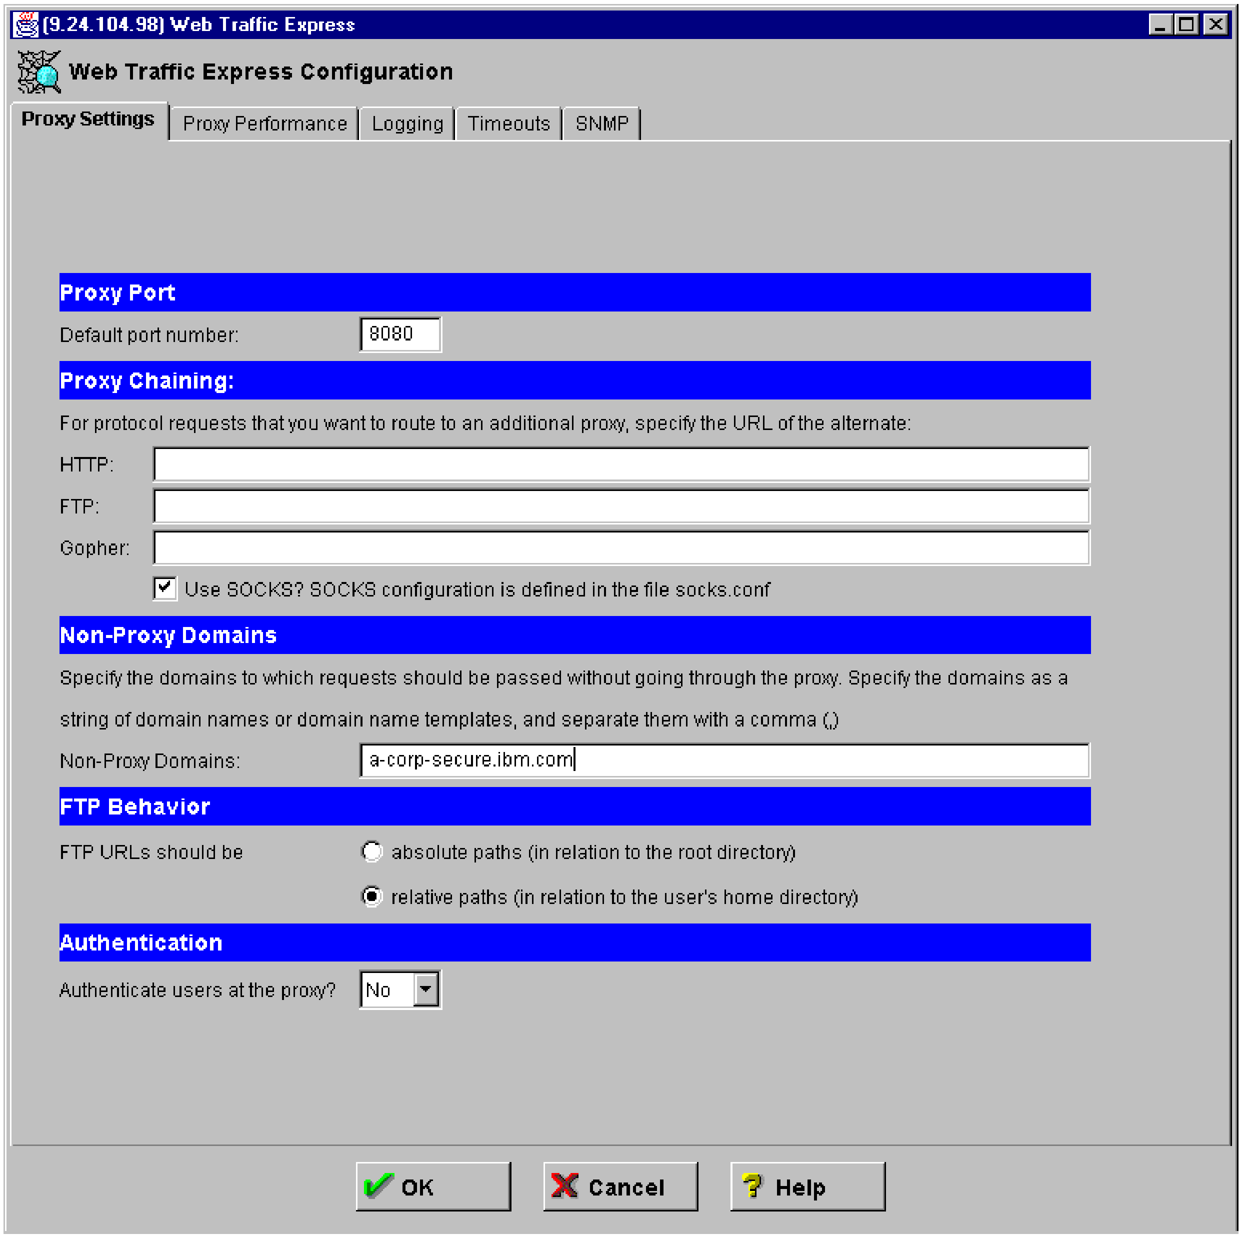

Figure 49. Proxy Settings tab

Proxy port

Clients need to know the IP address of the firewall running this HTTP proxy and the TCP port number to connect to. The default TCP port 8080 is a good choice for a proxy server to distinguish it from the actual HTTP servers running on the Internet at TCP port 80. This port must not conflict with TCP ports used by other applications on the firewall.

Proxy chaining

In the scenario described in “Chaining” on page 90, the firewall may be an intermediate host, requiring further upstream processing to reach the target servers. The upstream server may be either a SOCKS server or another HTTP proxy server.

By default the option to use SOCKS is enabled on the proxy settings tab. The SOCKS client configuration file /etc/socks.conf is specified. However, it initially contains only comments. Unless this file is edited, chaining will not occur.

In the following example, proxy chaining via SOCKS is enabled by simply editing the SOCKS configuration file /etc/socks.conf. The sample shows the "sockd" entry is used to indicate the upstream SOCKS host, socks-server.ibm.com. The "direct" entries are used to indicate hosts or subnets that can be directly reached. Use direct entries to specify your local or secure subnets so that traffic sent to them is not needlessly routed via a SOCKS server.

```
direct 9.0.0.0 255.0.0.0
direct 10.0.0.0 255.0.0.0
direct 172.16.0.0 255.240.0.0
direct 192.168.0.0 255.255.0.0
sockd @=socks-server.ibm.com 0.0.0.0 0.0.0.0
```

Figure 50. Client SOCKS configuration file /etc/socks.conf

If your upstream gateway is another HTTP server instead of a SOCKS server, deselect the SOCKS check box and specify the upstream proxy URL.

If the upstream proxy hostname is proxy7.au.ibm.com listening on port 8080, then you would specify the HTTP field as follows:

```
http://proxy7.au.ibm.com:8080/
```

Non-proxy domains

When chaining to an upstream HTTP proxy, this field is used to specify what DNS domains will be reached directly, instead of via the chained upstream proxy. This is equivalent to the "direct" keyword in the /etc/socks.conf file, defined “Chaining” on page 90.

We recommend that you place your secure side local domain in this field to prevent increased load on the upstream proxy.

Local domains

Local domains that can be accessed directly should not reflect off a proxy server. This increases load on the proxy server and consumes bandwidth.

FTP behavior

This field will determine if FTP URLs retrieve information via a relative or absolute path name relative to the logged-in user’s ID.

The following example uses the following FTP URL:

```
ftp://paul@somewhere.service.ibm.com/etc/services
```

This URL would be entered at the user’s browser that is configured to access the HTTP proxy. See .

If the user name "paul" on the host somewhere.service.ibm.com has a home directory /home/paul, then using absolute paths would retrieve the following file:

```
/etc/services
```

Using relative paths would retrieve the following file:

```
/home/paul/etc/services
```

Authentication

Authentication allows you to control who has access to the HTTP proxy. The default value of NO allows anyone access to the proxy.

To enable authentication, follow these steps:

1. Change the Authenticate users at the proxy setting to YES.
2. Create a firewall user. See “Basic configuration” on page 19 for detailed examples of adding users to the firewall.

   The following attributes of the user definition must be set:

   - Authority Level set to Proxy User
   - "Secure HTTP" set to Firewall Password

   The following image illustrates the setting of the required fields in the user definition:

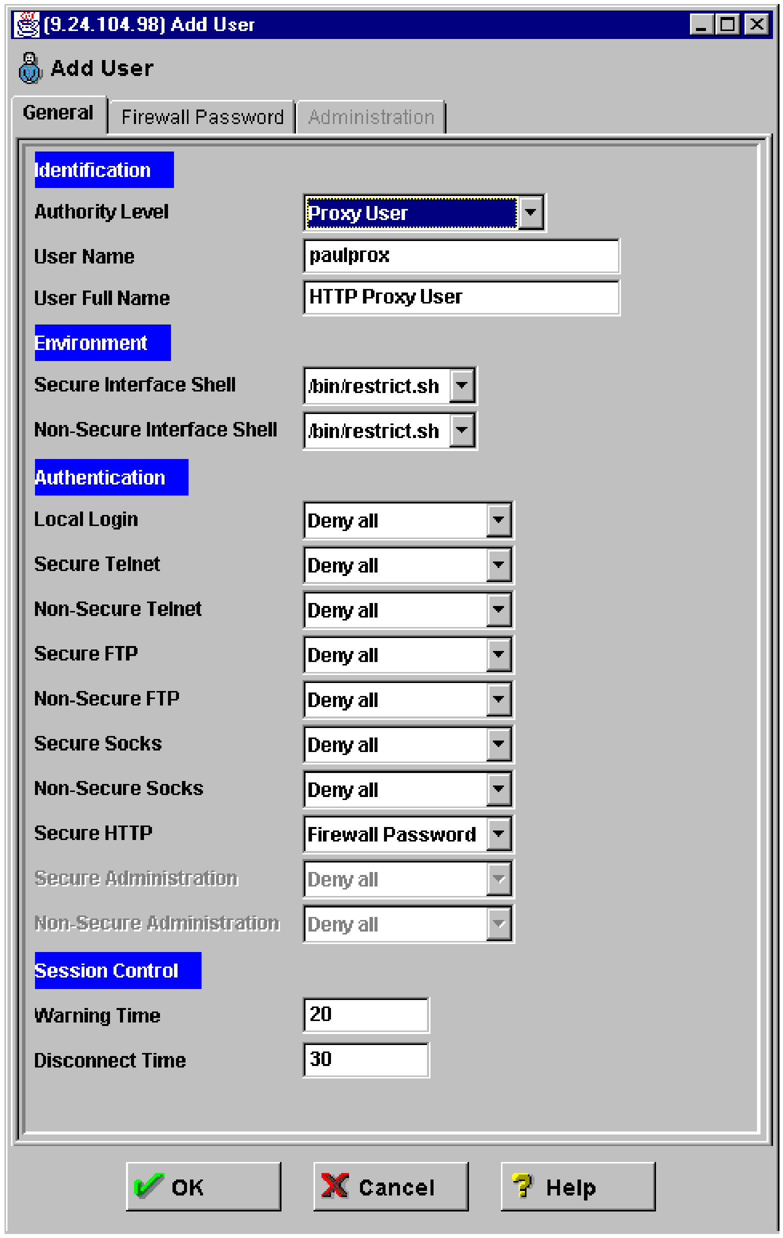

Figure 51. Creating a proxy HTTP user

When the user tries to access a URL using the HTTP proxy, the browser must supply to the HTTP proxy a valid username and password. In the case of the Netscape browser, a challenge/response dialogue window is presented:

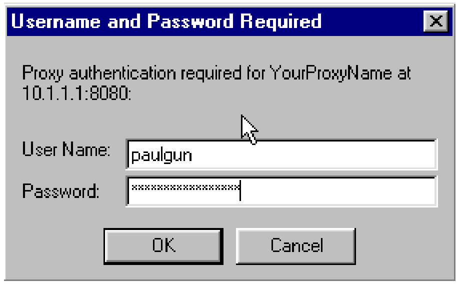

Figure 52. Client HTTP proxy user authentication

If the correct password is supplied the user's browser will remember the user and password. This will remain in effect until the user exits the browser.

Shared Browsers

Remember to exit your browser application to cause it to forget the current authentications. Unless you do this, someone else who uses your browser will inherit all the authentications you have used. This is a concern in a shared cyber-cafe style environment.

6.1.2.2 Proxy performance

The Proxy Performance tab modifies the performance characteristics of the HTTP proxy. Use of specific AIX measurement tools and benchmarks before and after modification of any parameter is required to ensure that changes do not adversely affect proxy performance.

Performance Measurement

Do not modify performance values unless you are prepared to quantify operating system measurements and bandwidth consumption before and after the change.

AIX performance skills are required.

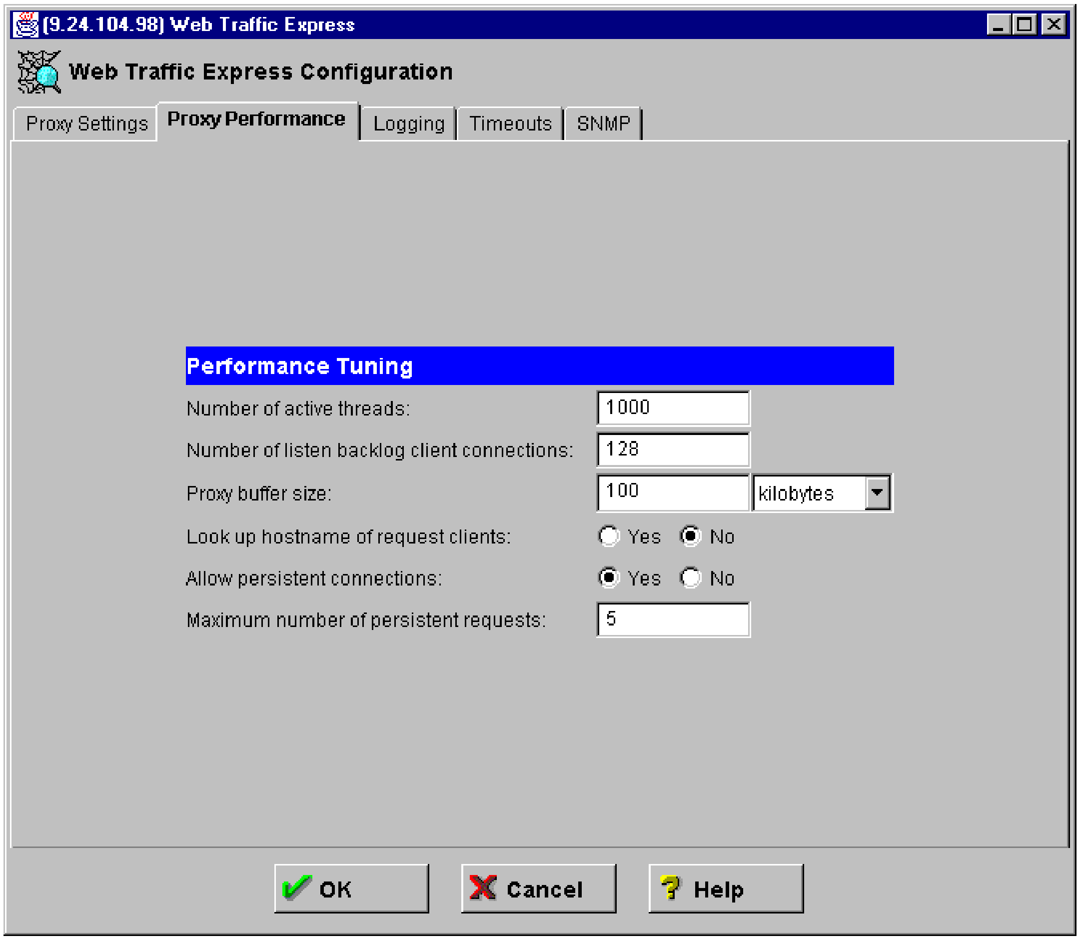

Figure 53. Proxy Performance tab

Number of active threads

This represents the total number of concurrent HTTP proxy tasks that will be working on requests. This will effectively throttle the number of requests that can be serviced by the HTTP proxy.

We recommend that you increase this value to accommodate the total population of proxy users. Planning considerations for this value include:

1. Total user population

   This represents the total number of users in your organization who will be using HTTP proxy. For example, your organization may have 1000 potential users.

2. Concurrency of user activity

   The total active users who are using HTTP proxy at any given point in time rarely approaches the total population. We can express this as a percentage of the total population, for example 10%.

3. Threads consumed per browser

   This indicates the total number of TCP sockets that are opened during a user's session. This is determined by three factors:

1. HTTP 1.1 server support

   If the HTTP Web server that is being accessed supports the HTTP 1.1 protocol, there is opportunity to reuse or cache socket connections. HTTP 1.0 Web servers do not support socket caching.

2. HTTP 1.1 client support

   The Web browser client must also support HTTP 1.1 to cache socket connections. We found Netscape Communicator Version 4.6 supported HTTP 1.1 by default, while Internet Explorer Version 5 required an optional parameter to be enabled to support HTTP 1.1 as discussed in Chapter 6.1.5.3, “Microsoft Internet Explorer support” on page 125.

3. Multi-threaded parallelism of Web clients

   High performance clients may support the operation of several parallel socket connections in order to increase performance. Each of these parallel sockets may itself be an HTTP 1.1 persistent connection.

   We have conducted some tests using various client environments to examine how many parallel socket connections are opened. Each of these open sockets will consume HTTP proxy thread resources. The number of parallel socket requests from a client will represent a factor to multiply by the total number of concurrent active users to determine the total number of threads required.

   The results were obtained by running a network trace of a browser opening a test HTML page we had created. This page contained 12 unique imbedded GIF images within the HTML document. We counted the number of socket open requests and looked at the TCP packet timestamps to determine if the socket connections were used in a sequential manner or in parallel. The following table summarizes the results:

Table 12. Number of utilized sockets required for various clients

Browser	Operating System	HTTP 1.1 support	Number sockets opened	Parallel
Netscape 4.05	AIX	always on	3	yes
Netscape 4.6	Microsoft Windows 98	always on	4	yes
Internet Explorer 5	Microsoft Windows 98	default off	2	yes
		turned on	2	yes
NCSA Mosaic 2.4	AIX	not supported	13	no

4. Number of imbedded objects per page

   The specific mix and number of imbedded images contained within an HTML document will potentially affect the thread workload. Large, complex pages with many inline images will consume more resources than single HTML pages without any inline images.

5. Time user spends on a page

   Persistent client connections to the Web server may if a user spends too long on a particular page. If a persistent timeout occurs, the penalty will be the creation of a new set of sockets and threads used to download the next Web page.

For example, if your site contains 1000 Netscape Version 4.6 users and 10% of these will be concurrently accessing the HTTP proxy during peak periods, we recommend you increase the number of threads to at least 400. They will provide capacity for 10% of 1000 multiplied by 4 or the potential parallelism of Netscape Version 4.6.

Increasing this value will place more potential load on the system. The load can be spread across multiple processors in an SMP system.

This can be increased up to the maximum value defined in the AIX kernel, SOMAXCONN. The default value for this TCP parameter in the AIX kernel is 1024. For example, it can be increased to XXXX connections by using the `no -osomaxconn=XXX` command.

This can be permanently increased by placing the above command at the end of the /etc/rc.net file.

Number of LISTEN backlog client connections

When no more threads are available to service requests, the TCP layer can hold new requests in a backlog queue. The queued requests remain here until threads become available to service the request or the connection time-outs.

The default backlog queue size is 128 potential client requests.

Increasing this value will allow HTTP proxy to accommodate greater spikes in the workload without increasing the number of threads.

If the size of the spike should exceed the size of the backlog queue, then the browser would get the message `TCP connection failure, try again later`.

Proxy buffer size

This value indicates how much data is buffered by HTTP proxy when receiving dynamically created output data. This is also known as the "chunk" size in the WTE proxy. HTTP servers using common gateway interface (CGI) scripts may output such dynamic data.

When the memory buffer is filled a "chunk" of data is returned to the client. HTTP proxy continues to return "chunks" of data to the client until there is no more data.

Increasing this value is generally not recommended as the proxy will appear to the end user to be unresponsive, where as in reality the proxy is attempting to buffer the data before returning it to the client. On slow WAN links this problem worsens.

Chunk Size and Perception

Unfortunately many HTTP sites do not return the actual byte count or size in the dynamically output HTTP stream.

This means the client system or HTTP proxy (acting on the clients behalf) must keep reading data until it times out or the end user aborts the download.

Placing a huge buffer between the client system and the server can potentially delay the client's receipt of data, while the buffer cache is filled. If this gets too long often users will prematurely abort the connection. It is perceived as more reliable for end users if they receive more chunks of data, rather that one huge chunk that has taken a long time to download.

Lookup hostname of request clients

When client requests are processed, this setting will control if DNS lookups are made of the client IP address.Changing this value from the default of "no" is usually detrimental to performance.

The only likely scenario where this option would be turned on is when offline processing of the logs is performed without access to a DNS server that can resolve the IP addresses.

Allows persistent connections

The HTTP 1.1 protocol definition allows client TCP socket connections to be cached or reused during the download of a Web page.

This significantly improves response time to the end user and is enabled by default.

How this works is easily described by first discussing the HTTP 1.0 protocol:

We will briefly explain how the HTTP 1.0 protocol behaves. An HTML Web page may look like the following:

```
<title>A-Corp-ibm-com</title>
<img src=glitzy-title.gif>
<h1>Welcome to A-CORP!</h1>
<p>We offer the following services:
<a href=abou.html><img src=about.gif></a>
<a href=prod.html><img src=products.gif></a>
<a href=serv.html><img src=service.gif></a>
<a href=supt.html><img src=support.gif></a>
<a href=help.html><img src=help.gif></a>
<a href=cont.html><img src=contact.gif></a>
```

Figure 54. A sample HTML Web page

To display this HTML page, a browser will need to separately open eight different files: the HTML wrapper and each of the imbedded GIF files.

Each file will create a unique TCP socket connection consisting of the following phases. Each will result in a separate TCP packet being sent across the network:

1. Request

2. Request acknowledge
3. Connection negotiation
4. Send data packets
5. Acknowledge each packet
6. Finish request
7. Finish request acknowledge
8. Finish acknowledge

The number of network packets required to transmit this simple HTML page is therefore approximately 64, assuming all data fits into one packet.

HTTP 1.1 persistent connections overcome this inefficiency by allowing a single TCP socket connection to transmit all of the data packets required for the entire HTML page.

It is recommended that this value be changed from the default value of "yes" only when an upstream HTTP server is not HTTP 1.1 compliant or the client environment does not support HTTP 1.1. See “Microsoft Internet Explorer support” on page 125 for tips on Internet Explorer configuration. If you are chaining to an HTTP 1.0-only proxy, you will need to change the option discussed at “HTTP 1.0 compatibility” on page 115.

Using persistent connections will require fewer sockets to be used on the HTTP proxy. Each connection will be open for a longer period than connections used in HTTP 1.0 requests.

Maximum number of persistent requests

When using a persistent connection as defined above, this value will determine the maximum number of requests to be serviced over a single persistent connection. Usually you will want an entire Web page to be downloaded over a single persistent connection. The structure of a Web page is typically HTML tags and imbedded inline images or other MIME types. The maximum value here will be the total number of these HTML documents and imbedded images in your largest document.

We suggest increasing this value to 8 to cater to the average number of inline files contained in a Web page as shown in Figure 54 on page 100.

6.1.2.3 Logging

The logging tab configures HTTP logging, useful for reporting and diagnostic purposes.

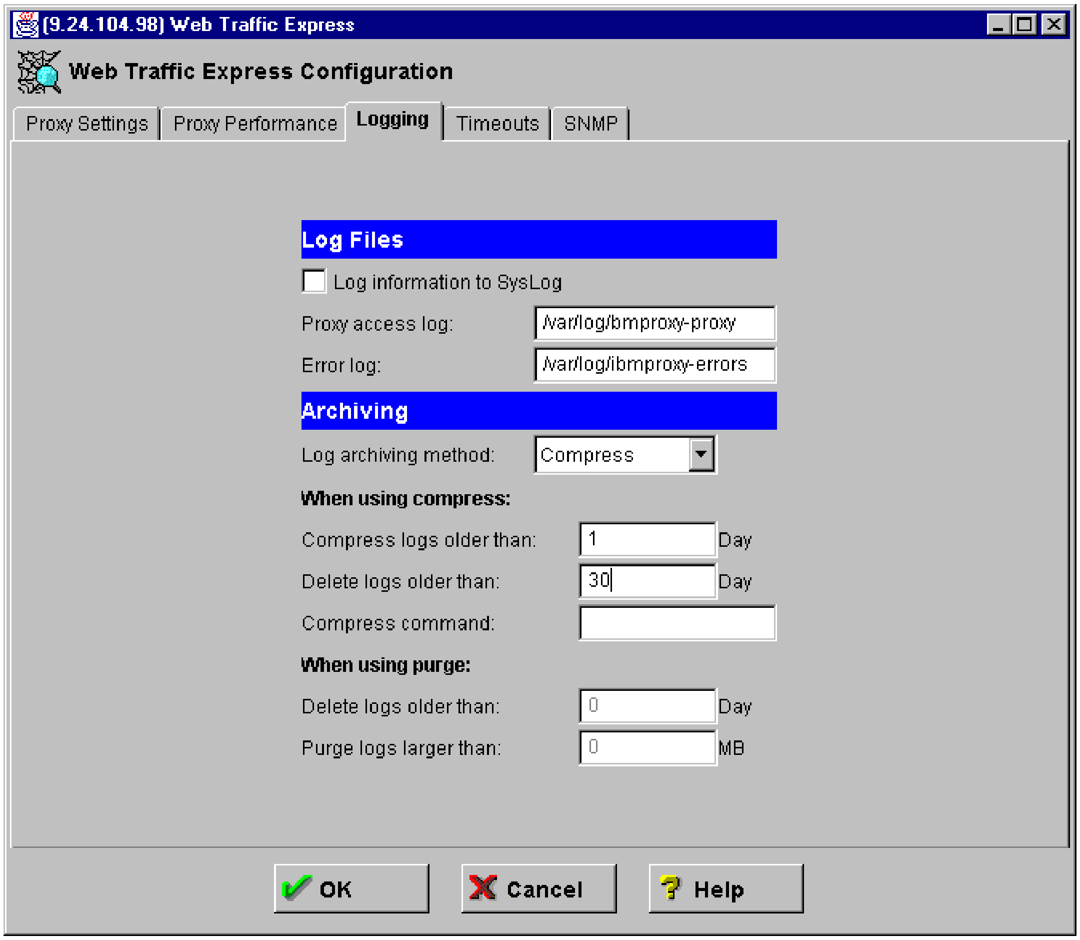

Figure 55. Proxy Logging tab

Log files

As described in Chapter 2, “Installation” on page 7, a dedicated logging file system /var/log is created.

We recommend that the HTTP logs be diverted to this file system instead of consuming limited space in the /usr file system.

Change the proxy access log and error log to:

```
/var/log/ibmproxy-proxy
/var/log/ibmproxy-errors
```

Using syslog to capture the HTTP logs would not be recommended unless you are using a consolidated syslog logging server and you have reporting tools that understand the syslog format.

Logging file system

Create a dedicated logging file system, /var/log, to capture your logs. The default location writes to the /usr file system that is normally close to 100% full. The /usr filesystem is designed only to contain binary commands for installed products and should not be used for temporary logging files.

We do not recommend using the /var filesystem. If this ever reaches 100% capacity, it will adversely affect the operating system. Using log management will contain the size of logging files; however, this is not always guaranteed in the case of a denial of service (DOS) attacks on your firewall that should capture all the denied packets.

The logged data includes the client IP address, URL and HTTP response code. The format is a common HTTP log format that can be imported into various tools. The following log fragment shows some typical Web transactions:

```
10.2.1.3 - - [02/Sep/1999:01:00:15 -1000] "GET
http://www.webshopper.com/graphics/buttons/showproducts.gif HTTP/1.0" 200
667
10.2.1.4 - - [02/Sep/1999:02:00:16 -1000] "POST
http://www.pc.ibm.com/msprotect/ncommerce3/ExecMacro/ccadmin.d2w/report
HTTP/1.0" 200 29858
10.2.1.5 - - [02/Sep/1999:02:00:19 -1000] "GET
http://ads.fairfax.com.au/image.ng/Params.richmedia=yes&site=sold&adspace=
469x60&loc=top HTTP/1.0" 302 -
10.2.1.4 - - [02/Sep/1999:02:00:19 -1000] "GET
http://static.wired.com/advertising/blipverts/WebMD/aging.gif HTTP/1.0" 200
12292
```

Each URL that is logged by the HTTP proxy can be broken down into individual fields. In the first URL in the above sample proxy log, each field is described as follows:

1. Client IP address

   In this case the client browser is running on the IP host, 10.2.1.3

2. Time stamp

   The date, time and time zone is recorded in each entry.

3. HTTP request

   The HTTP request is enclosed in double quotes (") and is composed of three subfields:

   1. HTTP Method

      The most common type of request is an HTTP GET request. This is commonly used by Web browsers to download HTML pages from Web servers. The sample log above also includes an example of the HTTP POST method that is used by Web browser to send information to a Web server. This might be used in the case of a user filling out a form contained within an HTML page.

   2. URL

      The full URL that is requested is logged here.

3. HTTP version

   High-performance Web servers use HTTP/1.1. HTTP/1.0 is also used by many Web servers and is described in more detail in “HTTP 1.0 compatibility” on page 115.

4. HTTP response code

   A successful HTTP transaction is indicated by the response code 200. Other codes are useful for diagnosing problems as described at “Interpreting HTTP response codes” on page 117. In the sample log above the third entry from client IP address 10.2.1.5 returns a code 302. This code indicates a URL redirection has occurred. That scenario is fairly common and does not indicate a failure. Assuming non-200 error codes are errors is not correct and we encourage you to understand the different classes of response codes.

5. Data size in bytes

   The size of data returned by the Web server is shown. We have seen in some rare cases the Web server is returning an unknown size or "-" value. This is usually associated with output from dynamic CGI programs. This is discussed further in “Proxy buffer size” on page 99.

In addition to the proxy access and error logs, we recommend changing the logging directory for the local document access log. The following activities will create entries in the local document access log:

- Advanced function client configuration files. See “Automatic” on page 123.
- Advanced function trace and reporting. See “Activity statistics” on page 118.

Edit the configuration file /etc/ibmproxy.conf and update the "AccessLog" entry. The default value is:

```
AccessLog        /usr/lpp/internet/server_root/logs/ibmproxy-log
```

Change this to:

```
AccessLog        /var/log/ibmproxy-log
```

This will ensure these logs are always written to the large dedicated logging filesystem you have created.

Archiving

The HTTP logs can quickly fill even a large dedicated logging file system.

The log management tools provided with HTTP proxy allow you to compress or purge log files as they reach a certain age.

We have found that a successful strategy to manage the logs is to enable the following parameters:

Field	Value
Method	Compress
Compress Age	2 days
Delete Age	30 days

Field	Value
Compress Command	`tar -cvf /var/log/archive/log%%DATE%%.tar %%LOGFILES%% ; compress /var/log/archive/log%%DATE%%.tar`

After a log file reaches two days old, it will be compressed. This allows easy examination of the current log and yesterday's log. Older logs can be uncompressed as required. After 30 days the log is deleted. This is a reasonable value compared with the size of the compressed log files. This also is dictated by your security policy regarding audit records.

The `compress` command shown is compatible with the Web Traffic Express Version 2 product. This can also be a custom shell script that you have written that is invoked by HTTP proxy at midnight.

6.1.2.4 Timeouts

The Timeouts tab (see Figure 56) alters specific time out values within the HTTP proxy. We do not recommend changing timeouts unless you can specifically prove modification is required.

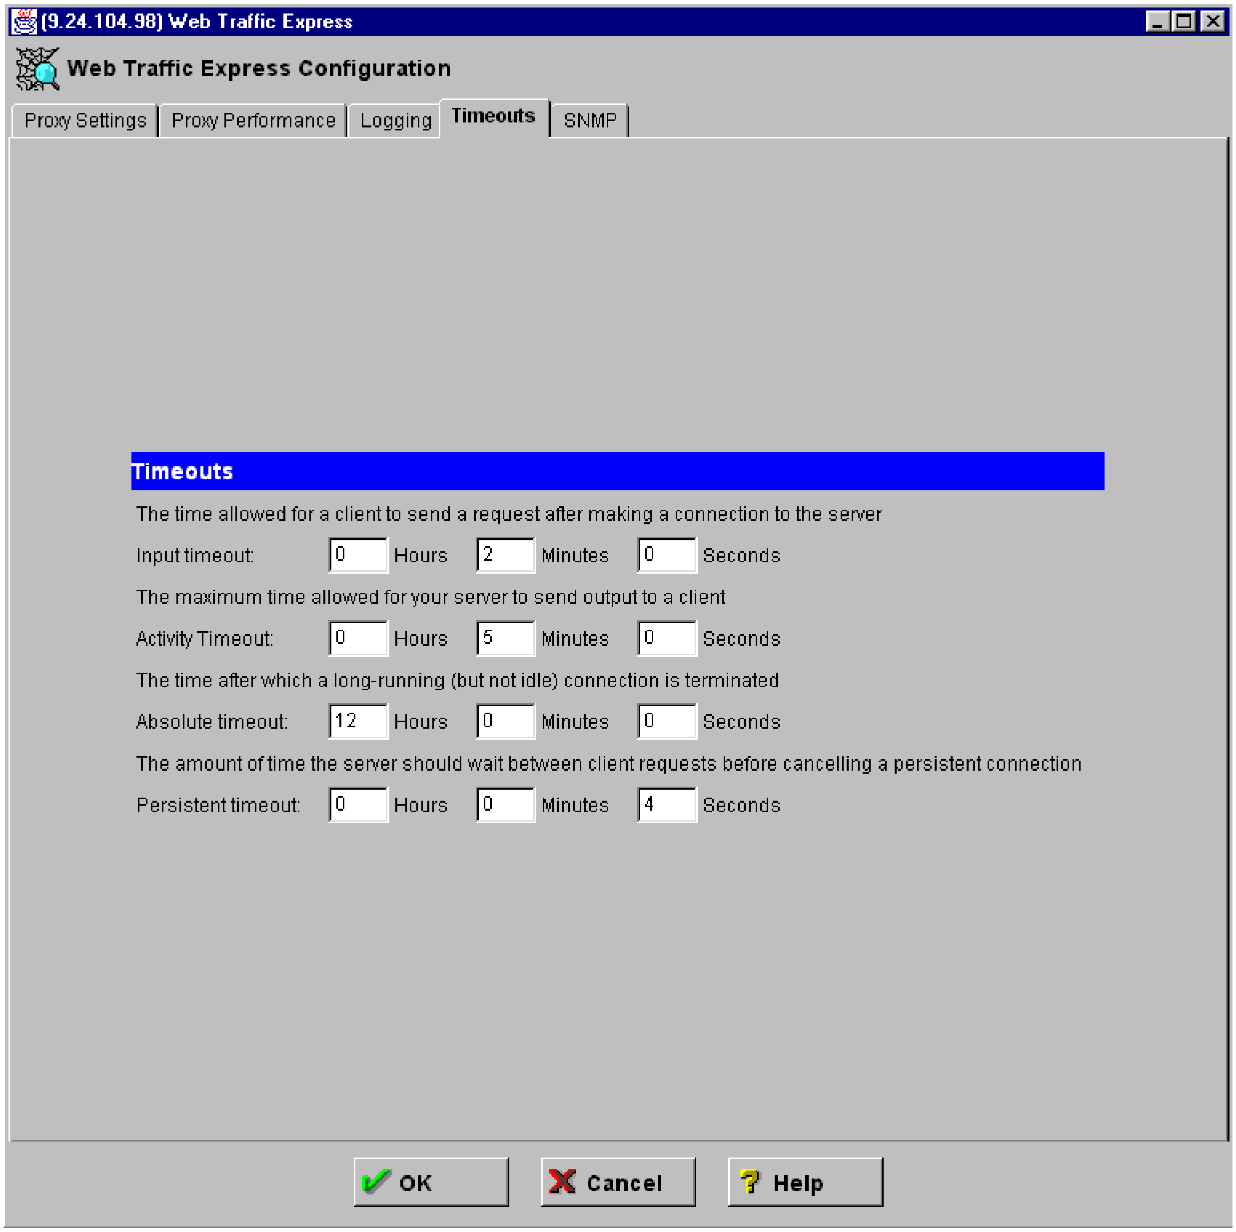

Figure 56. Proxy Timeouts tab

Adjusting the timeout values may be appropriate in the following cases:

- Users are reporting time out error messages appearing from the HTTP proxy
- You are constrained on available bandwidth
- Remote servers are experiencing large latency.

Timeout values that are too short may prematurely terminate normal, slow connections.

The main reason for having the time out parameters is to prevent system resources from being unnecessarily consumed for connections that have become dormant for too long without any data being sent or received.

Field tests in Australia connected to the Internet via slow WAN links had forced us to increase the timeout values as follows:

Time out parameter	Default			Recommended for slow connections		
	Hours	Minutes	Seconds	Hours	Minutes	Seconds
input	0	2	0	0	5	0
output	0	5	0	10	0	0
absolute	12	0	0	20	0	0
persistent	0	0	4	0	0	60

It is not at all unreasonable to use such a large absolute timeout in this environment. Consider a large company in Australia of say 1000 users connected to the Internet using a modest 2 Mbps bandwidth WAN. If our security policy permits an unrestricted download file size, then it can take several hours to download say a 100 MB file from a US-based FTP site. The absolute timeout will need to be adjusted for these situations. Premature timeouts are extremely wasteful of bandwidth in this case and will result in data retransmission.

6.1.2.5 SNMP

Network management of the HTTP subsystem is possible using the Simple Network Management Protocol (SNMP). An HTTP specific management information base (MIB) is available.

Network managers can view the HTTP-specific MIB information from a management console such as one provided by Tivoli.

To use this feature follow these steps:

1. The SNMP subagent must already be started as described in “Firewall Management” on page 367.
2. Add the following information to the end of the configuration file /etc/snmpd.peers:

   ```
   "dpid2" 1.3.4.1.4.1.2.3.1.2.2.1.1.1 "dpid_password"
   ```

3. Add the following information to the end of the configuration file /etc/snmpd.conf:

   ```
   smux 1.3.4.1.4.1.2.3.1.2.2.1.1.1 dpid_password
   ```

4. Enable SNMP in the SNMP window checkbox

Figure 57 illustrates the SNMP configuration window:

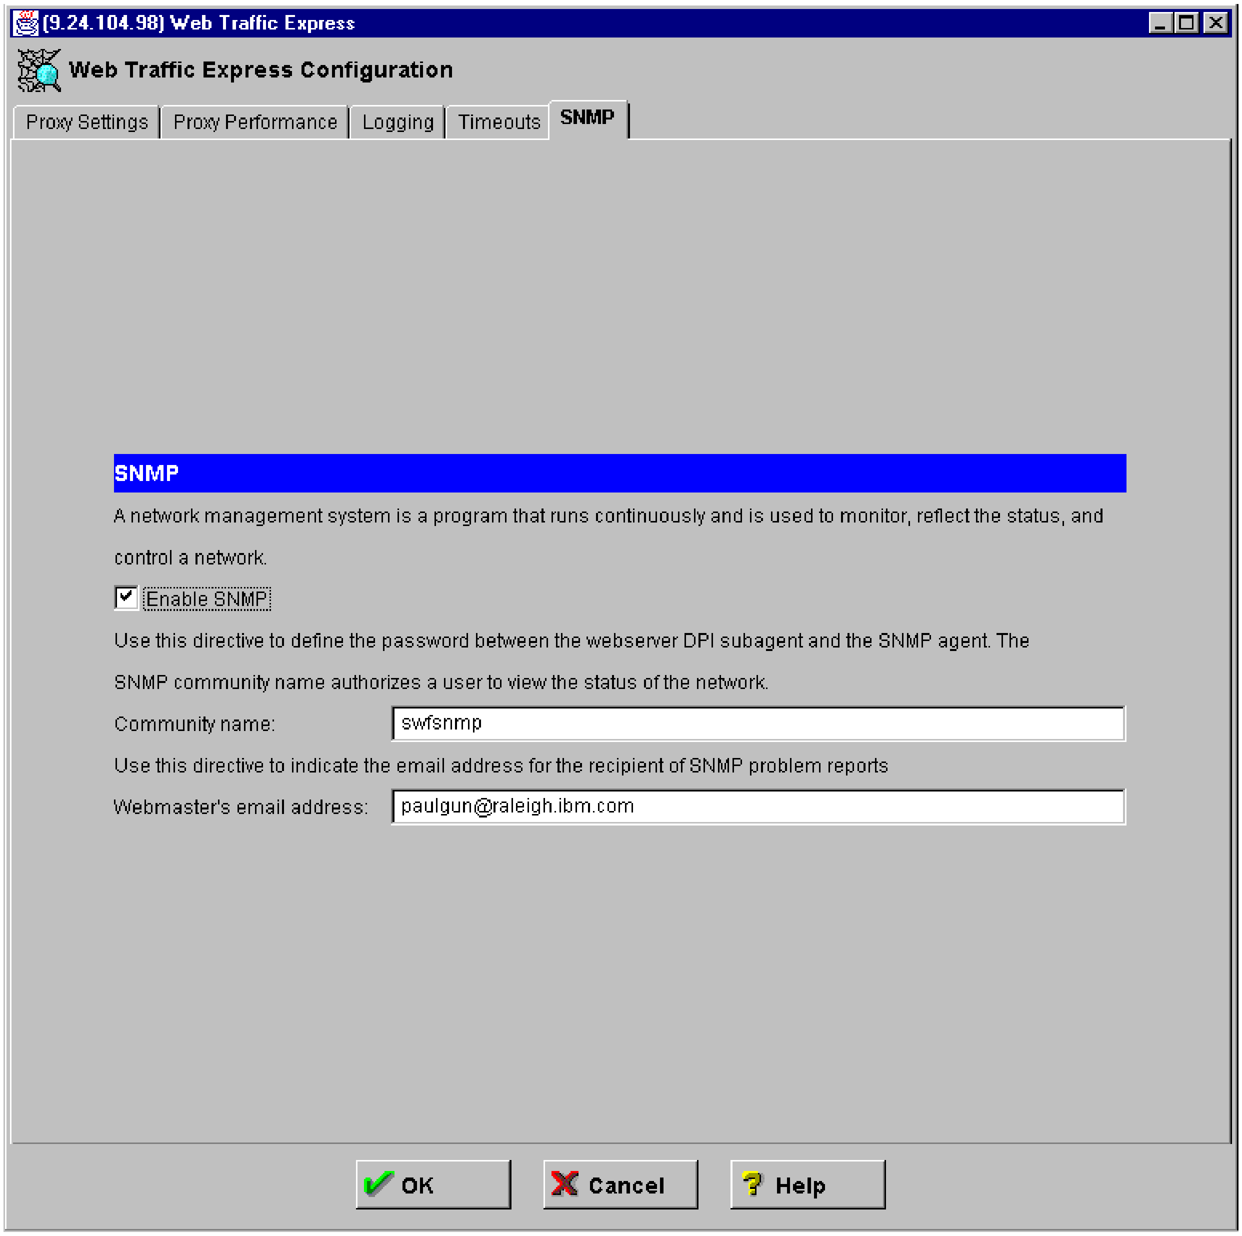

Figure 57. SNMP tab

5. Update the community name as defined in /etc/snmpd.conf

   The above window shows the default firewall community name.

> **Community Name**
>
> The default community name for the HTTP component is "public". The firewall SNMP subagent is configured to use a default community name of "swfsnmp". Both of these should be changed to a non-guess able value. Information about the firewall such as resource usage or interface traffic should not be considered public knowledge by any robust security policy definition.

6. Restart daemons in order:

4. snmpd
5. dpid2
6. ibmproxy

The following commands are used to restart all the daemons in sequence after they have been stopped:

```
startsrc -s snmpd
dpid2
startsrc -s ibmproxy
```

Rebooting the firewall can achieve the same result.

7. Verify MIB Information

The HTTP SNMP MIB information can be directly queried using the primitive snmpinfo command. The following example retrieves the port number that HTTP is using:

```
snmpinfo -m get -c swfsnmp -h localhost 1.3.6.1.4.1.2.6.154.1.1.1.1.1.8.1
1.3.6.1.4.1.2.6.154.1.1.1.1.1.8.1
```

We have found the following MIB objects to be useful:

Table 13. HTTP SNMP useful objects

Object ID	Description
1.3.6.1.4.1.2.6.154.1.1.1.1.1.12.1	Throughput in connections per second
1.3.6.1.4.1.2.6.154.1.1.2.1.1.3.1	Bytes received total
1.3.6.1.4.1.2.6.154.1.1.2.2.1.2.1.3	Number of HTTP error responses in the 400-499 range (transient failure) issued by HTTP proxy
1.3.6.1.4.1.2.6.154.1.1.2.8.1.2.1.404	Number of HTTP errors specifically returned as error 404 (not found)
1.3.6.1.4.1.2.6.154.1.1.2.4.0	Total time outs on HTTP proxy
1.3.6.1.4.1.2.6.154.1.1.3.1.1.5.1	HTTP status 1 = up 2 = down 3 = halted 4 = congested. No inbound sessions can be processed

8. Remote monitoring

The SNMP manager can query the firewall to monitor the health of the HTTP component.

The MIB definitions are available in the following file:

/usr/lpp/internet/server_root/Docs/IBMwwwmib.my

6.1.3 Advanced options

The following advanced options allow more control over the HTTP proxy than provided in the configuration client.

These options are not directly accessible via the GUI interface.

6.1.3.1 Manual proxy control

Manually restarting the proxy is usually required after manually editing the configuration file /etc/ibmproxy.conf.

Stopping

If the HTTP proxy needs to be stopped immediately, issue the following command:

```
stopsrc -s ibmproxy
```

Verify the proxy has stopped by issuing the following command:

```
lssrc -s ibmproxy
```

Starting

If the HTTP is currently inoperative, it can be started by issuing the following command:

```
startsrc -s ibmproxy
```

Automatic restart

If HTTP proxy stops for any reason, it can be desirable to have the process automatically restart.

This can be enabled via the automatic restart option on the HTTP proxy subsystem. Issue the following command to enable automatic restart:

```
# chssys -s ibmproxy -R
0513-077 Subsystem has been changed.
```

This may be useful for large sites to maintain continuous availability for clients.

Autostart on reboot

By default, the HTTP proxy process is stated on reboot via the following entry in the /etc/inittab configuration file:

```
rcibmproxy:2:wait:/etc/rc.ibmproxy > /dev/console 2>&1 # Start HTTP daemon
```

Disable permanently

If you will not be using the HTTP proxy and wish to permanently disable it from restarting on system boot, comment out the `rcibmproxy` entry in the file /etc/inittab with the colon ":" comment character:

```
:rcibmproxy:2:wait:/etc/rc.ibmproxy > /dev/console 2>&1 # Start HTTP daemon
```

6.1.3.2 Site Blocking

HTTP proxy has the ability to restrict or censor access to Internet sites. This can be used to prevent employees from accessing sites that are inappropriate to your security policy.

Freedom of speech

Some public services such as schools or libraries should consult local government regulations regarding freedom of speech laws. It may be in violation of local laws to allow or not to allow certain types of content such as controversial political, religious or pornographic content.

Enabling filtering on your HTTP proxy will involve some additional performance overhead, as part of each document will need to be scanned.

There are two types of URL filtering:

- By URL pattern
- By document content

Filtering by URL pattern only allows restriction of sites by matching patterns in the URL. Content filtering considers the actual content of the HTML pages.

URL Pattern

Patterns that match absolute URLs or wild card URLs using the "*" character can be formed.

In the following example we wish to forbid access to sites that contain the word "violence" or "gambling" in the URL. To do this add the following lines to the /etc/ibmproxy.conf file:

```
Fail    *violence*
Fail    *gambling*
```

After editing the configuration file you will need to refresh the proxy with this command:

```
refresh -s ibmproxy
```

Any attempt by the user to access pages that match the template will fail. The following sample URLs would match the Fail rule:

```
http://www.gambling.casinoabc.com/
ftp://ftp.gambling.casinabc.com/casino_payoffs.txt
https://www.casinoabc.com/gambling_tips.html
```

The user will see the following error page:

```
Error 403 - Access forbidden by rule.

Explanation: Either the file requested is specifically blocked by a Fail
directive or it does not match any of the files that are allowed to be
accessed according to other request mapping directives.

Action: No action is required.

URL: /
```

Content

Content filtering is a more sophisticated scheme that considers the actual content of an HTML page. This is achieved indirectly by evaluating the page content beforehand, then assigning a label to the content according to a classification scheme. This is performed by the author of the document by including meta tags inside the HTML document or by an independent organization called a label bureau.

The content filtering system is called the Platform for Internet Content Selection (PICS). One of the actual classification schemes supported by the HTTP proxy is the RSACi scheme, managed by the Internet Content Rating Organization (ICRA).

Sites rated with RSACi are encouraged to identify themselves with the following icon:

Self-classified sites contain an HTML META tag inside the document. The following example shows the self-classified RSACi rating label for the IBM corporate Web site, `http://www.ibm.com`:

```
<META HTTP-EQUIV="PICS-Label"
CONTENT='(PICS-1.1 "http://www.rsac.org/ratingsv01.html"
l gen true comment "RSACi North America Server"
 by "epc@www.ibm.com" for "http://www.ibm.com/"
on "1997.07.05T21:46-0500"
r (n 0 s 0 v 0 l 0))'>
```

The important piece of the label is the rating itself. From the META tag above, the rating given indicates no or a level zero (0) amount of: nudity (n), sexual (s), violent (v) or inappropriate language (l) content. Level four (4) is the highest rating for each category. We downloaded from `http://www.w3.org/PICS` the rating label definitions. We have included a part of this to illustrate the definition of violence:

```
((PICS-version 1.1)
 (rating-system "http://www.rsac.org/ratingsv01.html")
 (rating-service "http://www.rsac.org/")
 (name "The RSAC Ratings Service")
 (description "The Recreational Software Advisory Council rating service.
Based on the work of Dr. Donald F. Roberts of Stanford University, who has
studied the effects of media on children for nearly 20 years.")
 (default (label-only true))
 (category
  (transmit-as "v")
  (name "Violence")
  (label
   (name "Conflict")
   (description "Harmless conflict; some damage to objects")
   (value 0))
  (label
   (name "Fighting")
   (description "Creatures injured or killed; damage to objects; fighting")
   (value 1))
  (label
   (name "Killing")
   (description "Humans injured or killed with small amount of blood")^
(value 2))
  (label
   (name "Blood and Gore")
   (description "Humans injured or killed; blood and gore")
   (value 3))
  (label
   (name "Wanton Violence")
   (description "Wanton and gratuitous violence; torture; rape")
   (value 4)))
```

For further information on the ratings schema, refer to the RSACi Web site `http://www.rsac.org/ratingsv01.html`.

Other rating schemes are possible via the flexible PICS system; however, our example will use the RSACi system.

Accuracy of Ratings

The accuracy of the rating is at the sole discretion of the person supplying the PICS ratings tags within the HTML document.

Unfortunately for our lab we did not have access to a label bureau to supply ratings of content. Therefore we have used a very simple example that blocks access to a sample site "gamblers.casinoabc.com". We edited the PICS definition in the /etc/ibmproxy.conf file as follows:

```
DefinePicsRule "RSAC Example" {
    (PicsRule-1.0
        (
            serviceinfo (
                        name "http://www.rsac.org/ratingsv01.html"
                        shortname "RSAC"
                        available-with-content "YES"
                        )
            name        (
                        rulename "RSAC Example"
                        description "Example rule using the RSAC system to
block access to sites."
                        )
            passURL     ("http://www.ibm.com/*")
            failURL     ("http://gamblers.casinoabc.com/*")
ibm-javelin-extensions (
                               active "yes"
)
Filter ( Pass '((RSAC.v < 3) && (RSAC.s < 3) && (RSAC.n < 3) && (RSAC.l <
3))' )
)
    )
}
```

Figure 58. PICS configuration entry in /etc/ibmproxy.conf

Users attempting to access the site `http://gamblers.casinoabc.com/` will be presented with the following HTTP error page:

```
Error 403 - Blocked by filtering rule.

Explanation: There is a PICS rule associated with your request and
the rule has blocked it.

Action: No action is required.

URL: http://gamblers.casinoabc.com/
```

This will occur when a URL matches either an explicit fail rule or the content matches the RSACi filter rule.

For further information on constructing PICS filter rules, see the URL:

`http://www1.raleigh.ibm.com/pics/PicsRULZ.html`

SSL and content

Internet sites using the HTTPS or SSL protocol are sending encrypted data to and from the client browser. The contents of this data are not decrypted by HTTP proxy on the firewall; it is transparently proxied through.

HTTP proxy will by default check the home page of the site using HTTP and apply the ratings found here to all of the HTTPS documents found within the site.

We have found both Netscape Version 4 and Internet Explorer Version 5 directly support a native RSACi content rating service within the browser. On Netscape this is available by clicking **Help -> Netwatch**. Internet Explorer users will find this by clicking Internet **Options -> Content Advisor.**

Users who have access to the configuration menus of their browser could potentially override these settings, thereby overriding the site security policy. We recommend using RSACi on HTTP proxy, not the end-user's browser.

6.1.3.3 Customizing error pages

If a client that is using HTTP proxy receives an error, an HTML error page from HTTP proxy is returned to the client.

For example, the following screen appears:

```
Error 403 - Access forbidden by rule.

Explanation: Either the file requested is specifically blocked by a Fail
directive or it does not match any of the files that are allowed to be
accessed according to other request mapping directives.

Action: No action is required.

URL: /
```

The meaning of these errors is discussed in “Interpreting HTTP response codes” on page 117.

These error response pages are HTML files stored in the following directory:

`/usr/lpp/internet/server_root/pub/errorpages/`

We suggest adding the following information to the pages as appropriate:

- Helpdesk contact information
- Link to your security policy

We added the following HTML text to the end of the file:

`proxynotauth.htmls`

```
<p>Please contact HELPDESK on x4489 for any queries.
</BODY></HTML>
```

Figure 59. Customizing HTML error messages

6.1.3.4 HTTP 1.0 compatibility

Some environments may not support HTTP 1.1 clients or chained proxy servers. If you need to force the HTTP proxy to use the inefficient HTTP 1.0 protocol, edit the /etc/ibmproxy.conf configuration file and add the following option:

```
SendHTTP10outbound *
```

6.1.4 Diagnostics

In addition to the logging facilities shown at “Logging” on page 101 and the network management function shown in “SNMP” on page 107, there are some useful diagnostic HTML pages that can be used for ad-hoc queries. These pages contain HTTP statistics (Figure 62 on page 118) and a trace facility (Figure 63 on page 120).

Before these HTML pages can be viewed, an administrative user must be defined as per the following section:

6.1.4.1 Administrative user

The administrative user definition is different to the proxy user definition as described in “Authentication” on page 94.

To add an administrative user that will be used to view the diagnostic pages, follow these steps:

1. Create a directory for the password file. Use the following command:

   ```
   mkdir /usr/lpp/internet/server_root/protect
   ```

2. Create a password file using the `htadm` command:

   ```
   htadm -create /usr/lpp/internet/server_root/protect/webadmin.passwd
   ```

3. Add an administrative user, as illustrated in the following command (all on the same line):

   ```
   htadm -adduser /usr/lpp/internet/server_root/protect/webadmin.passwd
   paulgun secretpassword Paul Gunther
   ```

4. Add the password file to the PROT-ADMIN stanza in the file /etc/ibmproxy.conf. The following fragment of the file /etc/ibmproxy.conf illustrates the updated content:

```
Protection PROT-ADMIN {
          ServerId     Private_Authorization
          AuthType     Basic
          GetMask      All@(*)
          PutMask      All@(*)
          PostMask     All@(*)
          Mask         All@(*)
         PasswdFile   /usr/lpp/internet/server_root/protect/webadmin.passwd
}

Protect /admin-bin/* PROT-ADMIN
Protect /reports/*   PROT-ADMIN
Protect /Usage*      PROT-ADMIN
```

Figure 60. Password file user definition in /etc/ibmproxy.conf

5. Refresh the HTTP proxy server with the following command:

   ```
   refresh -s ibmproxy
   ```

   The following authentication prompt will appear the first time the diagnostic HTML pages are downloaded:

Figure 61. Advanced diagnostic user authentication

Non-Secure Authentication

Users cannot be authenticated at the non-secure interface.

6.1.4.2 Interpreting HTTP response codes

It is important to note that many HTTP errors generated by the HTTP proxy are external to the proxy itself and cannot be corrected by the firewall administrator. Examples of these would include the HTTP error response 404 "Not Found".

Fix the proxy!

Many end users see an HTTP error response code as interpreted by the HTTP proxy running on the firewall. This is often perceived by end users as a problem with the HTTP proxy or the firewall.

All end users and firewall administrators will need to become familiar with the response codes returned to determine the cause of the problem.

Firewall administrators should take careful note of all 500 class errors. These are generated when an internal HTTP error is produced. Examples of these may include resource limitations on the proxy such as running out of available threads.

The following table classifies all HTTP error response codes:

Table 14. HTTP response code ranges

Error Range	Error Class
100-199	Informational - Request received, continuing process
200-299	Success - The action was successfully received, understood, and accepted
300-399	Redirection - Further action must be taken in order to complete the request
400-499	Client Error - The request contains bad syntax or cannot be fulfilled
500+	Server Error - The server failed to fulfill an apparently valid request

6.1.4.3 Activity statistics

Some useful internal HTTP counters can be accessed by appending the following path to the HTTP URL:

```
/Usage/Initial
```

For example, if the HTTP proxy is running on the firewall at secure IP address 10.1.1.1 listening on TCP port 8080, then the full URL you need to view is:

```
http://10.1.1.1:8080/Usage/Initial
```

You can select **Refresh** on the following HTML screen at any time to update the values.

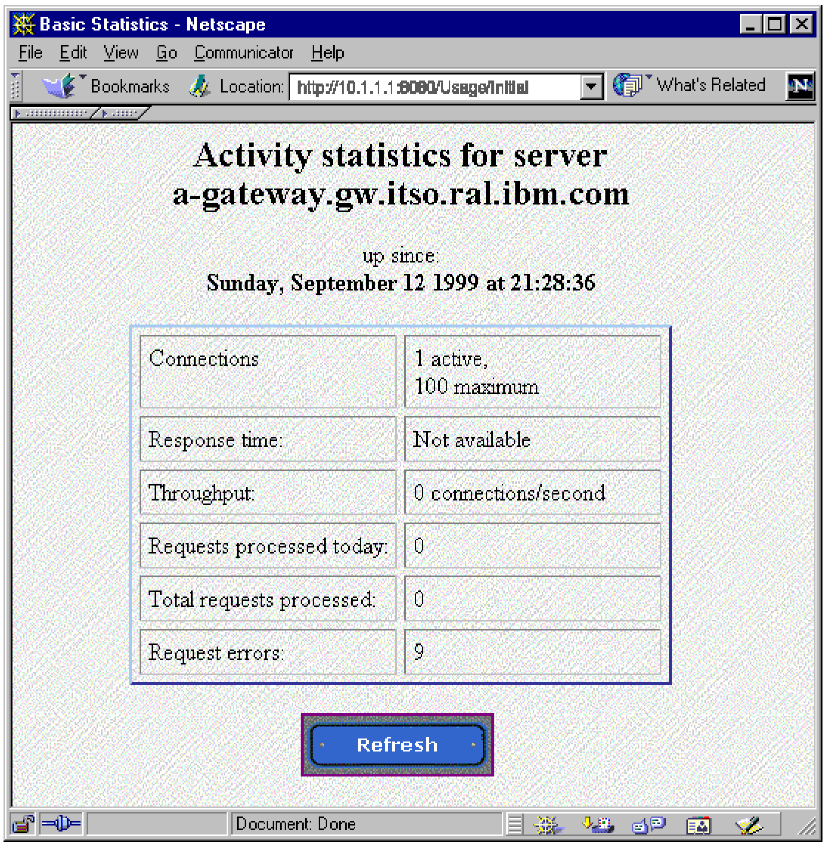

Figure 62. Basic statistics

This can be useful to determine you need to increase the number of threads or if excessive errors are occurring.

6.1.4.4 Trace

A trace function is available on HTTP proxy. This produces a detailed trace of execution of the HTTP proxy internal workings.

The output can only be interpreted by developers and your service organization for the purposes of problem determination. This is a last resort tool when all other methods have not resolved your problem.

If requested, follow these steps to gather a trace:

1. Ensure you have sufficient disk space to record a trace.

If you have created a dedicated logging filesystem, /var/log, you are ready to proceed.

2. Try to find a quiet time on the system to limit the amount of traffic going through the firewall.

   The output is quite large and limiting the traffic will make it easier to sift through the volumes of data to find the cause of your problem. It will cause problems if your trace file becomes too large to handle electronically via e-mail or FTP.

3. Append the following path name to the HTTP proxy URL:

   ```
   /admin-bin/trace
   ```

   If your HTTP proxy is running on the firewall at secure IP address 10.1.1.1 listening on TCP port 8080, then the full URL to activate the trace HTML is:

   http://10.1.1.1:8080/admin-bin/trace

   In the HTML screen that follows, you will need to specify a filename for the trace file and the subcomponents that you are interested in tracing:

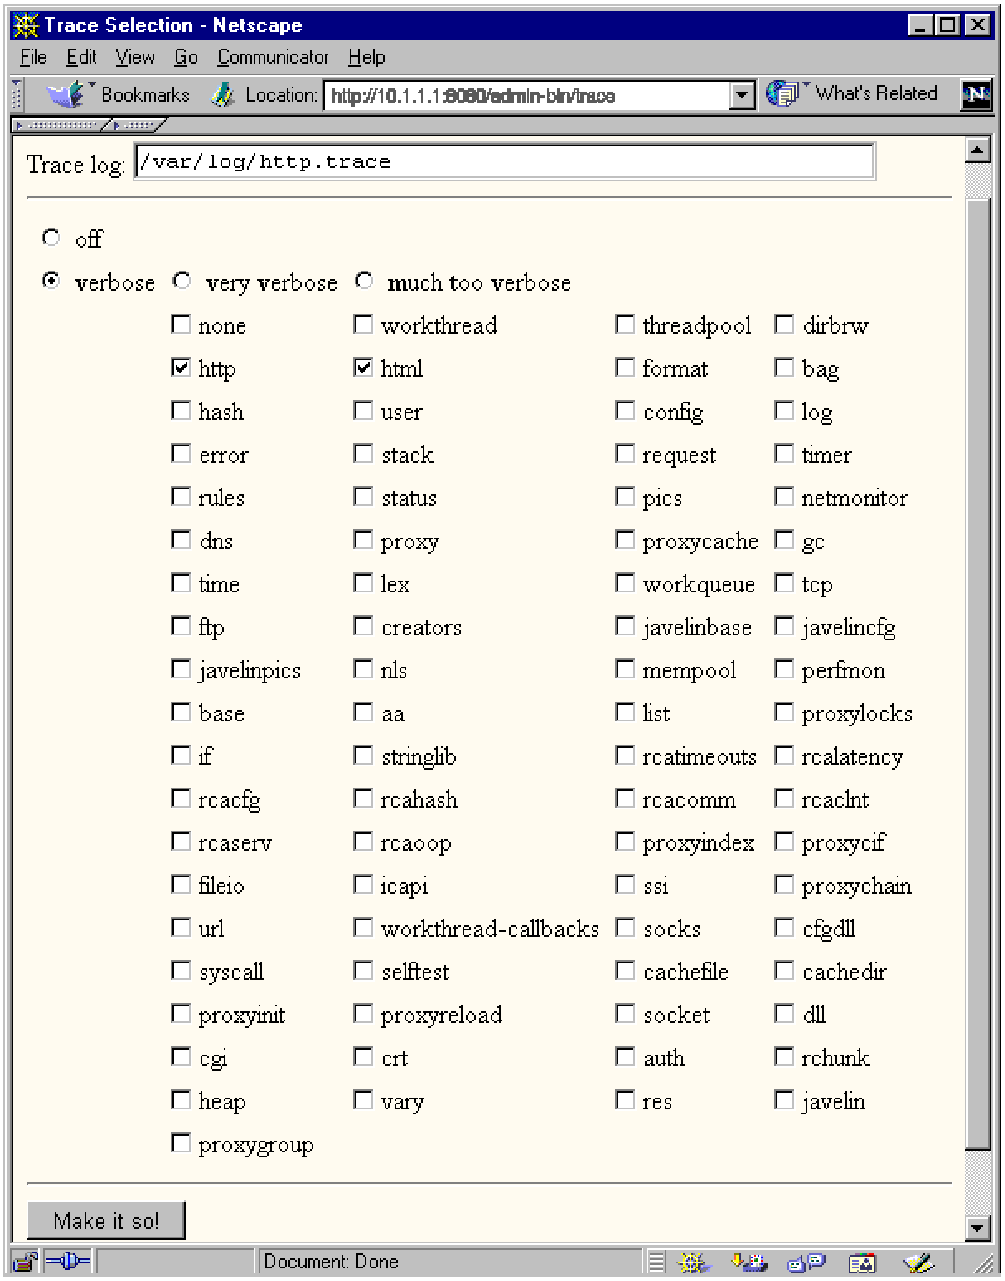

Figure 63. Enable trace

4. Selecting the "Make it so" button will start tracing.
5. At this point you will need to recreate the problem you are trying to trace
6. When you are done, turn off the trace from the previous screen.
7. Review the trace file output.

Check to make sure the file is not empty and appears to contain the events associated with your problem.

For example the following trace fragment details the HTTP headers and contains the correct hostname we were trying to access `home.netscape.com`.

```
Host: home.netscape.com
Connection: Keep-Alive, TE
Referer: http://home.netscape.com/escapes/search/netsearch_6.html
Accept: image/gif, image/x-xbitmap, image/jpeg, image/pjpeg, image/png
Accept-Encoding: gzip
Accept-Language: en
Accept-Charset: iso-8859-1,*,utf-8
Cookie: UIDC=198.133.22.73:0933091929:851638;
NGUserID=cdbcf74f-25896-933091929-8
TE: chunked
Via: HTTP/1.0 rs600031.itso.ral.ibm.com (IBM-PROXY-FW)
User-Agent: Mozilla/4.51 [en] (WinNT; U)

[00001415/935182175]: HTTP Proxy.. leaving socket 15 to host
home.netscape.com connected
```

Figure 64. A trace file /var/log/http.trace

8. Compress the tracefile then and it to your service organization.

6.1.5 Client configuration

This section describes how the client browser environment is configured to access the HTTP proxy.

We primarily used the Netscape browser product Version 4.5 in our lab.

There are two ways to configure the browser:

- Manual
- Automatic

6.1.5.1 Manual

Manual configuration requires each user to update a table of supported protocols within the browser. This mode of operation is not designed with central administration in mind.

Follow these steps to manually configure the Netscape browser:

1. **Edit Netscape Preferences**.
2. Select **Advanced** from the menu.
3. Select **Proxies** from the menu.
4. Select **Manual Configuration**.

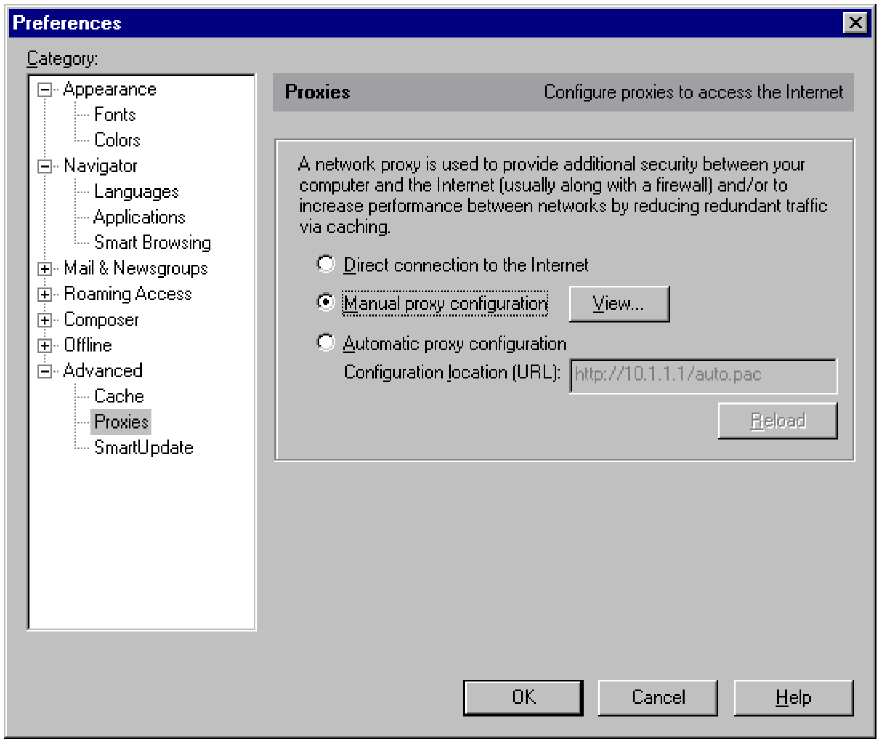

Figure 65. Select Manual proxy configuration

5. Enter the address of the firewall running HTTP proxy and the port number in all fields, except SOCKS.

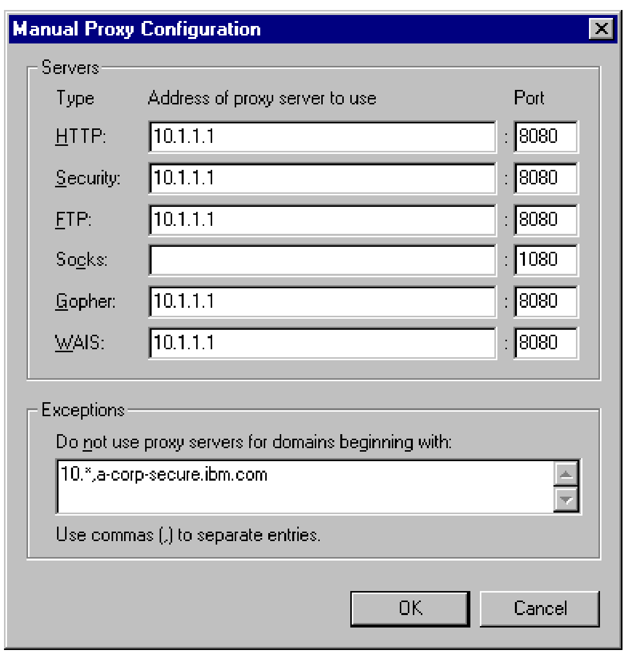

Figure 66. Enter address of HTTP proxy

Exceptions

The list of exceptions describes the networks should not go to the HTTP proxy. Enter the list of secure subnets by IP address or DNS domain name. This will prevent unwanted traffic bouncing off the firewall from the secure network back into the secure network.

6. Select **OK** to save the configuration.

6.1.5.2 Automatic

Automatic proxy configuration allows central administration of the client proxy configuration environment.

This provides the most flexible method to configure the client browser environment. This is achieved by use of a JavaScript function that is downloaded to the browser. The function examines the URL to be downloaded then calculates the proxy to service the request. This can be very sophisticated to include features such as load balancing and high availability.

To enable this feature, follow these steps:

1. The firewall administrator must first create the JavaScript function to be downloaded by the end user's browsers.

The following sample function will first check if the URL the user is accessing is either local or within the secure network. In that case the browser is instructed to go directly to the server. Next, the function checks to see what protocol the user is trying to access. If these are supported by the HTTP proxy, a list of HTTP proxy addresses is returned. If the first one is unavailable, the browser will automatically try to access the second backup site. Finally all remaining protocols are sent to the SOCKS server. This example demonstrates how specific protocols can return different proxy or SOCKS servers.

```
// Sample autoproxy pac file
function FindProxyForURL(url, host)
{
i = dnsResolve(host);
if (isInNet(i, "127.0.0.1", "255.255.255.255")
|| isInNet(i, "10.0.0.0", "255.0.0.0"))
{
return "DIRECT";
}
// need to go to proxy
if ( (url.substring(0,5) == "http:" ||
url.substring(0,4) == "ftp:" ||
url.substring(0,6) == "https:" ||
url.substring(0,7) == "gopher:") )
{
        return "PROXY 10.1.1.1:8080; PROXY 10.1.1.254:8080";
} else { //
// protocol not handled by HTTP proxy
return "SOCKS 10.1.1.1:1080";
}
} // FindproxyForURL
```

Figure 67. Sample autoproxy JavaScript function

2. The JavaScript function must be stored on an HTTP server available to the end users who will download this function.

We used another Web Traffic Express Version 2.0 server to act as an HTTP server. On this server we needed to configure two options in the configuration file /etc/ibmproxy.conf. The first option is the location of the pacfiles on the server expressed as an absolute path name:

```
PacFilePath   /usr/lpp/internet/server_root/pub/pacfiles
```

Next, we needed to make sure the HTTP server understands the .pac extension will correspond to the MIME type application/x-ns-proxy-autoconfig:

```
AddType  .pac      application/x-ns-proxy-autoconfig binary  1.0
```

We have placed this function in a file on the server called auto.pac in the directory specified above as PacFilePath.

3. Configure the browser to use the autoproxy configuration file "auto.pac"

The following window illustrates selection of automatic proxy configuration.

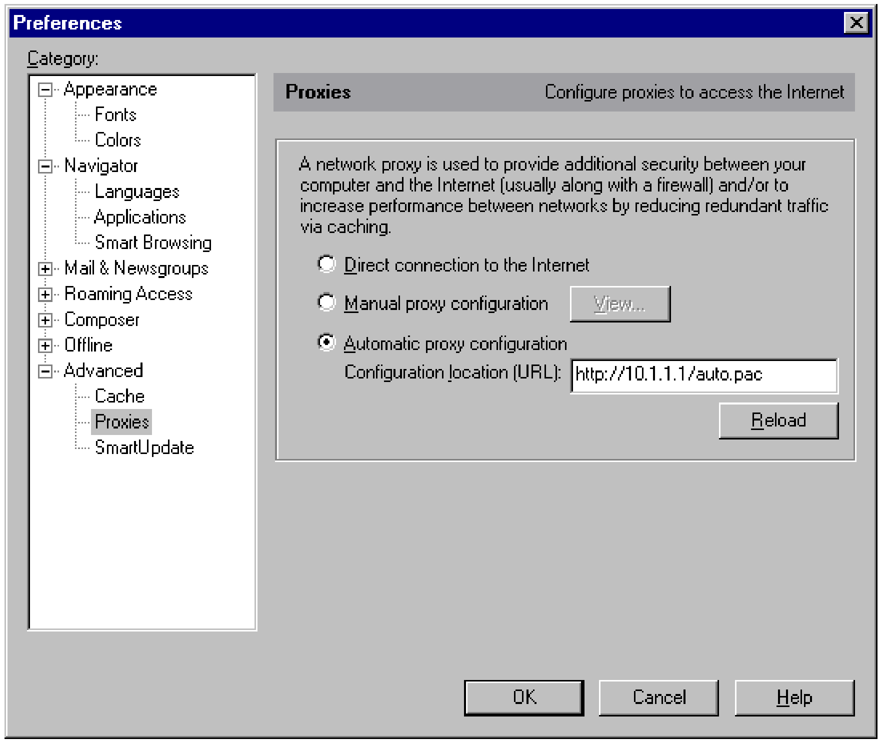

Figure 68. Automatic proxy configuration

In the event the administrator needs to modify the centrally configured auto.pac file, the end users will simply need to restart the browser to update the configuration.

6.1.5.3 Microsoft Internet Explorer support

All of the concepts in this section apply to Internet Explorer Version 5; however, Internet Explorer Version 4 can be configured in a very similarly fashion.

In order to configure Internet Explorer, click **Tools -> Internet Options...**. In Internet Explorer Version 4 you have to click **View -> Internet Options...**.

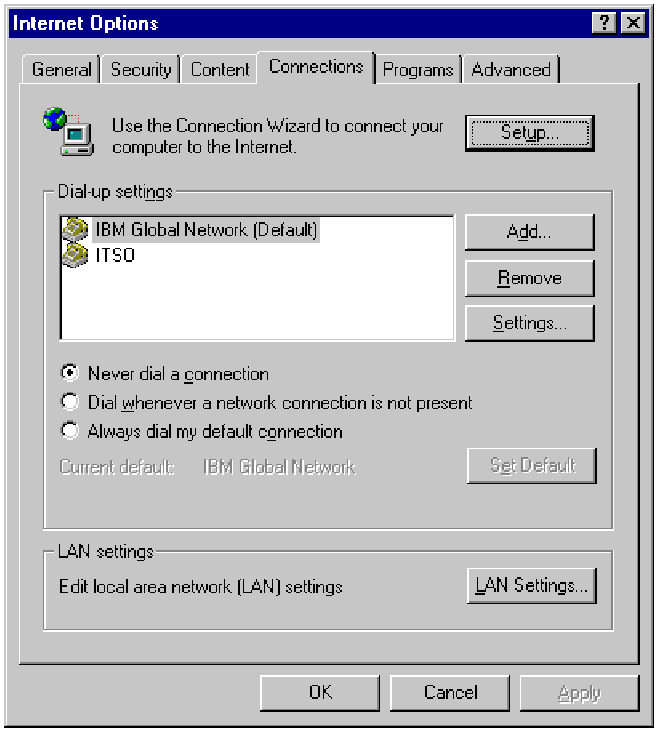

Figure 69. Internet Explorer Internet Options

You will now click the **Connections** tab and **LAN Settings**... You should now be in the window shown in Figure 70.

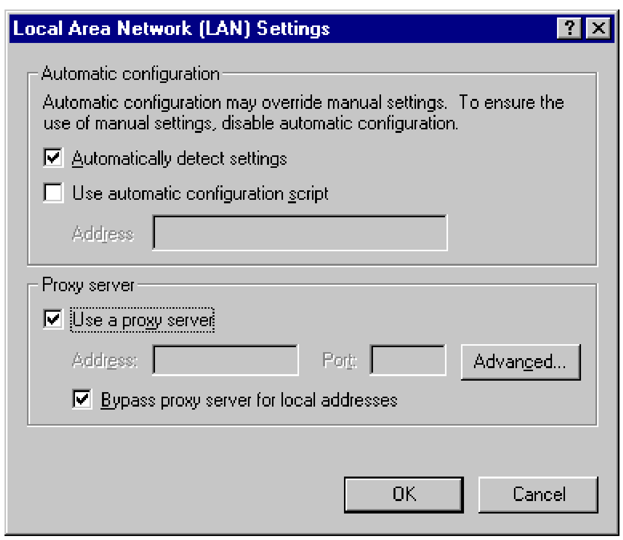

Figure 70. Internet Explorer LAN Settings window

Now select **Use a proxy server** and then click **Advanced**:

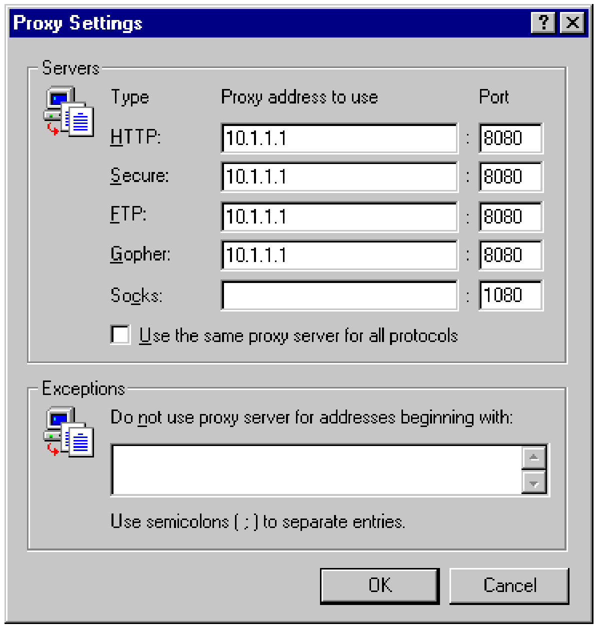

Figure 71. Internet Explorer Proxy Settings window

We filled in the fields as shown in Figure 71.

We found the following options not enabled by default in Internet Explorer Version 5.

- Use Web-based FTP
- Use HTTP 1.1 through proxy connections

We enabled both of these by clicking the **Advanced** tab shown in Figure 69:

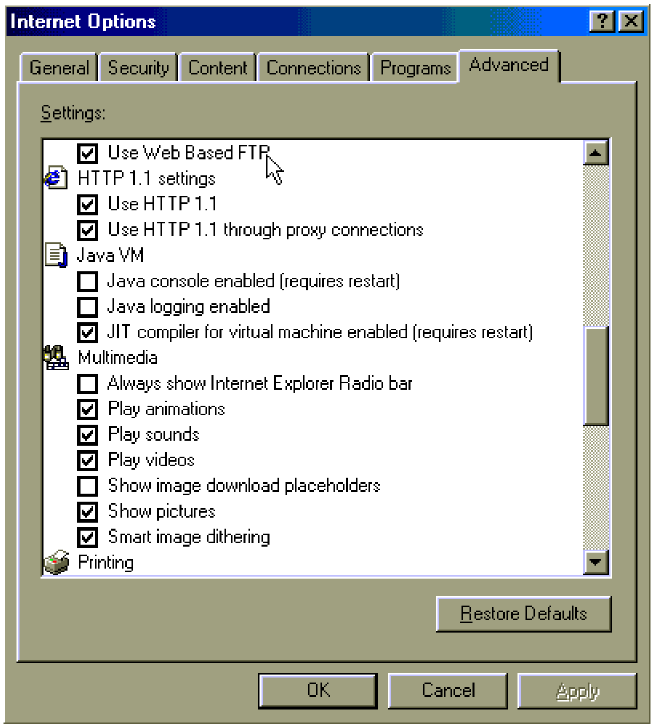

Figure 72. Enable FTP and HTTP 1.1 proxy

Automatic HTTP proxy configuration

Use the following procedure to enable automatic proxy configuration for the Internet Explorer Version 5 browser:

1. Select **Internet Options ...** from the Tools menu bar
2. Select the **Connections** tab
3. Select **LAN Settings**
4. Select **Use automatic configuration script** and deselect all other settings
5. Fill out the Address field as shown in Figure 73:

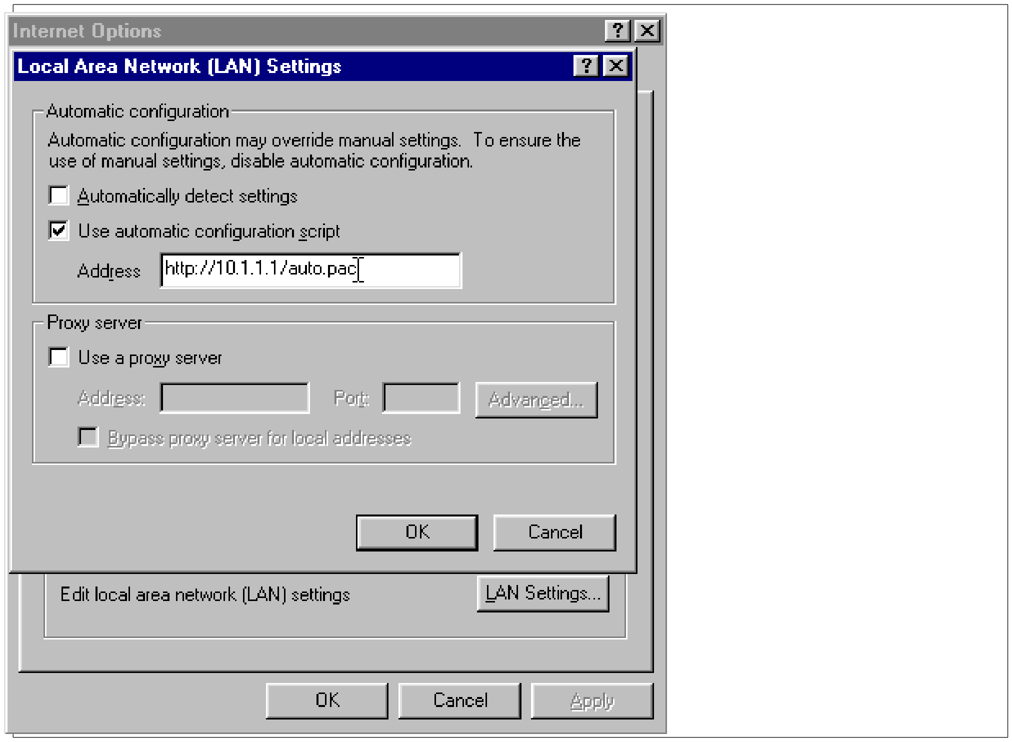

Figure 73. Automatic proxy settings for Internet Explorer Version 5

6.2 Telnet and FTP proxy

Additional TELNET and FTP proxies are available for clients that cannot be SOCKSified or cannot use a Web browser environment such as Netscape.

The proxy servers for Telnet and FTP replace the original Telnet and FTP servers supplied in the operating system.

The following excerpt from the file /etc/inetd.conf illustrates the new replacement daemons:

```
ftp stream tcp nowait root /usr/sbin/pftpd pftpd -ns
telnet stream tcp nowait root /usr/sbin/ptelnetd ptelnetd
```

Disabling the firewall-supplied FTP to use the original FTP daemon in AIX is not recommended.

This original ftpd server does not contain the authentication routines used in the firewall and allows any password to be used.

The Telnet and FTP proxies each operate in one of two modes:

- Authenticated
- Transparent

Authenticated users

From a security policy point of view, adding nonadministrative users to the firewall is discouraged. For this reason we recommend using transparent FTP, unless your security policy requires authentication.

Authenticated proxy users access the firewall in a restricted shell environment. They log in to the firewall operating system itself. It is then an additional task to ensure this restricted environment itself is secure, in addition to network security.

Users on the firewall also need maintenance:

- Remove unused accounts
- Reset forgotten passwords

6.2.1 Authenticated

To use the Telnet or FTP proxies in authenticated mode follow these steps:

1. Create proxy users on firewall
2. Clients login to firewall in a restricted shell environment

6.2.1.1 Creating proxy users

Creating users is described “Basic configuration” on page 19. The following image shows the parameters we have used in our example:

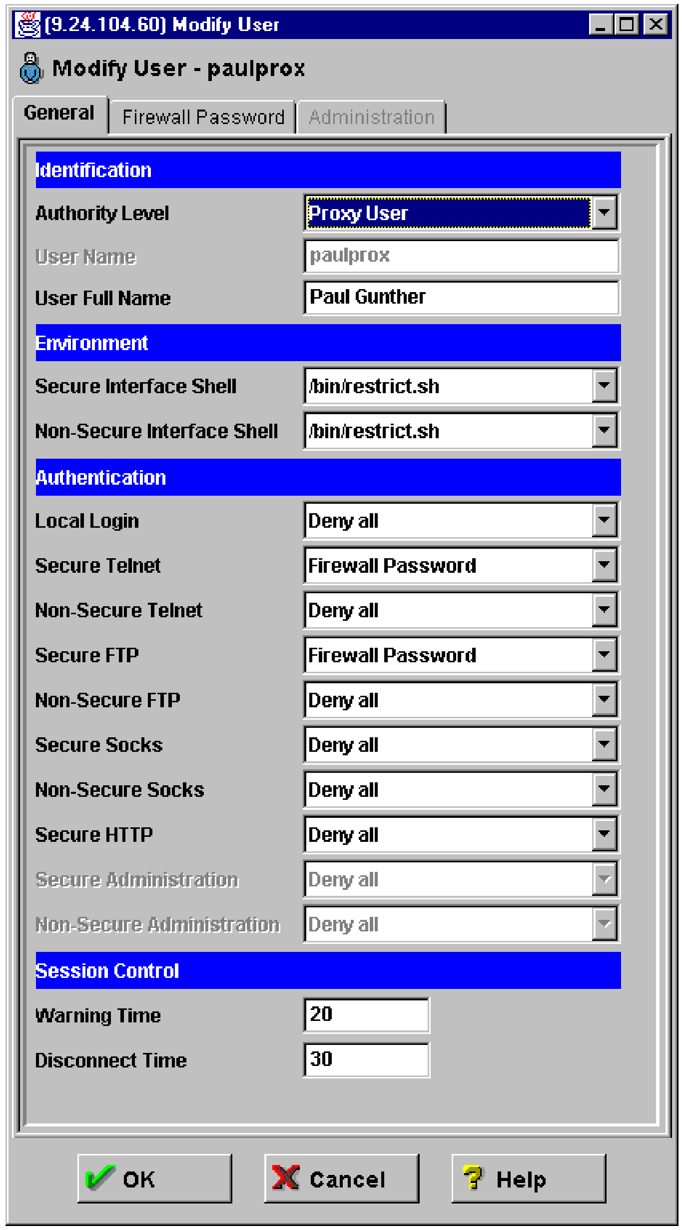

Figure 74. Adding a proxy user

6.2.1.2 Telnet

Using the telnet proxy involves first logging onto the firewall. At this point you will be within a restricted shell environment. Only a limited set of commands are available. From here you can telnet to another host beyond the firewall:

```
root@a-mail.a-corp-secure.ibm.com:/ > telnet 10.1.1.1
Trying...
Connected to 10.1.1.1.
Escape character is '^T'.


 telnet (a-gateway.gw.itso.ral.ibm.com)

Access Restricted
login:test
Password:
Only ITSO Firewall V4 Residents should Use this machine,
according to our Security policy
Last unsuccessful login: Tue Sep 14 03:06:45 EST 1999 on /dev/pts/5
Last login: Tue Sep 14 03:07:52 EST 1999 on /dev/pts/5

Currently logged in from 10.1.1.2
a-gateway.gw.itso.ral.ibm.com: help
Available commands are: passwd,telnet,tn,tn3270,ping,pingroute,exit,logout
70,ping,pingroute,exit,logout
a-gateway.gw.itso.ral.ibm.com: telnet b-gateway
Trying...
Connected to b-gateway.gw.itso.ral.ibm.com.
Escape character is '^T'.


 telnet (b-gateway.gw.itso.ral.ibm.com)

login:
```

6.2.1.3 FTP

Using the ftp proxy involves logging onto the firewall. From here you need to specify the target host on the other side of the firewall by use of the `quote site` FTP sub command. Next, use the `user` FTP sub command to specify the user ID on the remote system. You will be prompted for the password. The following example demonstrates the use of the FTP proxy:

```
root@a-mail.a-corp-secure.ibm.com:/ > ftp 10.1.1.1
Connected to 10.1.1.1.
220  a-gateway.gw.itso.ral.ibm.com FTP GATEWAY (Version 1.2 12/06/94 22:49:31) r
eady.
Name (10.1.1.1:root): test
331  Password required for test.
Password:
230  To specify destination, type "quote site remote.host.com"
ftp> quote site power.au.ibm.com
220 power.aixisc.au.ibm.com FTP server (Version 4.1 Mon Mar 23 10:20:45 CST 1998
) ready.
ftp> user paulgun
331 Password required for paulgun.
Password:
230 User paulgun logged in.
ftp> quit
```

6.2.2 Transparent

Transparent mode does not require any users to be defined on the firewall. A modified syntax is used to signify that transparent mode is used. In both cases, at the login prompt a username@remotehost syntax is used instead of the usual username. This will instructed the TELNET or FTP proxy to forward the request to the named remote host.

6.2.2.1 Enable Transparent proxy

To enable transparent TELNET or FTP proxy, select **Security Policy** from the main window:

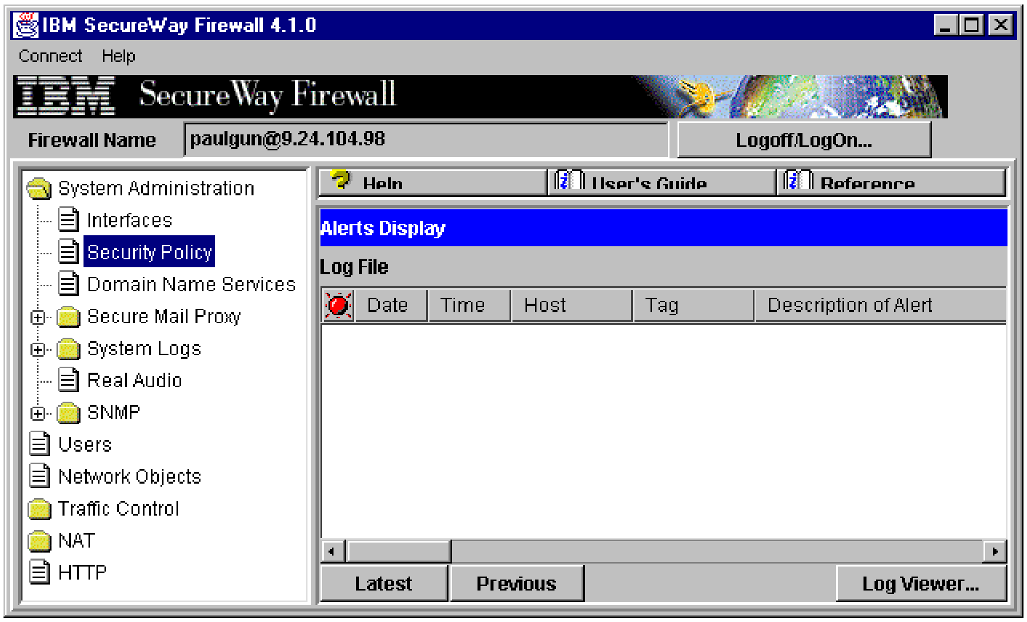

Figure 75. Select Security Policy

Enable the desired transparent proxies (Telnet and/or FTP) from the following window:

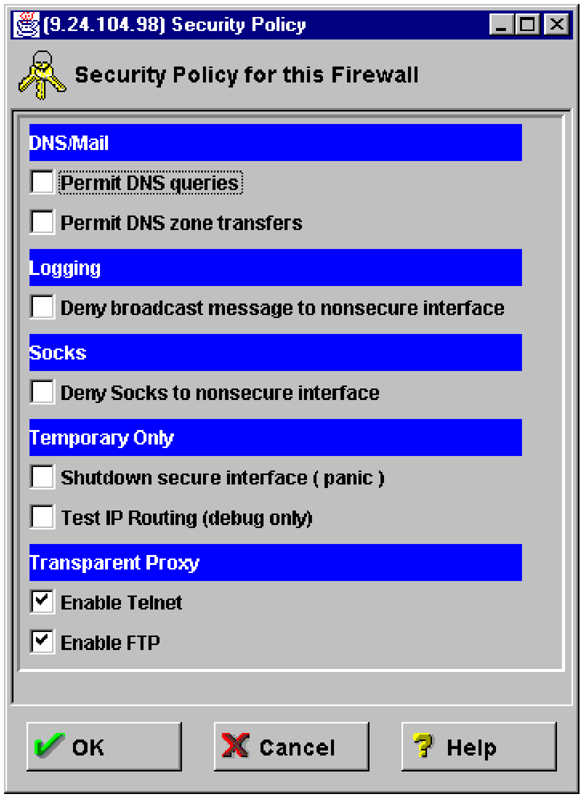

Figure 76. Enable Telnet and FTP transparent proxies

6.2.2.2 Telnet

With transparent telnet proxy enabled, the telnet server running on the firewall will look for a login name that includes the @ character to signify the request will be forwarded to the specified host. The following example of transparent telnet illustrates this point:

```
root@b-mail.b-corp-secure.ibm.com:/ > telnet gw
Trying...
Connected to gw.b-corp-secure.ibm.com.
Escape character is '^]'.


 telnet (b-gateway.gw.itso.ral.ibm.com)

login:paulgun@a-gateway.gw.itso.ral.ibm.com
Trying...
Connected to a-gateway.gw.itso.ral.ibm.com.
Escape character is '^]'.


 telnet (a-gateway.gw.itso.ral.ibm.com)

Access Restricted
login:paulgun
Password:
```

6.2.2.3 FTP

Like the Telnet proxy, the FTP proxy looks for the @ character in the login name to indicate the request will be forwarded to the indicated host. The following example demonstrates this:

```
root@b-mail.b-corp-secure.ibm.com:/ > ftp gw
Connected to gw.b-corp-secure.ibm.com.
220  b-gateway.gw.itso.ral.ibm.com FTP GATEWAY (Version 1.2 12/06/94 22:49:31) r
eady.
Name (gw:root): paul@proxy7.au.ibm.com
331 Password required.
Password:
230 User logged in.
ftp>
```

Chapter 7. 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.

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, 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.

The IBM SecureWay Firewall V4.1 for AIX provides the SOCKS server Version 5 protocol, which enables clients inside and outside the secure network to pass an authentication stage before accessing applications in the other network. The SOCKS protocol does not encrypt the data stream between the client and the server. Especially for inbound connections you should use additional encryption tools to protect the authentication and the data traffic.

The SOCKS V5 server also provides an authenticated generic proxy and the ability to proxy some UDP-based streaming audio and video protocols.

Here are a few examples of services that can be used through the SOCKS server:

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

Like the proxy servers, the SOCKS daemon is automatically started at boot time. In addition, a Watch Agent is provided to allow monitoring of the server. You can start this service manually if you wish (see 7.3, “Advanced configuration” on page 144for more information on the commands for managing the SOCKS server).

7.1 User authentication modes

IBM SecureWay Firewall V4.1 for AIX provides a smooth migration path in the form of three authentication modes so that customers can continue to use installed SOCKS V4 clients as they introduce SOCKS V5 clients (see Table 15 for more details).

Table 15. SOCKS authentication modes

Mode	Description	Tag
Permissive	The most permissive profile does not enable outbound authentication and permits any user, whether using a Version 4 or Version 5 client to connect. In this scenario, inbound connections are denied.	1
Intermediate	The migration profile allows SOCKS V4 users to pass unauthenticated, but requires SOCKS V5 users to authenticate. Inbound SOCKS V4 connections are required to authenticate. This is the default profile.	2
Strict	The most secure profile requires that all users use SOCKS V5 clients and provide valid authentication.	3

You cannot define the authentication mode with the configuration client. You have to edit a file (/etc/security/explode.cfg). Use the tag shown in Table 15 to indicate your desired profile. Figure 77 shows an example of the file explode.cfg.

```
# This file controls certain aspects of how the IBM Firewall
# constructs (explodes) configuration files.
#
# socks5profile
# Determines the "authentication profile" to be used for building
# the socks5 config file (socks5.conf)
# 1 = permissive mode, no authentication is performed
# 2 = intermediate (migration) mode, socks5 clients must provide
#   credentials, but socks4 clients may pass
# 3 = strict mode, all clients must authenticate (socks4 is
#   disallowed)
socks5profile=2
```

Figure 77. SOCKS user authentication profile explode.cfg

Be sure to regenerate the connection rules after editing this file to reflect any authentication changes (the daemon is automatically restarted).

SOCKS V5 clients can be authenticated by any of the supported authentication schemes (see 3.2.4.2, “Authentication methods” on page 31) except the user-supplied method. However, clients must support the Challenge-Response Authentication Method (CRAM) protocol in order to use strong authentication (see 7.3.2, “Modules” on page 146). Otherwise only user ID and password schemes can be applied.

7.2 Configuring SOCKS services

When the firewall is installed, the SOCKS server is enabled, but there are no connections defined. For SOCKS clients to use the SOCKS server, you must first configure SOCKS using the configuration client.

To set up SOCKS services on IBM SecureWay Firewall V4.1 for AIX, you need to:

1. Build a connection from the secure network to the firewall SOCKS server.
2. Build a connection from the firewall SOCKS server to the non-secure network.
3. Specify SOCKS server configuration for connections between the secure network and the non-secure network.

The first two steps are similar to configuring proxy connections. But the third step is unique for SOCKS connections. All SOCKS server configuration entries are placed in the file /etc/security/socks5.conf.

Currently only a subset of SOCKS V5 configuration is supported via the configuration client. The SOCKS V5 server supports a full SOCKS V5 configuration via file editing (see “Advanced configuration” on page 144 for details).

7.2.1 Connections

First you create a connection from the secure network to the firewall SOCKS server. Figure 78 shows the Configuration Client GUI window to build this connection.

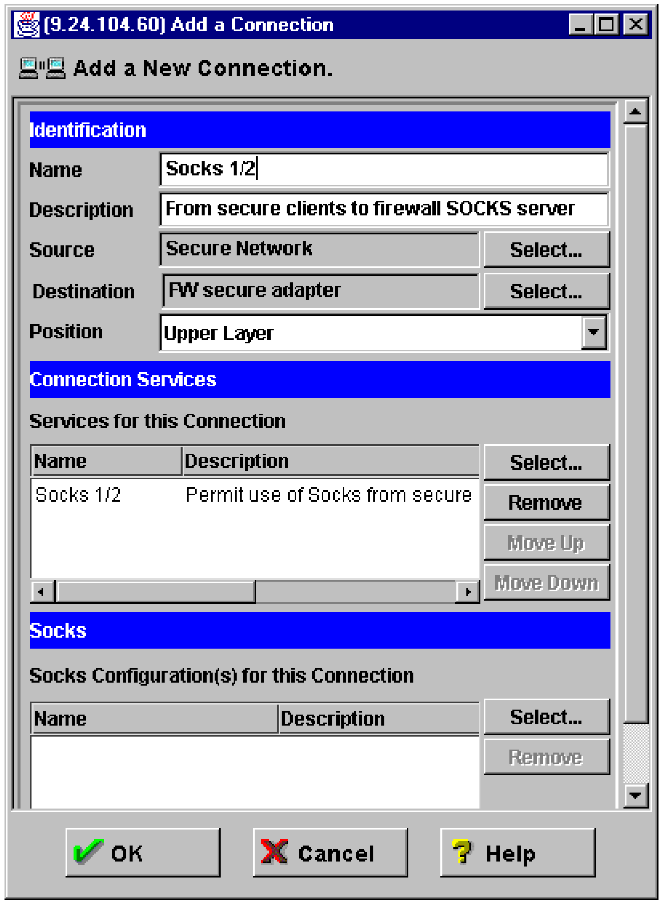

Figure 78. Connection from secure network to SOCKS server

In this example, the entire secure network is allowed to connect to the firewall SOCKS server on port 1080.

Next, you create a connection from the firewall SOCKS server to the non-secure network. Figure 79 shows the Configuration Client GUI window to build this SOCKS outbound connection.

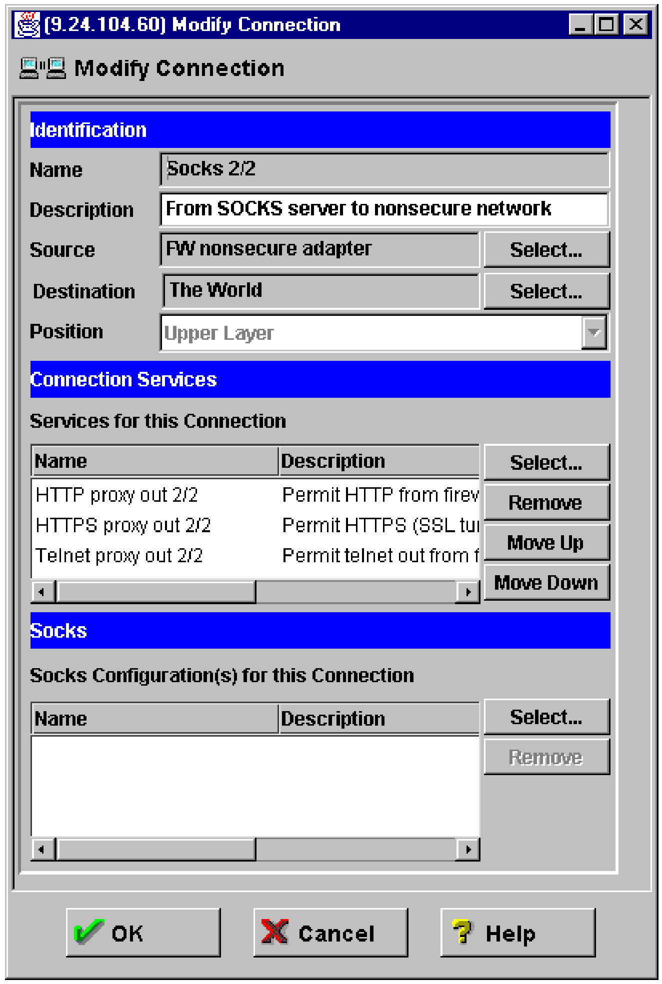

Figure 79. Connection from SOCKS server to non-secure network

In this example, the permitted outbound services are HTTP, HTTPS and telnet.

As the last step you specify a SOCKS server configuration for connections between the secure network and the non-secure network. Note that you build a connection with the secure network as the source object, the non-secure network as the destination object, and the SOCKS templates as the services. Figure 80 shows the Configuration Client GUI window for such a connection.

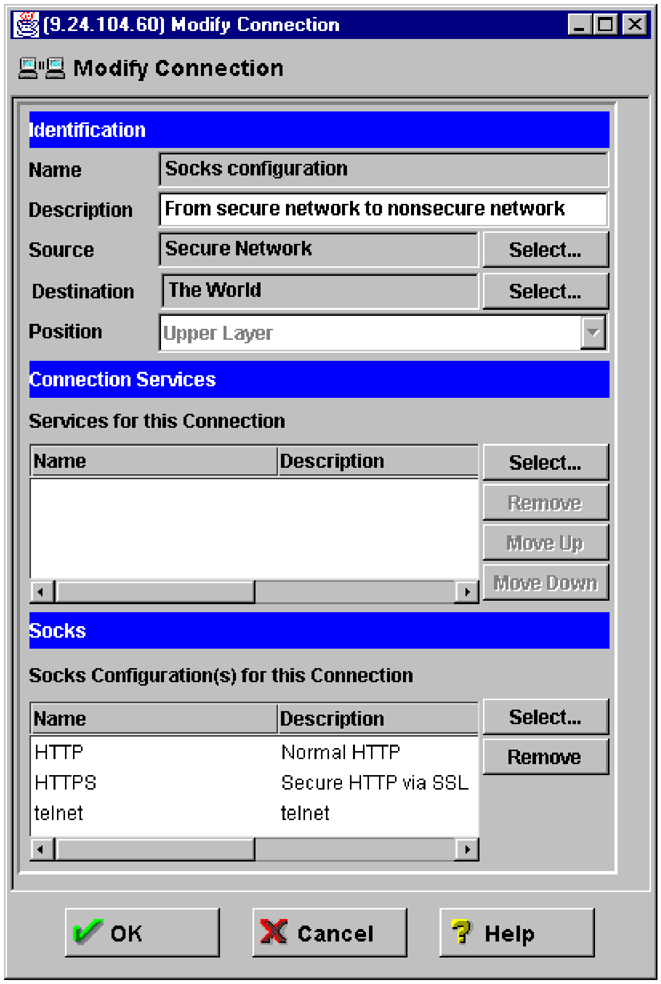

Figure 80. SOCKS configuration from secure to non-secure network

Click **Select** (on SOCKS panel in Figure 80) to get a list of predefined SOCKS templates. In this example, we selected three templates (HTTP, HTTPS and Telnet). The firewall provides several standard SOCKS templates.

You can expand the SOCKSified services by creating new templates. In the configuration client window, double-click the folder **Traffic Control**, then double-click **Connection Templates**. Inside this folder, double-click **SOCKS**. The window in Figure 81 is displayed.

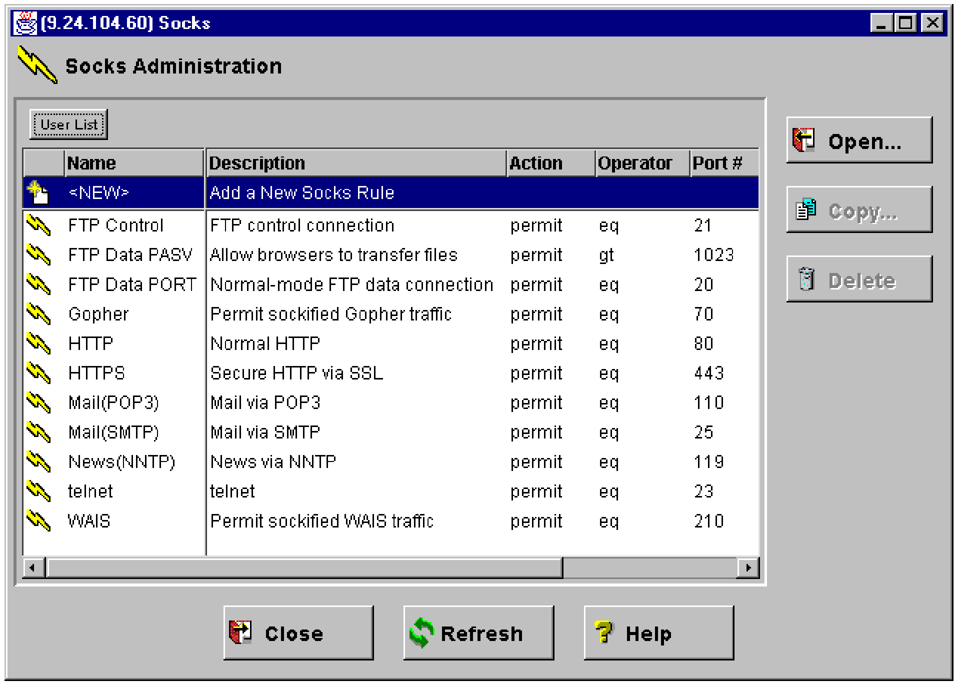

Figure 81. SOCKS Administration window

To create a new template, click **NEW** and then click **Open**. The following window is displayed.

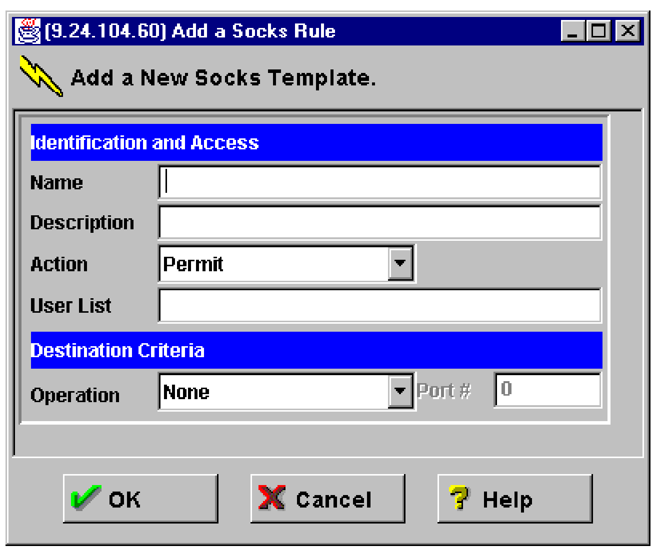

Figure 82. Adding a SOCKS Template

The fields to be entered are similar to those for the filter rule definition. The description for these fields are as follows:

Name Is the name of the template.

Description Describes the function of the template.

Action — Is the action to take if a session request matches the conditions in the filter definition. Possible values are *permit* (allow the session) or *deny* (refuse session establishment).

User List — Lists the user IDs which this configuration applies to. These are the IDs on the originating host, and they must be listed separated by commas and without blanks.

Operation — this is a logical operator code that represents the logical operation to be performed on the port number. Possible values are *Equal to, Not equal to, Less than, Greater than, Less than or equal to* or *Greater than or equal to*.

Port # — this is the number of a port. The port number is used with the Operation field to establish a relationship that must be met. For example, if you enter the Operation Greater Than and Port Number 23, then the port number must be greater than 23 for the rule to be invoked. If this pair is omitted, the line applies to all destination port numbers.

7.2.2 Filter rules

Just for completeness we show you here the filter rules generated from the configuration files for the SOCKS connections we configured in 7.2, “Configuring SOCKS services” on page 139.

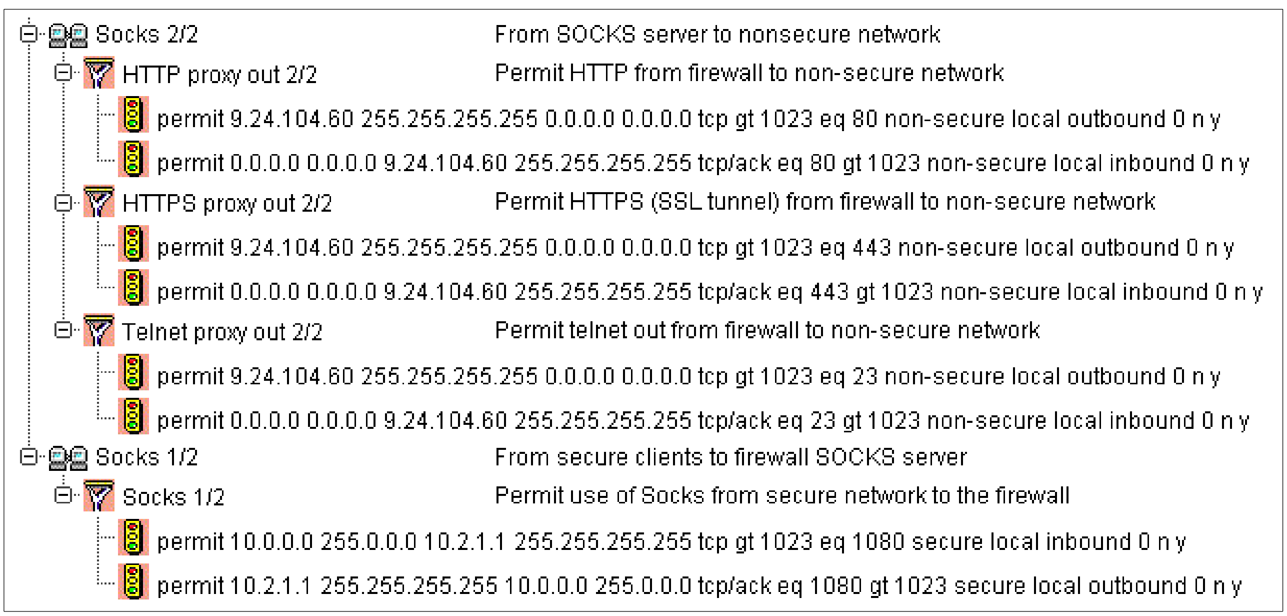

Figure 83. Filter rules for FTP, HTTP, and Telnet using SOCKS

The pair of rules on the bottom allows the SOCKS clients to contact the SOCKS server on port 1080. The other rules are protocol-specific and allow the firewall to contact the external servers. These are the same rules used for the proxy server.

7.3 Advanced configuration

The SOCKS configuration file is generated from Configuration Client GUI objects. Depending on your needs, you have three configuration options:

GUI Configuration You can do the most important configuration with the configuration client. No additional file editing is needed in this case. You get common defaults for all other configuration options.

Manual Configuration If you want to do all configuration by hand, edit the SOCKS V5 configuration file /etc/security/SOCKS5.conf. But you may never use the configuration client for SOCKS configuration because generating from GUI objects, all your handmade changes will be lost.

Mixed Configuration If you want to use both configuration methods, edit the file /etc/security/SOCKS5.header.conf. and do all other configuration with the Configuration Client. The SOCKS5.header.conf file is used as a header while generating the SOCKS5.conf file from the GUI objects.

The SOCKS V5 configuration file of the firewall contains the following sections:

- Variables
- Modules
- Routing
- Authentication
- Proxies
- Access Control

The next sections provide an explanation of the major settings. Please refer to the *IBM SecureWay Firewall Reference Version 4 Release 1, SC31-8418* and `http://www.aventail.com/` for a full SOCKS V5 configuration.

7.3.1 Variables

The following table shows the default values applied to the SOCKS server

Table 16. SOCKS V5 default behavioral values

Variable	Description	Default Value
SOCKS5_NOREVERSEMAP	Don't look up names to go with addresses (works faster this way).	1
SOCKS5_NOSERVICENAME	Don't look up service names to go with port numbers (works faster).	1
SOCKS5_NOIDENT	Disable identd requests.	1
SOCKS5_DEMAND_IDENT	Ensures that a user name is always associated with connection request.	0
SOCKS5_USECLIENTSPORT	Use client port (necessary for proxying streaming-UDP).	1
SOCKS5_BINDPORT	Change the inbound TCP port.	1080
SOCKS5_RECVFROMANYONE	Allow unsolicited UDP messages.	1
SOCKS5_MAXCHILD	Set maximum number of concurrent children.	64

The format for setting variables is as follows:

```
set variable value
```

Where `variable` is a variable from Table 16 and `value` is a defined value.

The `timeout` variable has another syntax, which is different from the SOCKS V5 standard, as follows:

```
timeout <time> minutes
```

An example of a timeout definition is as follows:

```
timeout 15 minutes
```

7.3.2 Modules

This section defines the modules used for authentication:

- server_password_IBM (ibmpwd)
- server_cram_IBM (ibmcram)

The first one is used for user ID/password authentication. The second one is used for strong authentication. It supports the Challenge Response Authentication Method (CRAM). The authentication method null is defined by default. You do not need additional filter rules for the authentication methods.

If you want to use your own authentication methods, you have to register your modules in this section. The module definition format is as follows:

```
module stub filename options
```

Where `stub` is a module dependent name prefix for accessing function names, `filename` is the path and name to the module, and `options` are the options for the module. The options can be omitted.

The following shows the predefined authentication modules from the firewall.

```
module server_password_IBM /usr/lib/ibm_gwauthp.mod
module server_cram_IBM /usr/lib/ibm_gwauthc.mod
```

7.3.3 Routing

The routing information is used to assign the network interface of multihomed hosts with a network number, a mask and, if needed, a range of ports. The format is as follows:

```
route dest-address dest-port interface-address
```

Where `dest-address` is an IP address and mask combination to specify a network, `dest-port` is either a single port or a range of ports, and `interface-address` is the IP address of a network interface.

You can use dashes (-) to indicate that there are no restrictions defined for that specific parameter.

The following is an example from our scenario. Note that we are using dashes to indicate that there are no restrictions for the destination port parameter.

```
route 10.2.1.1/255.255.255.0 - 10.2.1.1
route 9.24.104.60/255.255.255.0 - 9.24.104.60
```

7.3.4 Authentication

The contents of this section depend on the authentication profile set with the explode.cfg configuration file. If you use the permissive profile, no authentication methods are applied but connections from the non-secure network are denied. This results in ban entries for all non-secure interfaces.

If you use the intermediate profile, SOCKS V5 clients are authenticated while SOCKS V4 clients can establish outbound connections without authentication. This results in an additional null authentication method for connections coming from the secure network.

And if you use the strict profile, only SOCKS V5 clients can establish connections.

The format is as follows:

```
auth source-address source-port auth-methods
ban source-address source-port auth-methods
```

Where `source-address` is an IP address and mask combination to specify a network, `source-port` is either a single port or a range of ports, and `auth-methods` is a list of the defined authentication methods.

The following shows an example of the intermediate profile.

```
auth 10.2.1.1/255.255.255.0 - null,ibmpwd,ibmcram
auth 9.24.104.60/255.255.255.0 - ibmpwd,ibmcram
```

The list of authentication methods is applied backwards to clients that are requesting connections. This means the SOCKS server tries the Challenge Response Authentication Method first and then the username password authentication method. If both methods fail, the connection request is rejected for the `9.24.104.60` interface. Our secure interface is 10.2.1.1. Therefore, SOCKS V4 clients are allowed to use SOCKS without authentication. This is indicated with the null authentication method for this interface.

The dashes (`-`) indicate that there are no restrictions for the parameter.

7.3.5 Proxies

This section contains the proxy rules that tell the server how to connect to the destination. If you don't specify a proxy line, the destination is connected directly. The format is as follows:

```
proxy-type dest-host dest-port proxy-address proxy-port
```

Where `proxy-type` is the type of proxy server used (`socks4`, `socks5` or `noproxy`), `dest-host` is a combination of a network address and a mask, `dest-port` is either a single port or a range of ports, `proxy-addr` is the IP address of the proxy, and `proxy-port` is the port used by the proxy.

The following shows an example of a chained SOCKS server.

```
noproxy 10.2.1.0/255.255.255.0 - - -
socks5 0/0 - 9.37.3.60 1080
```

The dashes (-) indicate that there are no restrictions on the corresponding parameters. The first line ensures that no proxy is used for the local network. The second line forwards all other requests to the outer SOCKS server, which is a Version 5 server.

7.3.6 Access control

The last section contains the SOCKS rules. SOCKS rules can either permit or deny traffic. The format is as follows:

```
permit auth cmd src-host dest-host src-port dest-port user-list
deny auth cmd src-host dest-host src-port dest-port user-list
```

Where `auth` is a list of authentication methods, `cmd` is a command pattern, `src-host` and `dst-host` are combinations of a network number and a mask, `src-port` and `dst-port` are either a single port or a range of ports, and `userlist` is a list of users.

The following shows an example of SOCKSified HTTP.

```
permit - - 10.2.1.0/255.255.255.0 0/0 - 80 -
deny - - - - - - -
```

The dashes (-) indicate that there are no restrictions for the parameter. The first line allows all users from the network 10.2.1.0 to request connections from any port to port 80 on any host with any defined authentication method. The second line denies any other connection.

s5.conf

We have been talking up to now of the main configuration file, SOCKS5.conf, which is located in /etc/security. An intermediate configuration file, s5.conf, is created directly from SOCKS5.conf and is also located in /etc/security. Whereas user-specified rules are saved in SOCKS5.conf, the SOCKS server only reads the resulting s5.conf file.

The SOCKS5.conf is a file that is produced/created automatically upon activating the sockd rules. The SOCKS server however, does not read this configuration file directly. An intermediate configuration file is created from SOCKS5.conf which is then read by the SOCKS server upon boot-up or whenever it is refreshed. This file is /etc/security/s5.conf, and has a different format. Both SOCKS5.conf and s5.conf can be edited manually. However, we advise that only authorized expert firewall administrators do this, especially if there is no special reason to edit the files by hand. A utility program called "fwS5convert" is executed automatically to create s5.conf.

Regenerating filter rules overwrites SOCKS5.conf and s5.conf. You need to run "fwS5convert SOCKS5.conf s5.conf" whenever SOCKS5.conf is edited manually to produce s5.conf. As we mentioned earlier, whenever filter rules are regenerated through the GUI or command line, SOCKS5.conf will be overwritten.

The SOCKS5.conf configuration file is produced from three IBM Firewall configuration files through an explosion process. These files include explode.cfg, SOCKS5.header.cfg and sockd.conf. Therefore, anything that should be saved in between rule explosions should be placed in SOCKS5.header.conf file.

7.4 Chaining proxy and SOCKS

This example describes a chaining of a caching proxy server and the firewall SOCKS server (see Figure 84).

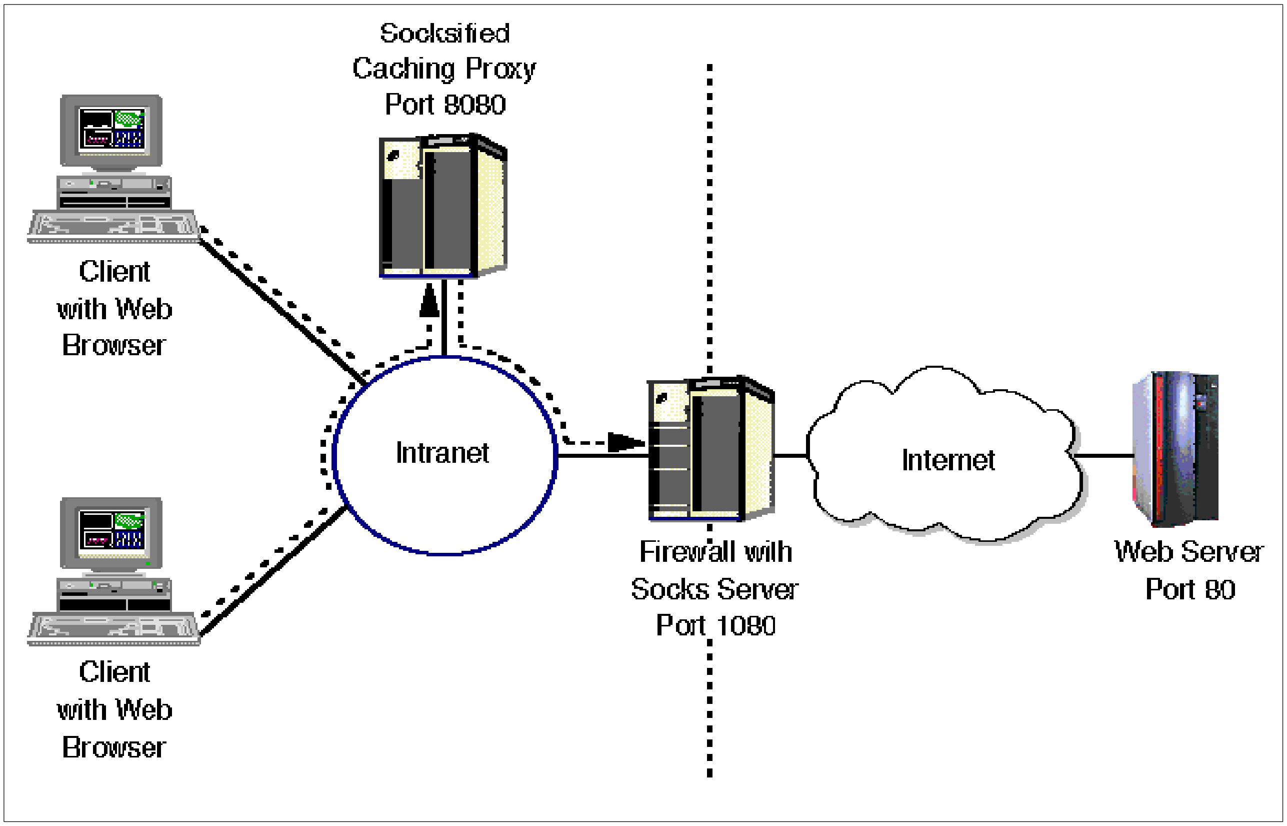

Figure 84. Proxy/SOCKS chaining example

We recommend that you use this way for Internet access. We use an internal Web proxy for the chaining proxy server, for example Web Traffic Express from IBM, to cache Web pages and FTP downloads. All internal clients point to this proxy. This proxy server must be SOCKSified to point to the firewall SOCKS server.

7.4.1 Chaining SOCKS Servers

You also can chain SOCKS servers, which means to SOCKSifying a SOCKS server. See *IBM SecureWay Firewall for AIX User's Guide Version 4 Release 1, GC31-8419* for a detailed description.

7.5 SOCKS client services

In order to make use of a SOCKS server, you need to have a modified SOCKSified client program that will direct the session to the SOCKS port on the server and handle the connect request/response sequence or you must SOCKSify the whole IP stack.

For general information about SOCKS, consult the SOCKS Web site at `http://www.socks.nec.com/`.

7.5.1 SOCKSified client programs

In general, Web browsers (such as the Netscape Navigator or the Microsoft Internet Explorer) provide built-in SOCKS support. SOCKSified versions of many other application clients are available from various Internet sites, for example:

Table 17. Sources for SOCKSified clients

Application	Information source
Netscape Communicator	`http://home.netscape.com/`
Microsoft Internet Explorer	`http://www.microsoft.com/ie/`
Notes 4.5 (and up) and Domino	Notes Documentation Database: Working with Lotus Notes and the Internet
IBM Web Traffic Express	`http://www.software.ibm.com/webservers/wte/`

7.5.2 SOCKSified IP stacks

Several manufacturers of TCP/IP implementations are incorporating SOCKS support into their products.

SOCKSified IP stacks are available from the following sources:

Table 18. Sources for SOCKSified IP stacks

Operating System	Applications	Product	Information source
AIX 4.3.3	Any TCP/IP application		Appendix 7.6.8, “AIX V4.3.3” on page 158
OS/2 Warp	Any TCP/IP application		Retrieve Software Updates from the Internet Connection Folder
OS/2 Warp		SOCKS5	`http://www.socks.nec.com/socks5.html`
Windows	Any WinSock Application	Aventail AutoSOCKS	`http://www.aventail.com/`
Windows	Any WinSock Application	Hummingbird SOCKS Client	`http://www.hummingbird.com/`
Windows	Any WinSock Application	NEC SOCKSCap	`http://www.socks.nec.com/sockscap.html`
UNIX	Shared Libraries	NEC runSOCKS	`http://www.socks.nec.com/how2socksify.html#runsocks`

7.5.3 SOCKSifying AS/400 clients

You can find information in the standard documentation about how to configure SOCKS support for AS/400. Log on to the AS/400 Information Center at `http://publib.boulder.ibm.com/pubs/html/as400/v4r4/ic2924/info/INFOCENT.HTM` and search for SOCKS.

You can find more information on this subject in the redbook *AS/400 Internet Security: IBM Firewall for AS/400,* SG24-2162 in Chapter 5 "Configuring SOCKS for the AS/400 system".

7.6 SOCKS connections example

In this section, we present you with some common connection examples for the SOCKS server. In the following examples, we show how to configure some clients to use the firewall as a SOCKS server. Note that some of the products listed here do not support SOCKS V5.

If you prefer, you can SOCKSify your stack (see “SOCKSified IP stacks” on page 150), so you do not need to perform any of these steps.

In order to use the SOCKS server, your users have to change their SOCKS pointer in their browsers to point to the firewall and the proper port.

SOCKS is protocol-independent, therefore you do not have to configure SOCKS in the browser for specific protocols such as HTTP or FTP.

See the configuration information for specific applications among the following. Most applications listed here provide only SOCKS V4. Refer to the documentation of each application to confirm which version of the SOCKS client they have.

7.6.1 Netscape Communicator Version 4

To configure Netscape Communicator to use a SOCKS server, click the main menu option **Edit** and then click **Preferences**. From the Category navigation open the menu **Advanced** and then click **Proxies**. Now select the radio button **Manual proxy configuration** and click **View**. The Manual Proxy Configuration window is displayed. Figure 85 shows a typical SOCKS configuration.

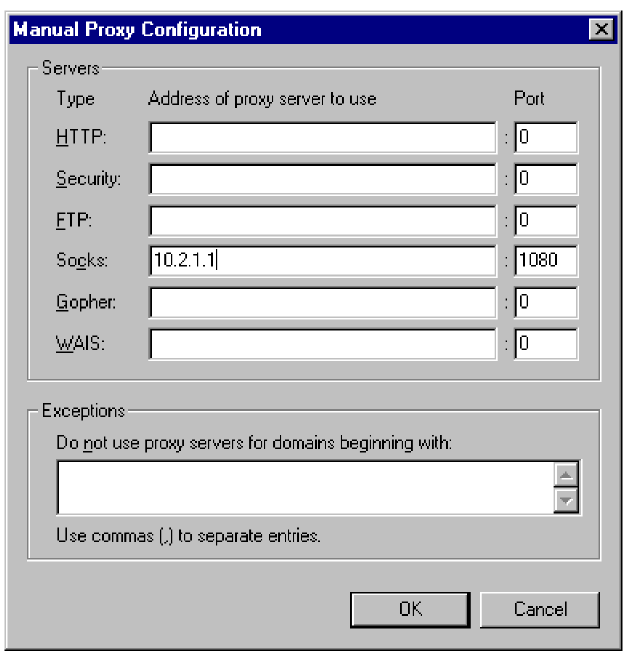

Figure 85. SOCKS configuration for Netscape Navigator

Put the hostname or IP address of the firewall in the SOCKS field and the respective port (in our example is 1080) in the Port field.

7.6.2 Microsoft Internet Explorer Version 5

Click the main menu option **Tools** and click **Internet Options**. The window Internet Options is displayed. Click the **Connection** tab. In the panel LAN settings click **LAN Settings**. The window in Figure 86 is displayed.

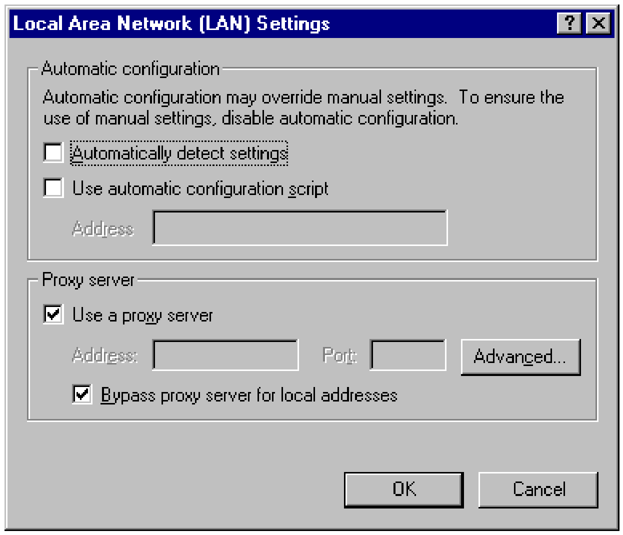

Figure 86. LAN settings window

In the Proxy Server panel, select the check boxes **Use a proxy server** and **Bypass proxy server for local addresses**, and click **Advanced**. The window in Figure 87 is displayed.

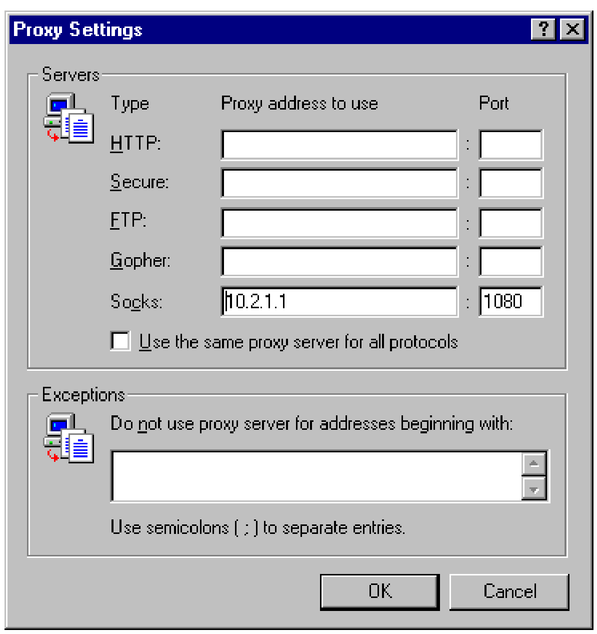

Figure 87. IE5 - Proxy settings window

Fill in the following fields in the Servers panel:

- SOCKS (in the column Proxy address to use): hostname or IP address of the firewall

• SOCKS (in the column Port): 1080

7.6.3 Microsoft Internet Explorer Version 4

Open the menu option **View** and click **Internet Options**. The Internet Options window is displayed. Click the **Connection** tab as shown in Figure 88.

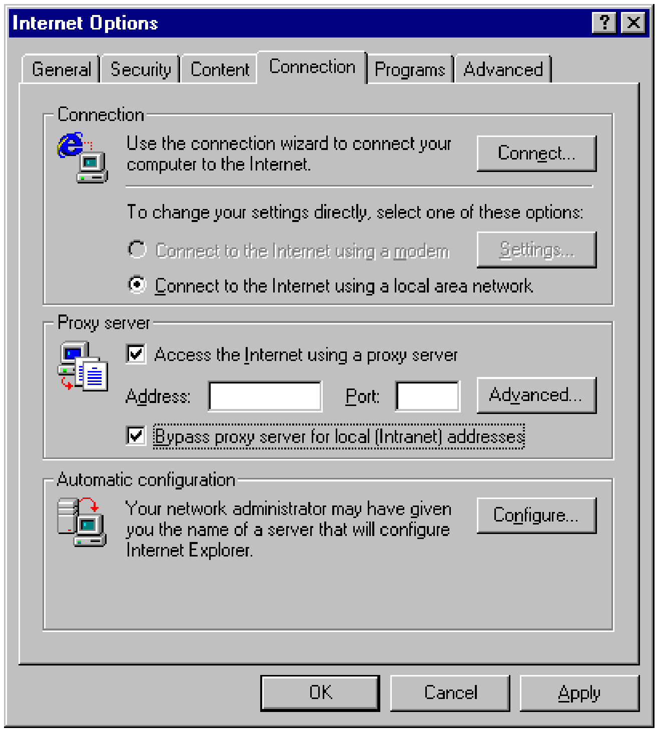

Figure 88. Internet options window

In the Connection panel select the radio button **Connect to the Internet using a local area network**. In the Proxy Server panel, select the check boxes **Access the Internet using a proxy server** and **Bypass proxy server for local (Intranet) addresses**, and click **Advanced**. The window in Figure 89 is displayed.

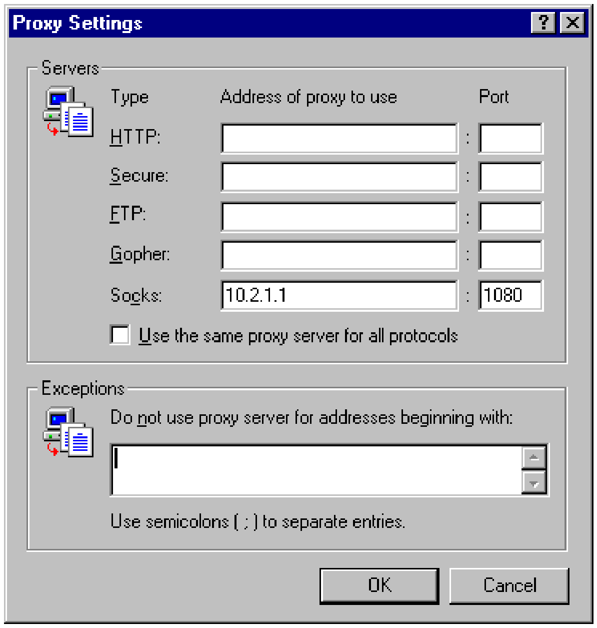

Figure 89. Proxy settings window

Fill in the following fields in the Servers panel:

- SOCKS (in the column Address of proxy to use): hostname or IP address of the firewall
- SOCKS (in the column Port): 1080

7.6.4 Lotus Notes Client Version 4

Locate the menu on the bottom right of your window as shown in Figure .

L

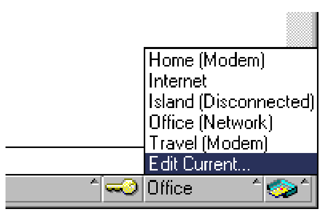

Figure 90. Lotus client menu

Click **Edit Current**, and the Location panel will be opened inside your client window. Locate the field **Web proxy**, and click the button shown in Figure 91.

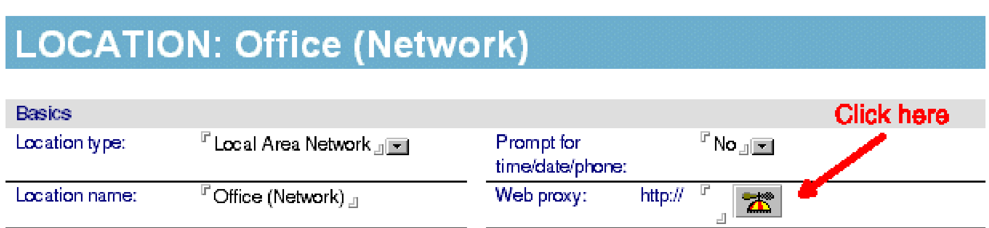

Figure 91. Location panel

The window shown in Figure 92 is displayed.

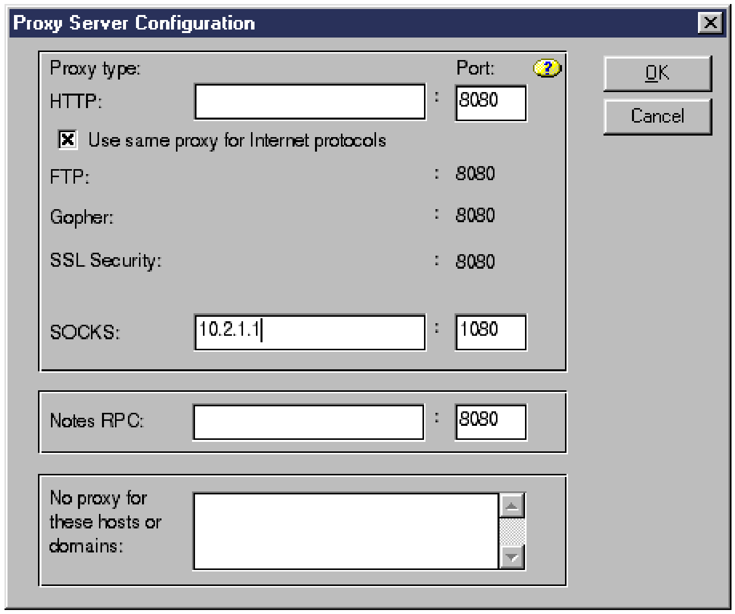

Figure 92. Proxy server configuration

Fill in the field SOCKS with the hostname or IP address of your firewall. If your SOCKS server is not listening on port 1080, change the Port field to the correct value.

7.6.5 NEC SOCKSCap V1

SOCKSCap can be used to SOCKSify applications that do not have a built-in SOCKS client. It can be downloaded from:

`http://www.socks.nec.com`

First you need to enter the information about the SOCKS server, and then you can add the clients you want to SOCKSify.

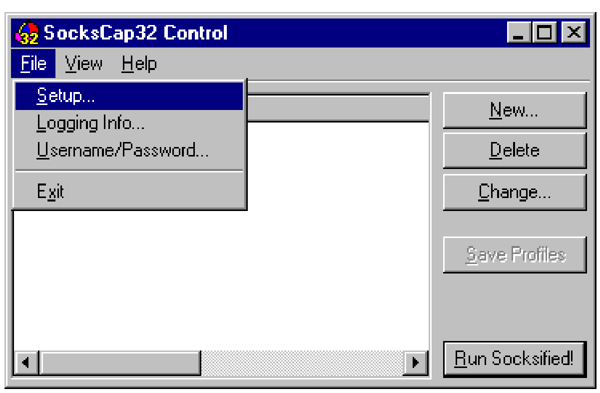

Figure 93. SOCKSCap main menu

In the SOCKSCap main window, click the menu option **File**, then click **Setup**. The window shown in Figure 94 is displayed.

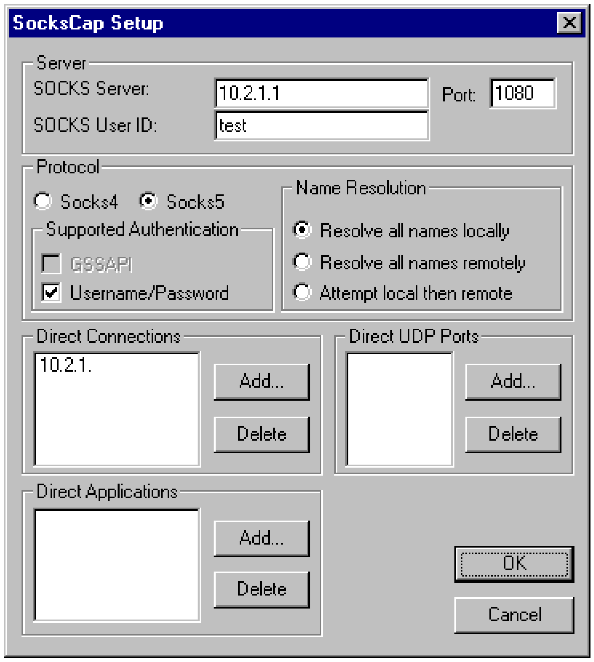

Figure 94. SOCKSCap setup window

Fill in the following fields:

- Server panel:

  SOCKS Server: hostname or IP address of the firewall

  Port: 1080

  SOCKS User ID: fill in with the respective user ID if you are doing user authentication

- Protocol panel:

  First, choose the version of the SOCKS server by selecting the radio button **SOCKS5**. If you are doing user authentication, select the check box **Username/Password**.

- If you want to avoid sending the Intranet requests to the firewall SOCKS server, add your internal network address in the list Direct Connections (for more information, see the help documentation of the product).

- If you selected the check box **Username/Password** in the setup window, which means you are doing user authentication at the SOCKS server, go back to the main window and click **File**, then click **Username/Password**. The window shown in Figure 95 is displayed.

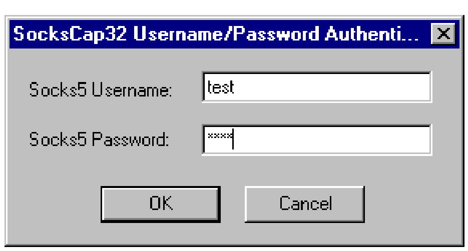

Figure 95. Username/Password input window

Enter the username and password for authentication in the firewall SOCKS server.

7.6.6 Hummingbird SOCKS client

The Hummingbird SOCKS client can be used to SOCKSify the entire TCP/IP stack of your machine. That means that any connection will be automatically redirected to the SOCKS server, even if the application you are using does not have a built-in SOCKS client.

You can find the guidelines for installing and configuring the Hummingbird SOCKS client at the following URL:

`http://www.hummingbird.com/products/socks/install.html`

Figure 96 shows an example of a basic configuration using a Hummingbird SOCKS client:

```
direct 10.0.0.0 255.0.0.0
sockd5 @=10.2.1.1 0.0.0.0 0.0.0.0
```

Figure 96. Example file c:\windows\system\socks.cnf

The `direct` line is used to indicate which addresses you can access without using the SOCKS server. You can use this keyword to specify your internal network address, so you do not have to use the SOCKS server to access internal addresses.

The `sockd5` line specifies the address of the SOCKS server (in this case, a SOCKS Version 5 server), and which addresses you want to access using the SOCKS server. This line indicates that all addresses should be directed to the SOCKS server (expect the `direct` line before the `sockd5` line).

7.6.7 ICQ

For detailed instructions on configuring the ICQ client to use a SOCKS V5 server, go to the following URL:

`http://www.icq.com/firewall/icqsocks5.html`

Figure 97 shows an example of the configuration using a SOCKS V5 server at the IP address 10.2.1.1:

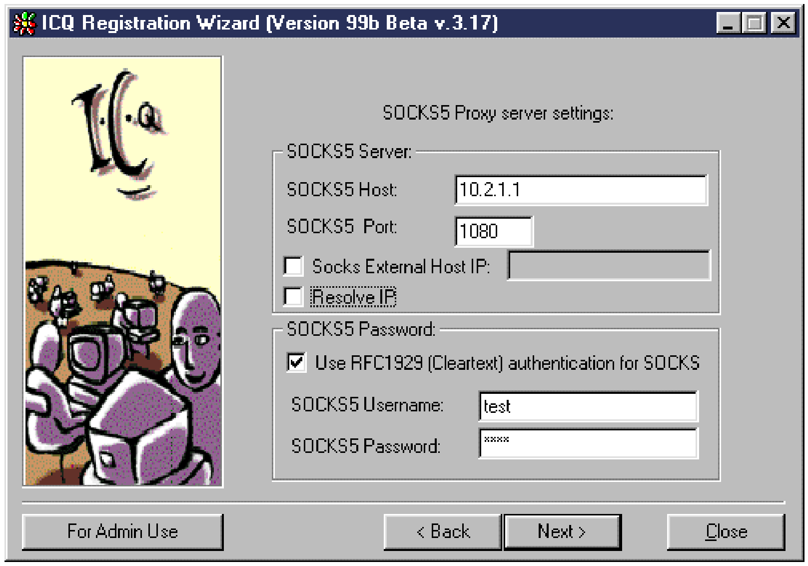

Figure 97. ICQ configuration window

7.6.8 AIX V4.3.3

AIX V4.3.3 contains SOCKS V5 client support built into the base operating system. This consists of a SOCKS client library API and an automatic SOCKSification feature. Automatic SOCKsification allows networked applications to use a SOCKS server running on the firewall without rewriting the code. This is similar to using a SOCKSified stack on a Microsoft Windows platform.

To configure an AIX 4.3.3 client as automatically SOCKSified, we set the AIX shell environment variable SOCKS5C_CONFIG to point to the configuration file /etc/socks5.conf. After setting this variable we could telnet to hosts on the non-secure network:

```
root@a-mail > export SOCKS5C_CONFIG=/etc/socks5c.conf
root@a-mail > telnet 9.24.104.60
Trying...
Connected to 9.24.104.60.
Escape character is '^T'.


 telnet (b-gateway.gw.itso.ral.ibm.com)

login:
```

The SOCKS client configuration file describes what subnetworks will pass through the specified SOCKS server. The following single line in /etc/socks5c.conf instructs the automatic SOCKSification feature to send requests that are destined for any host in the 9.0.0.0 Class A subnet to be SOCKSified and sent to the SOCKS server at IP address 10.1.1.1 on port 1080:

```
9.0.0.0/8 10.1.1.1:1080
```

7.6.9 AIX V4.3.2 and below

In AIX V4.3.2 and below the TCP/IP stack is not SOCKSified, so we must use client programs that have their own built-in SOCKS client.

Since most browsers and other Internet tools support SOCKS, we will show in this section how to use Telnet and FTP through a SOCKS server. In this case, the clients support only SOCKS V4.

First, create a file called /etc/socks.conf where you configure your SOCKS client, according to the following syntax:

```
direct secure-network network-mask
sockd @=firewall-ip-address destination-network network-mask
```

In our environment, we used the following configuration:

```
direct 10.2.1.0 255.255.255.0
sockd @=10.2.1.1 0.0.0.0 0.0.0.0
```

This means that any connection to hosts in the network 10.2.1.0 must be done directly, and for all other connections it must use the SOCKS server at the IP address 10.2.1.1.

After configuring this file, you need to download the commands for rtelnet and rftp (SOCKSified Telnet and FTP clients) from the following URL:

```
ftp://testcase.boulder.ibm.com/aix/fromibm/
```

When you need to connect to an external site, use rtelnet or rftp and the connection will be SOCKSified through the firewall.

7.6.10 RunSOCKS (for UNIX environments)

This application is similar to SOCKSCap (see 7.6.5, “NEC SOCKSCap V1” on page 155). It SOCKSifies applications that do not have built-in SOCKS clients. You can download this package from the following URL:

```
http://www.socks.nec.com
```

RunSOCKS is part of the SOCKS V5 product from NEC.

We downloaded this package, compiled it and installed it in a machine running Red Hat Linux V6 (kernel Version 2.2.10).

After installing, we created the file /etc/libsocks5.conf containing:

```
noproxy - 10.2.1. - -
socks5 - - - - 10.2.1.1
```

This means that all connections to the network 10.2.1.0 must be done directly without using the SOCKS server, and all other destinations must be redirected to the SOCKS server at the IP address 10.2.1.1.

Also, we declared the following environment variables:

```
SOCKS5_USER=test
SOCKS5_PASSWD=test
SOCKS5_SERVER=10.2.1.1
```

If you want to SOCKSify an application using RunSOCKS, you have to run it using the `runsocks` command:

```
# runsocks <command> <command parameters>
```

The following example shows how to SOCKSify a FTP connection (the ftp client is run using the command `ftp`):

```
# runsocks ftp ftp.software.ibm.com
```

For more information consult the online documentation provided in this package or check the URL:

`http://www.socks.nec.com/`

7.6.11 RealPlayer G2

In Real Player G2 Version 6.0.6.99, on top of SOCKSifying the application (for example, using SOCKSCap), you should manually configure the Real-Time Streaming Protocol (RTSP). Click on **Options**, then on **Preferences** and finally on the **Transport** tab:

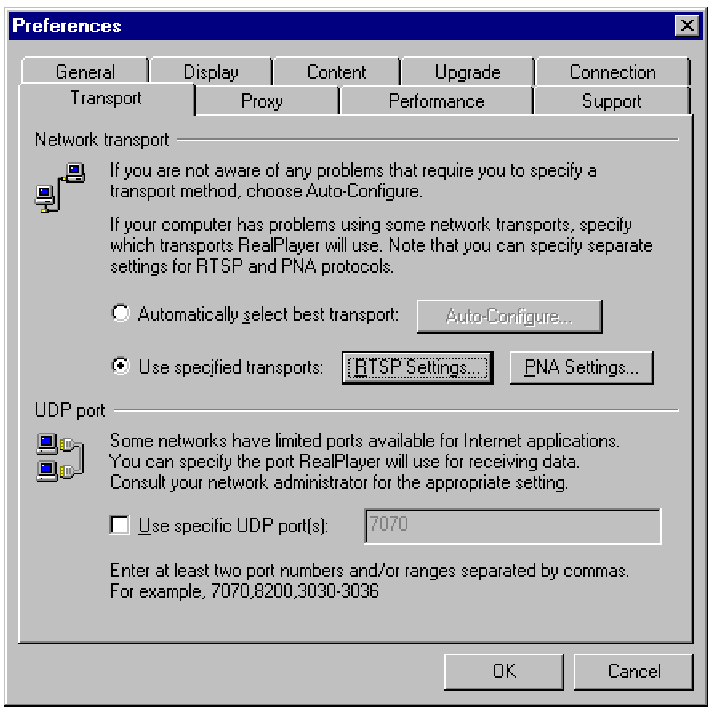

Figure 98. Configure RTSP transport

Using SOCKS we can only use TCP connection. UDP connections expect to be routed directly to the client system and cannot be used with the SOCKS protocol. The following figure illustrates configuration of the TCP-only transport:

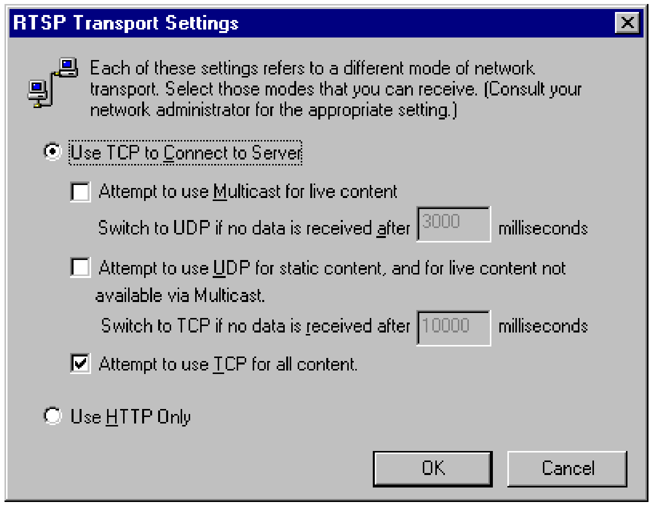

Figure 99. RTSP transport settings

In addition to transport settings we disabled all proxy settings as follows:

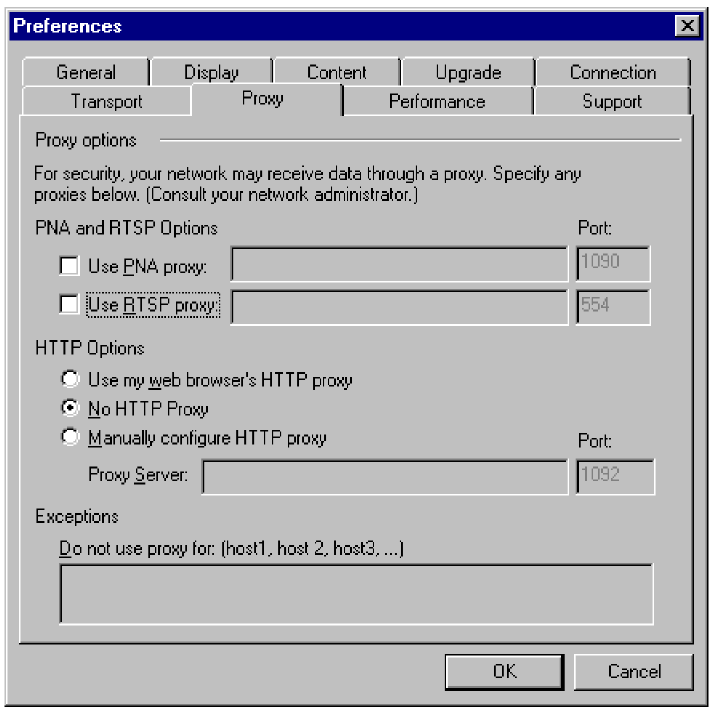

Figure 100. Disable all proxy settings

7.7 Using SOCKS traffic monitor

Along with the configuration client you receive a SOCKS Traffic Monitor application. This application allows you to monitor the current SOCKS traffic on an IBM SecureWay Firewall V4.1 for AIX or on an IBM SecureWay Firewall for NT in real-time. To use this application you must first start the SOCKS Monitor Service on the firewall you want to monitor.

Run the following command to start the SOCKS Monitor Service on your firewall:

```
# /etc/security/socks/bin/fwMonitor
```

Once the service is started you can start the Traffic Monitor[1] application on a remote Windows client. There is no AIX version of the Traffic Monitor, just a version for Windows NT/95/98. You must run this application on a PC. Do not forget to define a connection in the firewall for this traffic.

7.7.1 Connection for remote usage

To use the SOCKS Traffic Monitor from a remote PC running Windows NT/95/98, you have to define a connection from this machine to the firewall for the IP traffic.

Unfortunately there is no predefined service that we can use for that connection. Therefore, you will have to build this connection from scratch, starting with the two rules (see Chapter 4, “Packet filters” on page 43 for how to create rules and services).

The SOCKS Traffic Monitor uses port 5051/tcp. See *Guarding the Gates Using the IBM eNetwork Firewall V3.3 for Windows NT,* SG24-5209-01 for a detailed description of this process.

7.7.2 Starting the SOCKS traffic monitor from Windows NT

To start the application select **Start -> Programs** from the Windows NT NT/95/98 task bar. Select **IBM Firewall Client**. On the next menu select **SOCKS Monitor**. The window in Figure 101 is displayed.

[1] Traffic Monitor is the term used in the Configuration Client; in the firewall official documents it is called SOCKS5 Watcher.

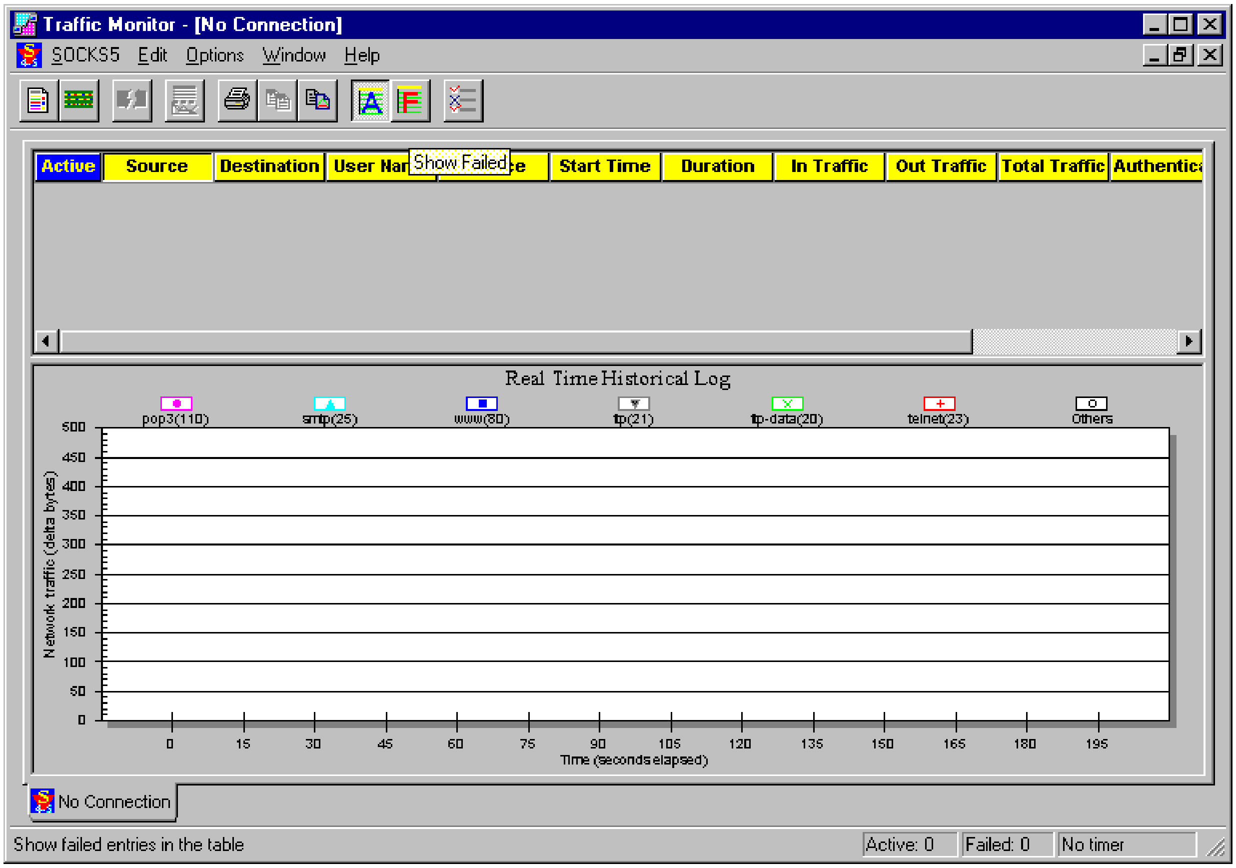

Figure 101. SOCKS traffic monitor without connection

From this window you can connect to any firewall with a SOCKS V5 server. To do so, select **SOCKS5** from the menu bar and **Connect** from the following pull-down menu. The application then asks you for the address of the SOCKS V5 server (see Figure 102).

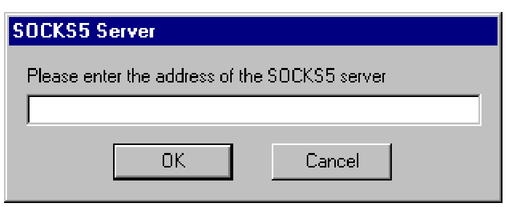

Figure 102. Connection to SOCKS V5 server

The window in Figure 103 is displayed when you enter the address and click the **OK** button. Please remember that the SOCKS Monitor Service must run on the firewall and you must have a valid IP filter to allow the traffic to connect to the SOCKS V5 server.

Figure 103 shows a sample output from the monitor.

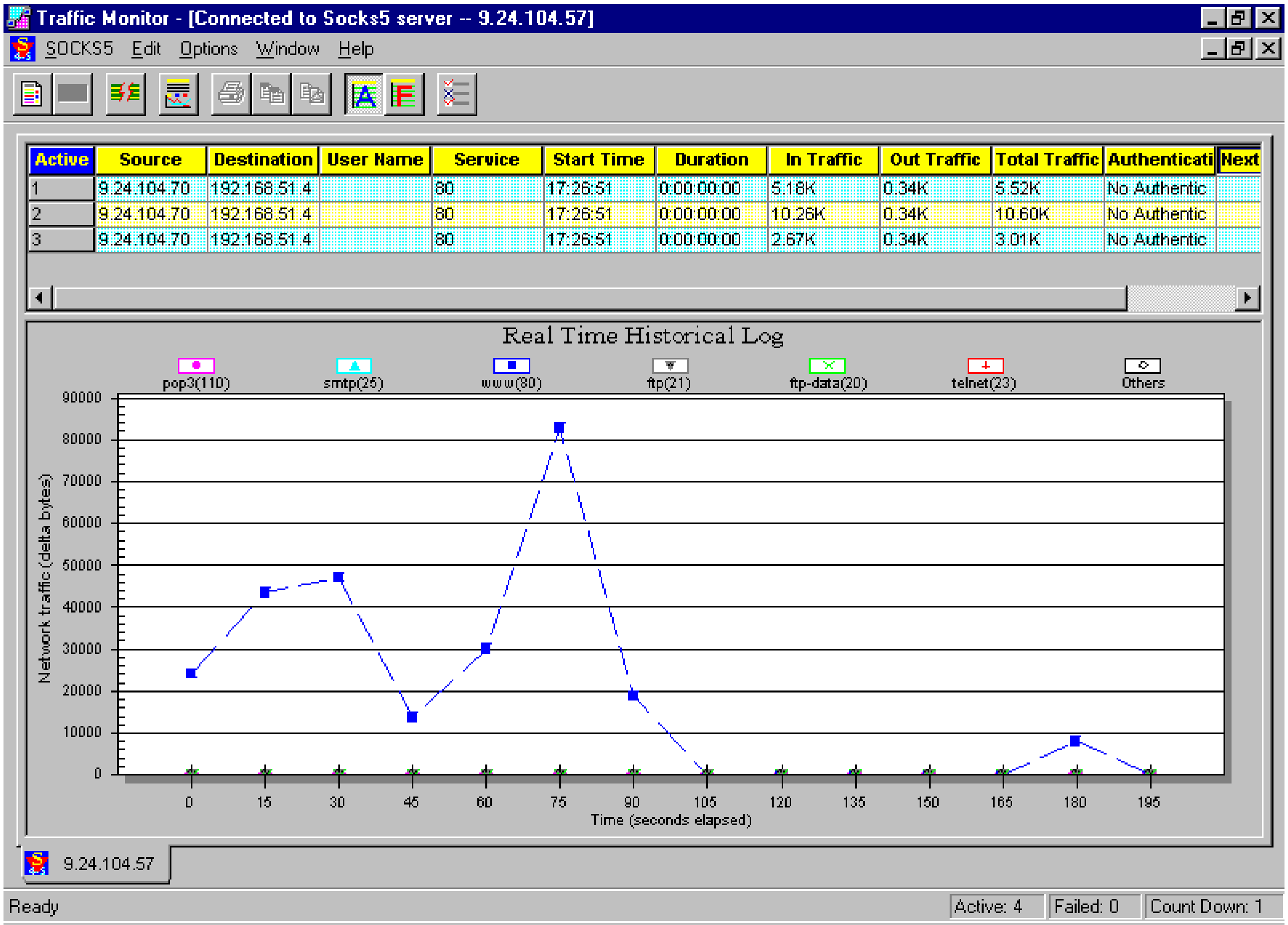

Figure 103. Sample traffic monitor output

The window is divided into two parts. The upper part shows you the current active sessions to the SOCKS V5 server with detailed data about the source and destination address, type of service, time, size, and authentication.

The lower part gives you a real-time historic log of the activity on the SOCKS V5 server. You see separate lines for every service monitored. The following services are monitored by default.

Table 19. Default monitored services

Type of service	Port
ftp	21
ftp-data	20
pop3	110
smtp	25
telnet	23
www	80
other	

You can also edit the monitored services.

7.7.3 Edit traffic monitor

To edit the properties of the SOCKS Traffic Monitor, select the **Options** entry from the Options menu. In the next window you have three tabs:

- Table
- Data
- Graph

Selecting the **Table** tab lets you edit the columns shown in the active connections table in the upper part of the main window.

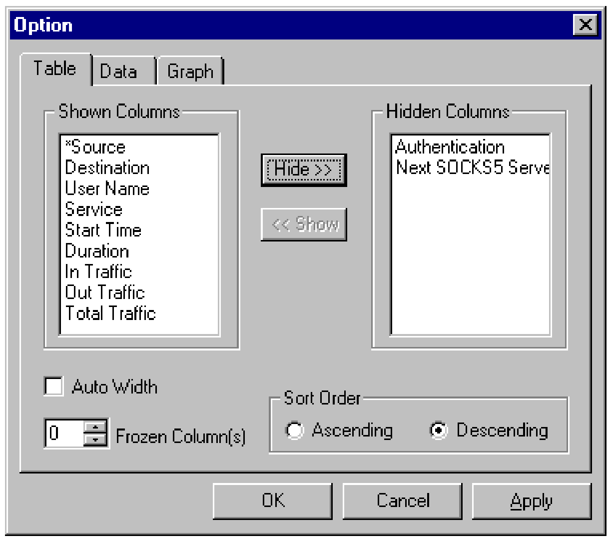

Figure 104. SOCKS monitor table settings

Choosing the **Data** tab lets you customize the data format shown in the active data connections table.

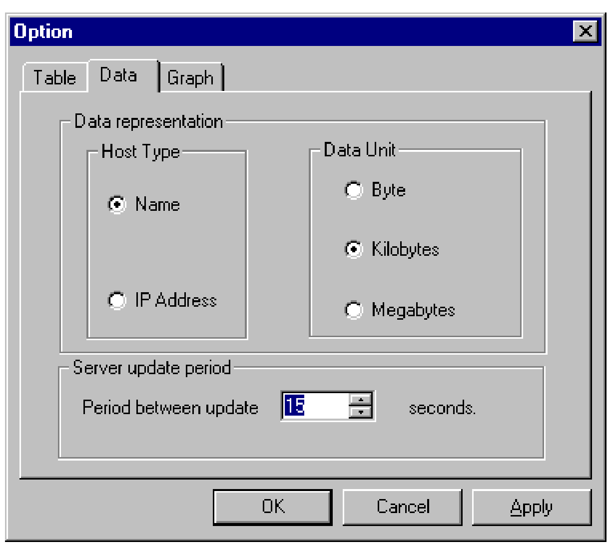

Figure 105. SOCKS monitor data settings

Selecting the **Graph** tab enables you to edit the services that should be monitored. You can add new ones, and edit and delete existing ones.

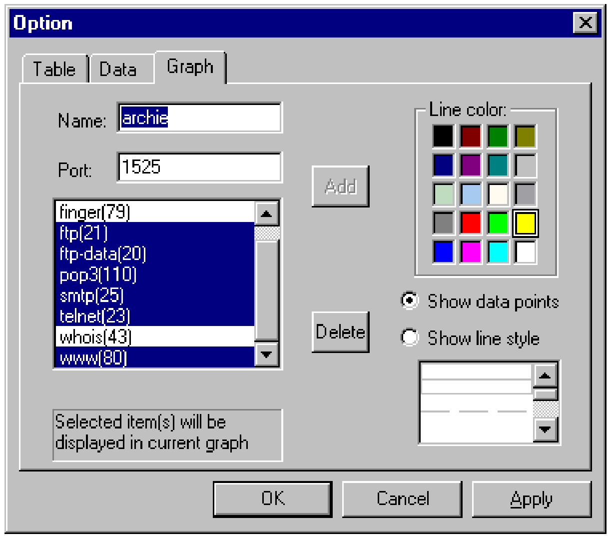

Figure 106. SOCKS monitor graph settings

You also can define and change the line type and the color used for that kind of traffic in the monitor window.

Chapter 8. Secure Mail Proxy

One of the reasons why an organization would want to connect its network to the Internet is mail access. Through the Internet, people in an organization can exchange mail with the rest of the world. IP mail is transmitted via the simple mail transfer protocol (SMTP), which is a simple client/server architecture allowing store and forward or direct delivery.

Usually people want to have free access in and out of the secure network for mail traffic. To accomplish this, the network administrators have to connect the mail servers to the Internet so any machine is able to establish a session with it. Even if the administrator installs a firewall to protect the network, this machine is still exposed to attacks, incoming connections need to be allowed from the Internet to be able to receive e-mail.

The idea of a mail relay is to avoid a direct session from being carried across the firewall gateway. It hides the internal mail gateway from the non-secure network. Only the firewall mail server is advertised outside the secure network, which is much more resistant to attack than the real mail server.

IBM SecureWay Firewall V4.1 for AIX uses its own secure mail gateway called Secure Mail Proxy. Secure Mail Proxy obstructs attempts to subvert the mail server by validating that certain SMTP commands are properly formed before they are relayed to the secure mail server. A user who telnets to the firewall on port 25 is unable to damage the server.

8.1 How it works

The Secure Mail Proxy acts as a real-time gateway between two or more e-mail domains. In contrast with a traditional SMTP relay, messages are not stored on the firewall before being forwarded to the destinations. The SMTP conversation is interpreted as it happens, and the Secure Mail Proxy conversation is forwarded on to each of the necessary destination servers, command by command. For both incoming and outgoing mail, it only relays the message based on the Secure Mail Proxy configuration file /etc/security/fwsecuremail.cfg.

It also allows you to log information about the mail sessions and failure conditions.

8.1.1 SMTP proxy

Secure Mail Proxy is automatically started when we start the firewall. It acts like a mail exchanger, but it does not queue any inbound or outbound mail, nor does it store mail before formatting it. If a destination server accepts the mail, the mail packets are transferred to the destination server like a bidirectional pipe. If a destination server is not available or does not accept the mail, the mail is rejected.

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

8.1.2 SMTP commands

Secure Mail Proxy can understand all SMTP commands: EHLO, HELO, MAIL, RCPT, DATA, RSET, SEND, VRFY, EXPN, HELP, NOOP and QUIT. Before receiving an RCPT command, Secure Mail Proxy accepts EHLO, HELO, MAIL, RSET, HELP, NOOP and QUIT, but all others are rejected. After it receives the RCPT command it looks for the destination server according to the path defined by RCPT. If found, Secure Mail Proxy tries to connect to the destination mail server at TCP port 25. If the destination mail server cannot be found or a connection cannot be made, Secure Mail Proxy rejects the RCPT command and waits for another RCPT command.

If a connection is made, Secure Mail Proxy transfers subsequent commands to the destination mail server (except VRFY and EXPN to a secure mail server). VRFY is used to verify a user name in a mail server. The input string is a user name and the result is detailed information about the user and his/her mailbox. EXPN is used to expand a mailing list. The input string is the mailing list name and the multiple responses will be given containing the full name of the users and their mailboxes. It is dangerous to allow VRFY and EXPN commands because it discloses information of valid e-mail addresses to the outside world.

8.1.3 Multiple secure servers

You can set up multiple secure mail servers. Secure Mail Proxy determines the destination secure mail server by looking up the destination specified by the RCPT command. This destination is compared to the domains specified in the Public Domain Name field of the secure mail server setup.

If a name is found (which means that this domain is an internal domain known by Secure Mail Proxy), it tries to connect to the corresponding mail server. If more than one secure mail server is found, it will send it to the first available server. If you are using load balancing, it will balances the delivery between the servers.

If a name is not found, Secure Mail Proxy attempts to resolve the name as a hostname (this should happen only for external domains). If a hostname is not found, Secure Mail Proxy makes a call to DNS for a mail exchanger (MX record) that corresponds to the name and if found, connects to it.

If multiple MX records are found, the Secure Mail Proxy will follow the priority set for those records in DNS. If an MX record is not found, then the Secure Mail Proxy will try to find an A record.

The Secure Mail Proxy fans out messages to multiple domains with a limit of 32 domains. The default is 15 (see 8.1.7.1, “Proxy Characteristics” on page 175).

8.1.4 Incoming mail

In this section, we explain what happens in an SMTP session to incoming mail. Figure 107 illustrates the flow of incoming mail.

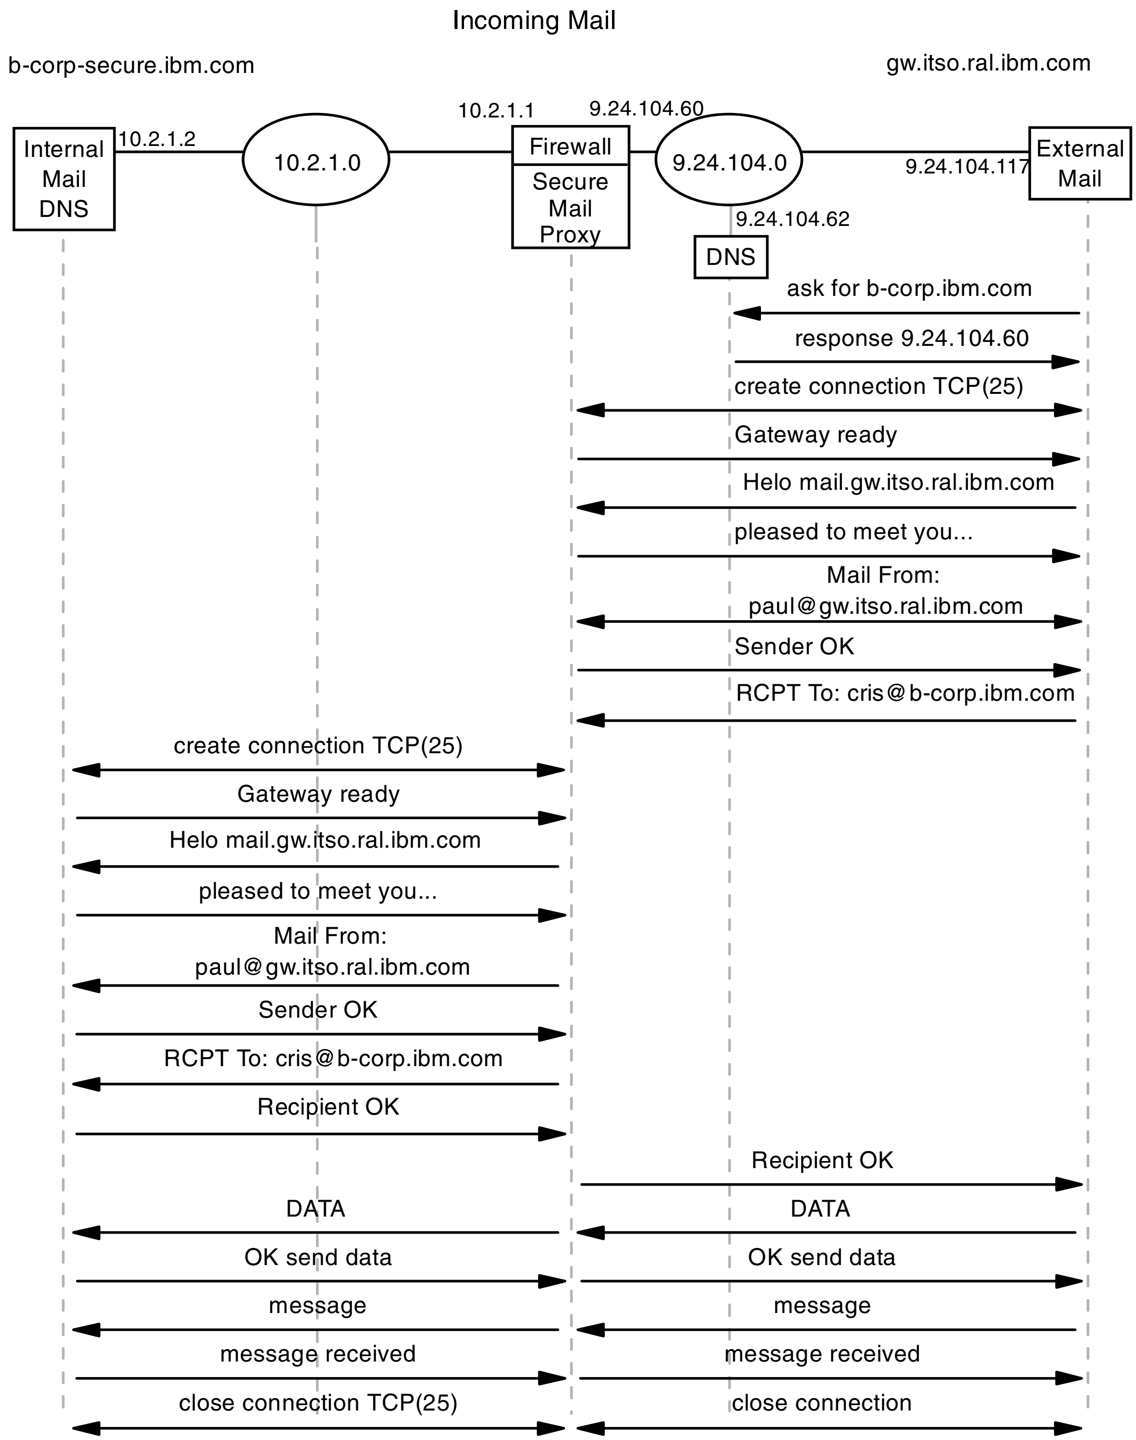

Figure 107. Incoming mail flow

We assume that mail has been sent from the client to the SMTP server, so the flow will begin from the SMTP server. The flow is based on the configuration that is installed in our lab. The mail is sent from paul@gw.itso.ral.ibm.com to cris@b-corp.ibm.com.

Step 1: The external mail server tries to resolve the destination domain. It queries the external DNS for the destination domain. It cannot be resolved by the name server, but the name server has an MX record that points to the IP address of the nonsecure interface of the firewall. The following is an example of the MX record.

```
b-corp.ibm.comINMX10b-gateway.gw.itso.ral.ibm.com.
b-gateway.gw.itso.ral.ibm.com.INA9.24.104.60
```

Step 2: The name server sends a response to the external mail server. The response contains the IP address of the non-secure interface of the firewall. The IP address will be used by the external mail server to establish an SMTP session.

Step 3: The external mail server and non-secure interface of the firewall establish a TCP session on port 25 (SMTP). In Figure 107 we do not give a detailed flow of how the TCP connection is established. As you know, it is established by a three-way handshake process (SYN, SYN/ACK and ACK).

Step 4: The external mail server sends the SMTP message: `HELO mail.gw.itso.ral.ibm.com`.

Step 5: The Secure Mail Proxy responds with `pleased to meet you....`

Step 6: The external mail server sends `MAIL From:paul@gw.itso.ral.ibm.com`. This describes the source address of the mail.

Step 7: The Secure Mail Proxy accepts the sender by replying `Sender OK`.

Step 8: The external mail server sends `RCPT To:cris@b-corp.ibm.com`. This describes the destination address of the mail.

Step 9: Before the Secure Mail Proxy replies to the message, it tries to resolve the destination address. Based on the configuration file of Secure Mail Proxy, it knows where the TCP connection should be established. In this case the internal mail server address is the destination IP address.

Step 10: The Secure Mail Proxy creates a TCP connection on port 25 from a secure interface to the IP address of the internal mail server.

Step 11: The internal mail server sends the response `Gateway Ready`, explaining that the SMTP gateway is ready.

Step 12: Secure Mail Proxy sends the SMTP message: `HELO mail.gw.itso.ral.ibm.com.`

Step 13: The internal mail server answers with `pleased to meet you....`

Step 14: Secure Mail Proxy starts to send the source address of the mail by sending `MAIL From:paul@gw.itso.ral.ibm.com`.

Step 15: The internal mail server responds with `Sender OK`.

Step 16: The Secure Mail Proxy sends `RCPT To:cris@b-corp.ibm.com` to describe the destination address of the mail.

Step 17: The internal mail server checks the domain name that is sent by the Secure Mail Proxy. If it matches the domain name list of the internal mail server configuration, then it will reply `Recipient OK`. Otherwise, it will reject the SMTP message by sending `domain b-corp.ibm.com unknown`. It will directly impact the response that should be sent by the Secure Mail Proxy to the external mail server.

Therefore, it is very critical to set up the internal mail server as a gateway of all domains in the secure network.

Step 18: The Secure Mail Proxy also sends a `Recipient OK` message to the external mail server.

Step 19: The external mail server sends the SMTP command: `DATA` to ask whether the Secure Mail Proxy is ready to receive data. And the Secure Mail Proxy also sends the SMTP command `DATA` to ask whether the internal mail server is ready to receive data.

Step 20: The internal mail server responds with `OK send data` to the Secure Mail Proxy, which indicates that the internal mail server is now ready to receive data. That response is also sent from the Secure Mail Proxy to the external mail server for the same purpose.

Step 21: Now, the data will flow from the external mail server to the internal mail server through the Secure Mail Proxy. This is like an application proxy.

Step 22: The data will be relayed line by line and every time the internal mail server receives one line, it will acknowledge the data by sending the SMTP message `message received`. This acknowledgment will be sent also from the Secure Mail Proxy to the external mail server.

Step 23: After all messages have been sent, both the external mail server and the Secure Mail Proxy will respectively close the TCP connections with the Secure Mail Proxy and the internal mail server.

8.1.5 Outgoing mail

Outgoing mail basically uses the same concept as for incoming mail. But now, Secure Mail Proxy should resolve the external mail server. The mail is sent from cris@b-corp.ibm.com to paul@gw.itso.ral.ibm.com. The outgoing mail flow is illustrated in Figure 108.

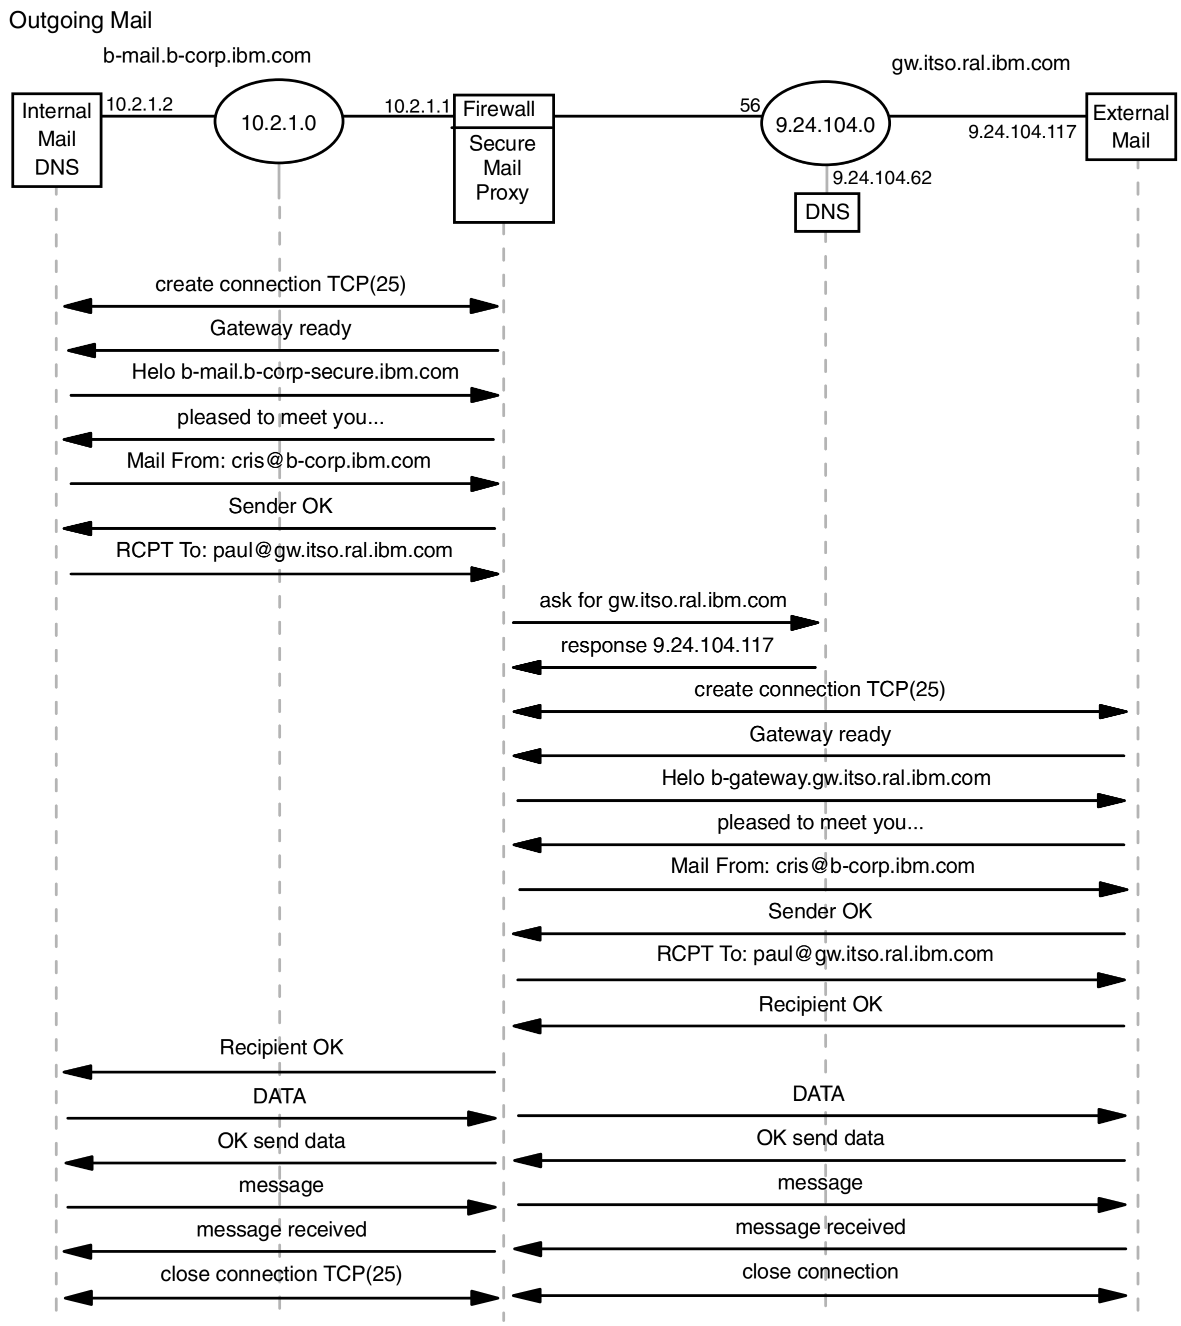

Figure 108. Outgoing mail flow

Step 1: First, we have to make sure that the configuration in the internal mail server is right. The internal mail server should relay all mail that is destined to the Internet to the Secure Mail Proxy. Then, it will create a TCP connection on port 25 with the secure interface of the firewall.

Step 2: The Secure Mail Proxy will answer the connection by sending the SMTP message `Gateway ready`.

Step 3: The internal mail server begins the SMTP conversation by sending `HELO b-mail.b-corp-secure.ibm.com`.

Step 4: The Secure Mail Proxy will reply with `pleased to meet you`.

Step 5: The internal mail server sends the source address of the mail using the SMTP message `MAIL From:cris@b-corp.ibm.com`.

Step 6: The Secure Mail Proxy answers the message with `Sender OK`.

Step 7: The internal mail server sends the destination address of the mail using the SMTP message `RCPT To:paul@gw.itso.ral.ibm.com`.

Step 8: The Secure Mail Proxy resolves first the destination domain from the destination address of the mail. This resolution uses the standard flow of the firewall. It queries the internal name server first. If the internal name server does not know the answer, it forwards the request to the firewall. The firewall queries the external name server. If the external name server cannot find an record for this name, then there is no valid answer for Secure Mail Proxy. Normally the DNS would answer with the address defined in an MX record.

We suppose that the external DNS finds an A record for this address. The external name server sends the response to the firewall, then the firewall sends it to the internal name server and the internal name server sends it back to the firewall. Now, the firewall knows the destination IP address of the external mail server.

Step 9: The Secure Mail Proxy creates a TCP connection on port 25 with the external mail server.

Step 10: The mail server responds to the connection by sending the SMTP message: `Gateway ready`.

Step 11: The Secure Mail Proxy starts the conversation with the SMTP command: `HELO b-gateway.gw.itso.ral.ibm.com`.

Step 12: The external mail server sends a response `pleased to meet you...`, which indicates that the external mail server is ready to begin the conversation.

Step 13: The Secure Mail Proxy sends the source address of the mail by sending the SMTP command `MAIL From:cris@b-corp.ibm.com`. Please notice that the Secure Mail Proxy has changed the source domain from itso.faire.com to faire.com. This is another major function of the Secure Mail Proxy: to hide internal domain names.

Step 14: The external mail server sends a response `Sender OK` to indicate that it accepts the sender.

Step 15: The Secure Mail Proxy sends `RCPT To:paul@gw.itso.ral.ibm.com` to describe the destination address of the mail.

Step 16: The external mail server determines whether this is the valid destination domain name or not. If it is the valid domain name, then the external mail server responds with the message `Recipient OK`.

Step 17: The Secure Mail Proxy also sends the same message (`Recipient OK`) to the internal mail server, which indicates that the destination name is well known.

Step 18: The internal mail server sends the SMTP command `DATA` to ask whether the Secure Mail Proxy is ready to receive data. And the Secure Mail Proxy also sends the same SMTP command to the external mail server.

Step 19: The external mail server responds with `OK send data` to the Secure Mail Proxy, which indicates that the external mail server is now ready to receive data. That response is also sent by the Secure Mail Proxy to the internal mail server for the same purpose.

Step 20: Now, the data starts flowing from the internal mail server to the external mail server through the Secure Mail Proxy.

Step 21: The Secure Mail Proxy relays the data line by line and every time the internal mail server receives the data, it will acknowledge the data by sending the SMTP message `message received` to the Secure Mail Proxy. This acknowledgement is also sent by the Secure Mail Proxy to the external mail server.

Step 22: After all messages have been sent, both the internal mail server and the Secure Mail Proxy will respectively close the TCP connections with the Secure Mail Proxy and the external mail server.

8.1.6 Overflow server

The overflow server is responsible for handling any messages which, due to errors, the Secure Mail Proxy was unable to handle. Messages are routed to the overflow server under these circumstances:

1. One or more receiving SMTP servers generate an error after the proxy begins to transmit the body of the message, while one or more receiving SMTP servers involved in the same transmission receive the message successfully. If the overflow server is not installed, the e-mail will be sent again to all recipients (the receiving mail server should be able to identify duplicate e-mail and discard it).
2. The e-mail being sent exceeds the Secure Mail Proxy's fan-out limit. The Secure Mail Proxy will only open a certain number of outbound connections for delivering the message to each destination mail server (fan-out). Destinations exceeding that limit will be forwarded to the overflow server.
3. If the destination server is known to DNS but it is not available, then the message is sent to the overflow server.

The overflow server can be co-resident with the firewall, or it can reside on a different computer. However, if it resides in a different computer, it should be placed on the secure side of the network (for examples of these configurations, see 8.3.5, “Lab scenarios” on page 201).

If you notice any of the following situations in your environment, we recommend you set up an overflow server:

- Your internal mail server does not limit the number of outbound recipients and it cannot handle temporary errors properly.
- Your internal mail server does not handle large mail queues, or it does not retransmit the e-mail when it cannot be delivered.
- You are noticing excessive number of duplicate messages.

You have basically three options for where to install the overflow server:

1. Use the internal mail server as the overflow server. In this case, the connection is established on port 25. This server must be able to handle the retransmission properly to avoid loop conditions.

2. Use of another mail server on the firewall. In this case, you can use sendmail to be the overflow server, but it cannot use the same port as the Secure Mail Proxy. This configuration also requires extra disk space for the queue management.
3. Use of a separate mail server on another machine. In this case, it is better to allow this server to send outbound messages directly through the firewall, to avoid loop situations.

Refer to 8.3.5, "Lab scenarios" on page 201 for practical examples of using the overflow server.

8.1.7 Additional security

In this section we discuss the features of the Secure Mail Proxy. These features provide more security to your environment, including options for protecting your internal server, avoiding mail SPAM, internal domain hiding, among others.

For more information about configuring the Secure Mail Proxy, see 8.4, "Advanced configuration" on page 214.

8.1.7.1 Proxy Characteristics

In the GUI main window, open the folder **System Administration**, then open the folder **Secure Mail Proxy** and double-click **Proxy Characteristics**. The window in Figure 109 is displayed.

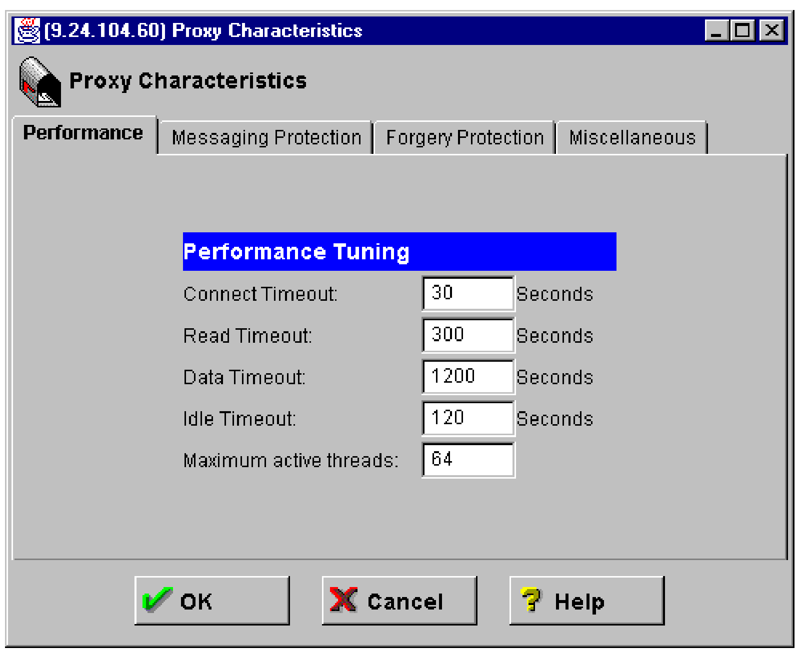

Figure 109. Secure Mail Proxy performance options

The options available inthe Performance tab are:

Connect Timeout: specifies the time to establish an IP connection to the remote server. The default value is 30 seconds.

Read Timeout: specifies the time it waits while trying to receive data from the remote server. The default value is 300 seconds.

Data Timeout: specifies the time to wait for the remote server response after finishing sending the data. The default is 1200 seconds.

Idle Timeout: specifies the time it waits for a command after the data was sent. For example, in case the remote server finishes sending the e-mail but does not close the session, it will time out according to the value of this field. The default is 120 seconds.

Maximum active threads: specifies the number of sessions that can be established at the same time. We recommend that this field is tuned according to the capacity of the internal mail server. For example, if the internal mail server can open up to 16 sessions at the same time, we recommend you use the same value here, so the Secure Mail Proxy will not try to open more sessions than the internal mail server can handle. The default value is 64.

The next tab is Messaging Protection. Click the tab **Message Protection** and the window shown in Figure 110 is displayed.

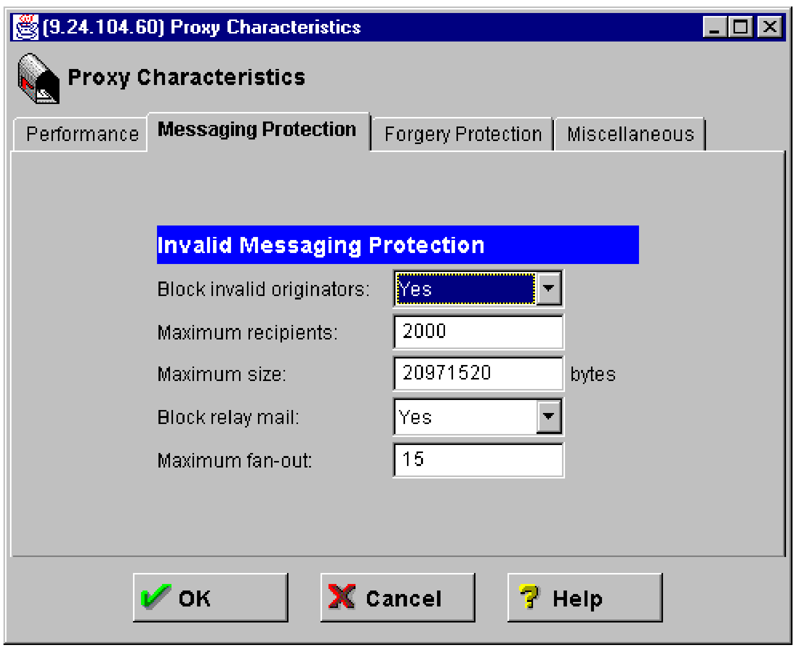

Figure 110. Secure Mail messaging protection options

The options available in this tab are:

Block invalid originators: The Secure Mail Proxy will only accept e-mails from a hostname that can be replied to. It looks up the hostname, using the hosts file and/or DNS (including the MX record). If this option is set to Yes and the Secure Mail Proxy does not find any record for this hostname, the connection is rejected. If it is set to Warn, it accepts the e-mail and adds an entry to the log. If it is set to No, it accepts the e-mail and does not log.

Maximum recipients: specifies the maximum number of recipients in a single e-mail. If an e-mail exceeds this limit, the Secure Mail Proxy sends an error message using the RCPT command. The default value is 2000.

Maximum size: specifies the maximum size in bytes of the body of the e-mail message (including attachments). The default is 20971520 bytes.

Block relay mail: if this option is set to Yes, it will allow only the servers from the local domain to use the secure mail server as mail relay. It will block any other machine. If it is set to Warn, it will allow the relaying from any machine, and it adds a log entry when the originator is not in the internal domain. If it is set to No, it will allow relaying for any machine without logging any warning.

Maximum fan-out: specifies the maximum number of outbound sessions generated by a single inbound session. The default is 15.

Click the **Forgery Protection** tab and the window shown in Figure 111 is displayed.

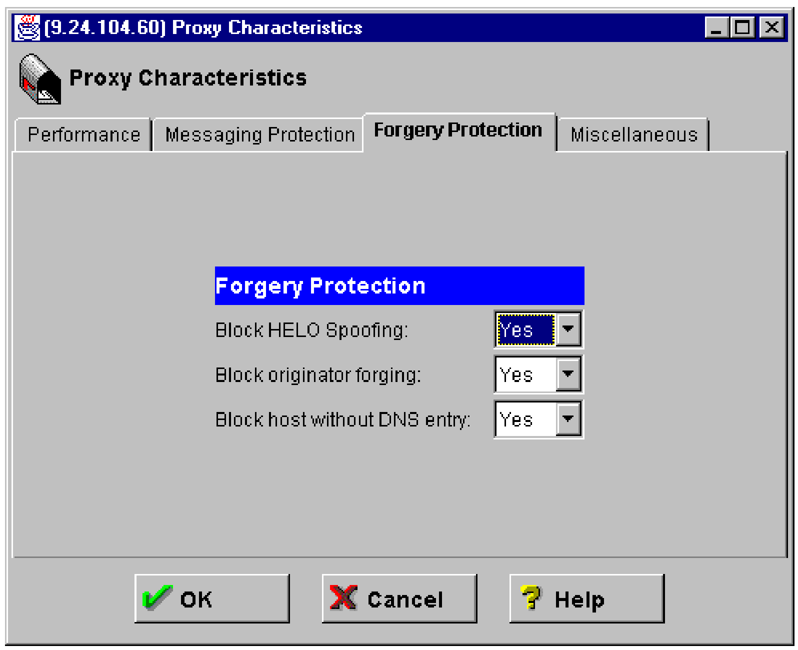

Figure 111. Secure Mail Proxy forgery protection options

All the options in this panel accept one of the following values:

- **Yes**: it blocks the e-mail according to the restriction of the option.
- **Warn**: it accepts the e-mail and adds an entry to the log file.
- **No**: it accepts the e-mail and does not log.

The options available in this panel are:

Block HELO spoofing: If you select Yes, the secure mail proxy will reject connections where the hostname of the sending system, as determined by reverse DNS, does not match the hostname provided by the sender in the HELO command.

Block originator forging: in case the originator domain belongs to the internal network, it checks if the IP address belongs to the internal network also. If it does not, it rejects the message.

Block host without DNS entry: it rejects messages from any host that cannot be identified by DNS.

Note

When you use the options Yes/Warn/No, remember that Yes and Warn may cause some extra overhead in the Secure Mail Proxy, because it has to do extra checking before accepting an e-mail. Using No means that you will have no overhead, but may need to turn some security options on.

We recommend you study carefully all options, and start by turning them all off, and then using Warn to do any tests you may need. Turn on just the ones you really need.

Click the **Miscellaneous** tab and the window shown in Figure 112 is displayed.

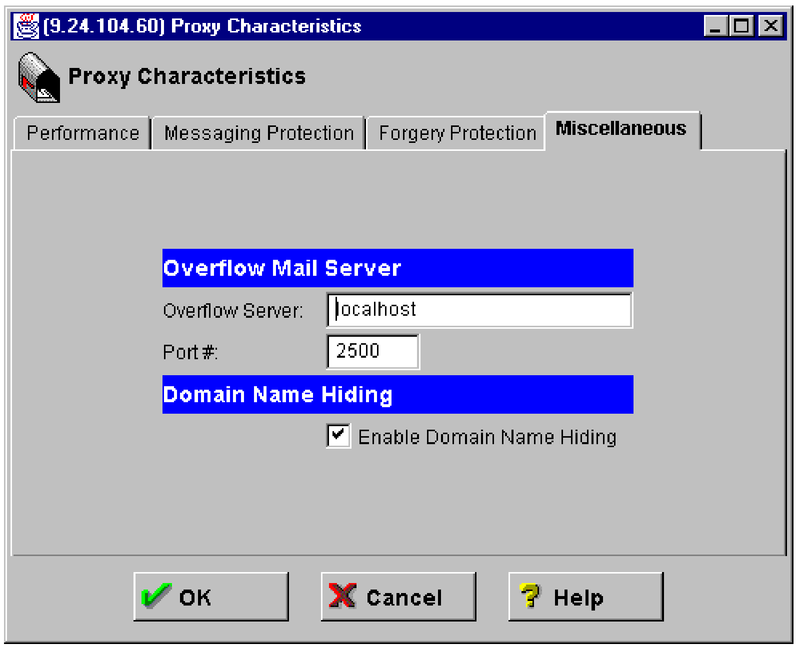

Figure 112. Secure Mail Proxy miscellaneous options

The options available in this tab are:

Overflow Server: specifies the hostname or IP address of the overflow server. The default value is localhost.

Port: specifies the port that the overflow server is listening for e-mail connections. The default value is 2500. For more information about the overflow server see 8.1.6, “Overflow server” on page 174.

Enable Domain Name Hiding: select this check box to enable the domain name hiding for all messages going through the Secure Mail Proxy. For more information on domain name hiding, see 8.1.7.3, “Domain name hiding” on page 179.

Disabling the overflow server

If you do not want to use an overflow server, enter 0 (zero) as port number.

8.1.7.2 Excluded mail domains

In the GUI main window, open the folder **Secure Mail Proxy** and double-click **Excluded Mail Domains**. The window in Figure 113 is displayed.

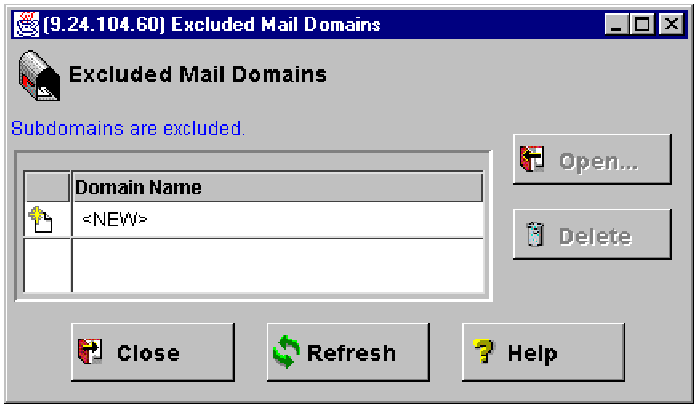

Figure 113. Excluded mail domains

This allow you to make a list of domains from which you do not want to receive e-mail (see 8.3.5.6, “Excluding mail domains” on page 213 for a practical example).

8.1.7.3 Domain name hiding

In the GUI main window, open the folder **Secure Mail Proxy** and double-click **Domain Name Hiding**. The window in Figure 114 is displayed.

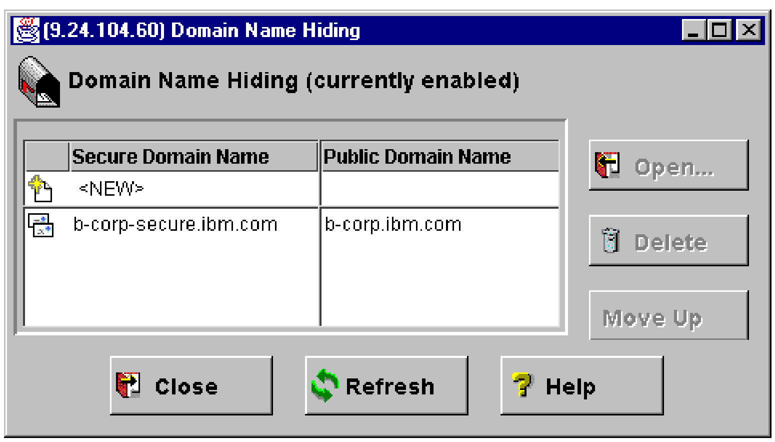

Figure 114. Domain name hiding

This feature is used to hide your secure domain name when sending e-mail to the Internet.

The Secure Mail Proxy replaces the secure domain name with the public domain name on all outgoing mail. If they are the same, no domain name hiding is done. The following headers will be rewritten by the Secure Mail Proxy:

- Received from
- Message-ID
- From
- To

- Reply-To
- Cc
- Bcc

You can modify this behavior. The file fwsecuremail.cfg has the parameter SMTPSB.HIDE_SECURE_NAMES that accepts the following values:

- Y: Yes, hide the secure domain by substituting each secure domain found in the header with its corresponding non-secure domain. Note that prior to this release the substitution was simple string substitution. With this release the substitution occurs for all subdomains of the secure domains.
- N: No, don't modify the header lines at all.
- R: "Received from:" header lines containing the secure domain will be removed and all other occurrences of the secure domain in the header will be substituted.
- A: All "Received from:" header lines will be removed and all other occurrences of the secure domains in the header will be substituted. This option was intended to simulate the original Safemail (the mail proxy of previous IBM firewalls versions) design. NOTE: If mail is allowed to flow into or out of your secure network without going through the firewall SMTP proxy, mail list servers or mail forwarding within your network could lead to mail loops. If this is the case, option A is not recommended. Instead, use either the option R to remove all internally generated "Received from:" lines, or the option Y to hide the secure domains but maintain the "Received from:" lines.

The Secure Mail Proxy does not rewrite inbound addresses set in RCPT commands. So, the internal mail gateway server is responsible for relaying inbound mail to the appropriate destination. The internal mail gateway must be configured to accept the public domain name as an alias for their private domain names.

You can modify this behavior. The file fwsecuremail.cfg has two parameters that apply to this situation: SMTPSB.RECIPIENT_REWRITE and SMTPSB.REWRITE_INBOUND.

- SMTPSB.RECIPIENT_REWRITE: Values supported: Y or N (default: N)
  - Y: Yes, change the public domain to its corresponding secure domain in "RCPT TO:" commands before sending them to corresponding secure mail servers.
  - N: No, don't modify the "RCPT TO:" commands.
- SMTPSB.REWRITE_INBOUND: Values supported: Y or N (default: N)
  - Y: Yes, change the public domain to its corresponding secure domain by substituting each public domain found in the header with its corresponding secure domain before sending mail to corresponding secure mail servers. The substitution occurs for all subdomains of the public domains.
  - N: No, don't modify the headers for mail sent to secure mail servers.

Refer to 8.3.5.4, “Domain name hiding” on page 206 for an example in domain name hiding.

8.2 Planning for your mail configuration

By nature, use of e-mail involves many different servers. This includes the sender, recipient, firewall, DNS servers, mail servers, and other intermediate hosts.

The configuration of the firewall is therefore the collaboration of several parties.

8.2.1 Why plan?

Careful planning and documentation is therefore critical to the success of the implementation for the following reasons:

- The implementation work involves many potentially different people who need to work from a common plan.
- Problem determination requires detailed knowledge of network layout. This is best described on paper, rather than verbally.
- Auditing of network elements is easier with written documentation. This may avoid the need for a physical walkthrough to reverse engineer the configuration.
- Peer review is facilitated by allowing the information to be transmitted and reviewed.
- High-level decision makers can easily understand the topology when working from a fixed set of documentation
- Network upgrades or modifications are easily described in relation to the existing components

The following sections in this chapter form the basis for discussion points that need to be documented prior to implementation work.

8.2.2 Functional overview diagram

A network diagram allows all of the functional elements to be described in the correction relationship to each other. A good diagram will describe the following attributes:

- Functional description of each element
- Hostname
- DNS domain names
- Interface IP addresses
- Interface types
- Interface subnetmasks
- Routes

See the Figure 115 on page 186.

8.2.3 DNS worksheets

SMTP e-mail requires DNS in order to function. Any problem with the DNS servers becomes a problem for the SMTP system. Knowledge of the DNS design is important when reviewing the configuration.

The following worksheets capture the required information:

8.2.3.1 Firewall

The firewall requires the following DNS information:

Table 20. DNS information required for a firewall

Zone	Item	Comment
Nonsecure	ISP DNS Server	IP address of the Internet Service Providers DNS server
	DNS A Record	Mapping of name to IP Address of non-secure interface
	DNS PTR Record	Reverse Mapping of nonsecure IP address to name
	DNS MX Record(s)	Mail Exchanger host and preference
Secure	Secure DNS Server	IP address of secure DNS server
	DNS A Record	Mapping of name to IP address of secure interface
	DNS PTR Record	Mapping of IP address to name of secure interface

Root nameservers

The default configuration of the DNS server on the firewall is a caching only nameserver. The specified non-secure DNS server provided by the ISP is queried to discover DNS entries. When this server is unavailable, the DNS server on the firewall will attempt to contact the root nameservers.

Your system will be potentially more reliable if you allow DNS traffic to flow from your non-secure interface to the root nameservers.

See Chapter 5, “Domain Name System (DNS) Service” on page 77 for more information.

8.2.3.2 Secure mail server

The following information is required by the secure mail server:

Table 21. DNS information required for secure mail server

Item	Comment
Secure DNS Server	IP Address of secure DNS server
Secure DNS Domain	DNS domain used to configure the secure domain

Item	Comment
Managed non-secure DNS Domain(s)	The non-secure domain name the secure SMTP server manages

8.2.3.3 Secure client

Clients in the secure network require the following information:

Table 22. DNS information required for secure clients

Item	Comment
Secure DNS Server	IP Address of secure DNS server used by secure clients
Secure DNS Domain	DNS secure domain the client and secure mail server are using.

Table 23. Bandwidth considerations

The network capacity of the network on the non-secure side of the firewall provides an important limit that affects the design of your firewall. In practice, the answers to the following questions will determine the design and relevant parameters used in the firewall.

8.2.3.4 Dedicated SMTP interface?

Will SMTP traffic require a dedicated interface to avoid contention with other non-SMTP traffic?

- What is the non-secure WAN connection speed to the internet?
- Assuming an average SMTP message size of 1500 bytes, what is the maximum rate of messages this connection will support?
- What is the maximum rate of messages that is expected?
- What other protocols will flow on the same shared connection (for example, HTTP traffic)?

8.2.3.5 Fan-out limit

Increasing the fan-out limit parameter is counter productive if the available bandwidth is already saturated.

The fan-out limit only comes into play when a large number of recipients or domains on the nonsecure side are receiving mail.

8.2.3.6 Timeouts

Increasing the SMTP timeout values is appropriate in the following cases:

- If the available bandwidth is already saturated.
- The fan-out limit is increased.
- Consider also the network latency when sending e-mail to your friends in Australia, Brazil or Korea. The time taken to connect to the remote server may exceed the default timeout values.

8.2.3.7 Determining bandwidth usage

Bandwidth usage on a live system can be determined by two methods:

- Router WAN utilization

Consult your ISP router support to access the WAN utilization transmitted across the interface. This is typically expressed as a percentage of the available network resource, for example: 95% of 128 Kbps.

- LAN interface volumes

  The netstat command includes the total traffic volumes transmitted or received across a LAN interface. The following command will display this information:

```
netstat -v

TOKEN-RING STATISTICS (tok0) :
Device Type: IBM PCI Token-Ring High-Performance Adapter (14101800)
Hardware Address: 00:04:ac:63:31:a4
Elapsed Time: 6 days 22 hours 0 minutes 3 seconds

Transmit Statistics:                      Receive Statistics:
--------------------                      -------------------
Packets: 45943                            Packets: 8255877
Bytes: 4400595                            Bytes: 4331854163
Interrupts: 45943                         Interrupts: 8229351
```

By taking a regular measurement of the transmit and receive statistics, it is possible to determine the amount of data sent during that period.

8.2.4 Storage worksheet

Persistent and temporary storage on disk is required for the following activities:

- Temporary storage during relay of SMTP messages
- Logging of error, warning and security messages
- Auditing of valid transactions
- Overflow server requirements if on the same host as the firewall

These storage requirements are additional to the installation of the base operating system (AIX) and the firewall software.

Review your storage

You will need to review the storage requirements of your firewall over time. We cannot predict your unique requirements. There is no such thing as a "typical" scenario. Review our experiences.

8.2.4.1 How long to keep logs

Your security policy will describe the minimum period over which to keep persistent logging data.

We recommend that logs be kept for a minimum of 30 days to allow sufficient time to detect and trace the history of malicious behavior.

8.2.4.2 Temporary storage

The temporary storage will need capacity to store the maximum number of concurrent SMTP connections multiplied by the maximum size configured.

8.2.4.3 Logging

The minimum recommended logging needs to record only error, warning or security conditions. During the normal course of processing SMTP e-mail, errors are to be expected, such as the mail server being down or the recipient is not known at the destination mail server.

A conservative guess on the capacity required is to assume 5% of the total message traffic may generate temporary or permanent errors.

8.2.4.4 Auditing

If additional logging is required to include valid transactions (see “Code selection” on page 217), calculate the total requirement based on the number of messages processed over the time period required to keep the logs.

We had recorded 320 bytes for a single SMTP transaction:

```
ls -l /var/log/mail
-rw-------   1 root     sys 320 Sep 03 18:11 19990903194234.LOG
```

Multiply this value (320 bytes) by the total number of messages processed per day.

8.2.4.5 Overflow server

If running on the same machine as the firewall, the overflow server needs capacity to store

- Audit logs
- Spool area while holding messages

The format of this data depends on the software used.

8.3 Case study

The configuration of the messaging environment may vary according to the needs of each company, resources, software, and so forth. That is why we tried to create a simple yet comprehensive environment. Using the same environment we created different scenarios, to provide examples of situations that are very common to most network administrators.

We would like to emphasize that these scenarios are just examples of real situations, so they cannot be considered as definitive solutions. Use them as suggestions, but always remember to study your needs and your resources to come up with a solution that best fits your needs.

8.3.1 Description

The following diagram shows the topology of the environment created for the scenarios we will discuss later on.

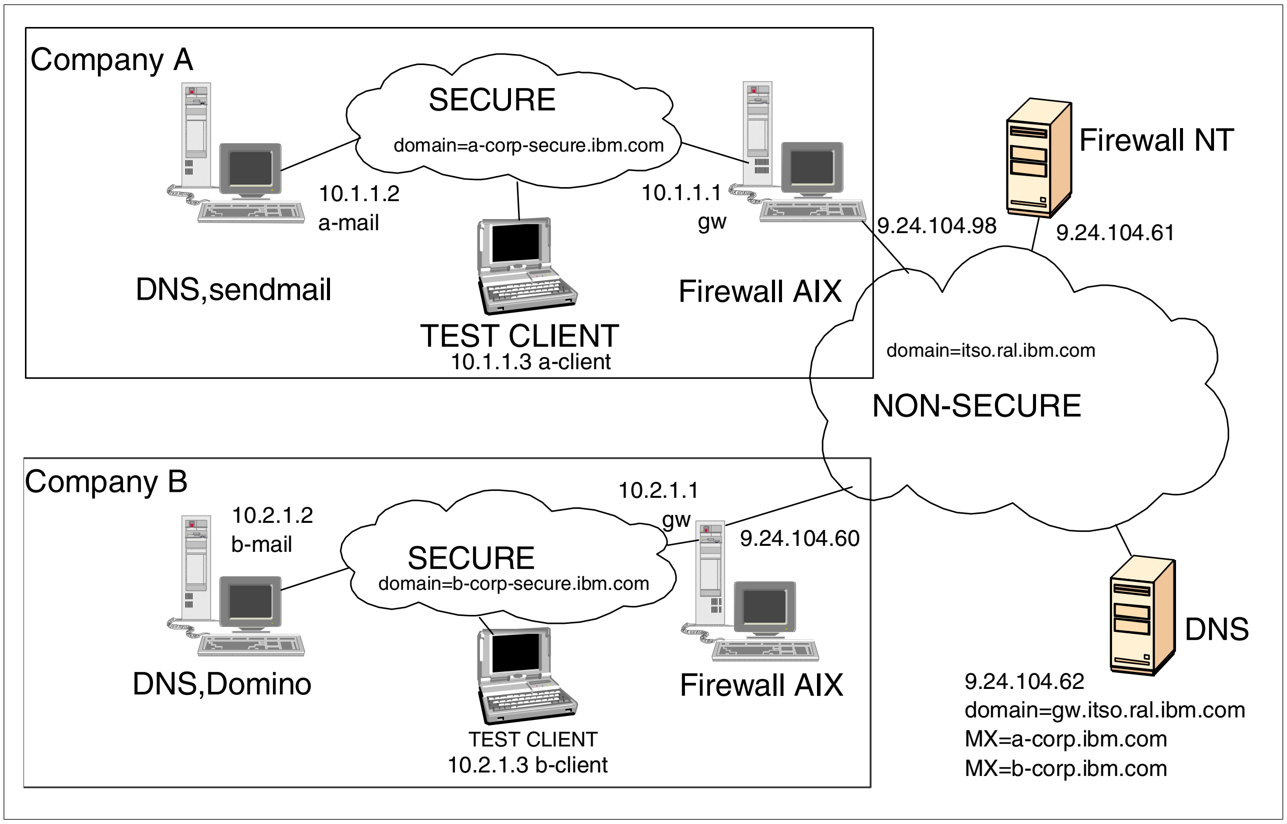

Figure 115. Diagram of the network used in the mail scenarios

We used two separate "secure" networks in our environment. These networks simulate a real-life situation of two separate companies, each one having its own firewall and internal servers, and connected to a non-secure network.

These two networks belong to two companies: Company A and Company B. The ISP that provides them access to the Internet is responsible for maintaining their external DNS configuration and routers. The IP address of the external DNS server is 9.24.104.62.

Company A uses the network address 10.1.1.0 (subnet mask 255.255.255.0) for their secure network. The secure domain is a-corp-secure.ibm.com and the nonsecure domain is a-corp.ibm.com. The mail domain is a-corp.ibm.com.

The machines on this network are:

gw: This machine is the firewall. The IP address of the secure interface is 10.2.1.1 and the IP address of the non-secure interface is 9.24.104.98.

a-mail: This is the internal mail server, running AIX 4.3.3, sendmail 8.9.3 and POP3.

It is also the internal DNS server, and it is authoritative for the domain a-corp-secure.ibm.com and forwards all other queries to the firewall. Its IP address is 10.1.1.2.

a-client: This is the mail client. It uses Netscape Communicator V4 mail client to download messages from the POP3 server and send messages to the SMTP server. Its IP address is 10.1.1.3.

The other network, 10.2.1.0 (subnet mask 255.255.255.0), belongs to Company B. The secure domain is b-corp-secure.ibm.com and the nonsecure domain is b-corp.ibm.com.

Company B recently incorporated Company-C, and it chose to keep a separate mail domain for the users of that company, which is c-corp.ibm.com. This is only an e-mail domain; there are no machines that actually use this domain.

The machines inside this secure network are:

gw: This machine is the firewall. The IP address of the secure interface is 10.2.1.1 and the IP address of the nonsecure interface is 9.24.104.60.

b-mail: This is the Company B internal mail server. It is running Lotus Domino V5, and it is using SMTP and POP3.

It is also the DNS server for this network, and it is authoritative for the secure domain b-corp-secure.ibm.com and it forwards all other queries to the firewall. Its IP address is 10.2.1.2.

c-mail: This is the Company C mail server. It is running Microsoft Exchange V5.5, and it is also using SMTP and POP3. Its IP address is 10.2.1.9.

b-client: This machine is the e-mail client, and it is running both Lotus Notes client V5 and Netscape Communicator V4. This e-mail client reads mail from a Lotus Notes server (domain b-corp.ibm.com). Its IP address is 10.2.1.3.

b-client2: This machine is the e-mail client, and it is running Netscape Communicator V4. This e-mail client reads mail from Microsoft Exchange server (domain c-corp.ibm.com). Its IP address is 10.2.1.8.

8.3.2 Mail servers configuration

Once you have all the information you need about your mail environment (see 8.2, “Planning for your mail configuration” on page 181), you can start setting up the configuration.

In a firewall environment we need to do all the following steps before configuring the firewall itself:

1. The mail server must support SMTP.
2. All e-mail with a destination address other than the local domain must be forwarded to the firewall.
3. All mail clients inside the secure network must send the e-mail to the internal mail server.
4. The internal mail server must be configured to accept messages with both secure and nonsecure domain in the e-mail address (in case these domains are not the same), except if you are using the recipient rewrite option (see 8.1.7.3, “Domain name hiding” on page 179).
5. You need to include MX records in the external DNS for all your non-secure mail domains, and the mail exchanger for those domains must be the firewall nonsecure interface.

It is not our purpose to show the complete installation and configuration procedure for the mail servers. If you already have a mail server running, or if you are going to configure your mail server from scratch, refer to the documentation of the product in how to configure it and how to do the changes necessary to use it in a firewall environment according to the preceding list of prerequisites.

In the following sections, we cover the specific configuration needed on sendmail and Microsoft Exchange in our environment.

8.3.2.1 Sendmail 8.9.3

For Company-A's mail server we used sendmail V8.9.3, which is shipped with AIX V4.3.3.

This is the server for the domain a-corp.ibm.com. It is working internally with the secure domain, which is a-corp-secure.ibm.com. When the e-mail is sent to the Internet, Secure Mail Proxy translates this domain into the nonsecure domain, which is a-corp.ibm.com, and translates it back to a-corp-secure.ibm.com when the e-mail comes in from the Internet.

The configuration of sendmail is done by editing the file /etc/sendmail.cf. We did the following changes to this file:

1. We added the secure domain a-corp-secure.ibm.com to the Cw definition, so sendmail accepts any e-mail to this domain as local:

   ```
   Cwlocalhost a-corp-secure.ibm.com
   ```

2. We defined the DS macro with the hostname of the firewall (this hostname must be resolved into the secure interface address). In our example, the hostname is gw.a-corp-secure.ibm.com, as follows:

   ```
   DSgw.a-corp-secure.ibm.com
   ```

3. We configured the masquerade function to use the secure domain in all outgoing e-mail. First, we added the name we want to hide (which is the local hostname:

   ```
   CMa-mail.a-corp-secure.ibm.com
   ```

   Then, we added the domain we want to masquerade as in the DM macro:

   ```
   DMa-corp-secure.com.br
   ```

   To enable the masquerade, we have to make sure that the S94 ruleset is uncommented:

   ```
   S94
   R$* < @ *LOCAL* > $*     $: $1 < @ $j . > $2
   ```

4. We also changed the DZ macro to hide the true version of sendmail that we are running:

   ```
   #DZ8.9.3
   DZ1.0
   ```

5. We changed the greeting message (the banner that appears when a host connects to the mail server), so it does not show much information about the sendmail and the environment. We removed the $v, which shows the version of AIX and sendmail that it is running. We kept $j (which is the hostname of the machine where sendmail is running), $Z (which we changed in step 4, so it does not show the real information) and $b, which is the current time and date.

If you prefer, you may do further changes in this greeting message. You can see our final greeting message below:

```
O SmtpGreetingMessage=$j Sendmail $Z; $b
```

6. Finally, we changed the format of the Received line that is added by sendmail, so it does not include information about the machine and the internal network:

```
HReceived: $?sfrom $s $.$?_($?s$|from $.$_)
        $.by $j ($Z)$?r with $r$. id $i$?u
        for $u; $|;
        $.$b
```

You could do more changes to make sendmail more secure, but in our environment we have the Secure Mail Proxy protecting it.The Secure Mail Proxy prevents those server from using certain SMTP commands, so they have limited access to this server.

8.3.2.2 Microsoft Exchange V5.5

This is the server for Company-C domain. Since this is only a mail domain (we are not using it as the DNS domain), we chose not to create a separate secure domain, so the public domain and the secure domain are the same: c-corp.ibm.com.

We are not covering here the installation and configuration of this product, so for information about it consult the vendor of this product. What we are covering are the changes you need to do in order to use this server in a firewall environment with the Secure Mail Proxy.

After installing it (remember that you have to install and configure SMTP also), we configured the server as described next.

Open the Microsoft Exchange Administrator tool, then click the **Configuration** button or press Ctrl+Shift+C. You will see the configuration options, as shown in Figure 116.

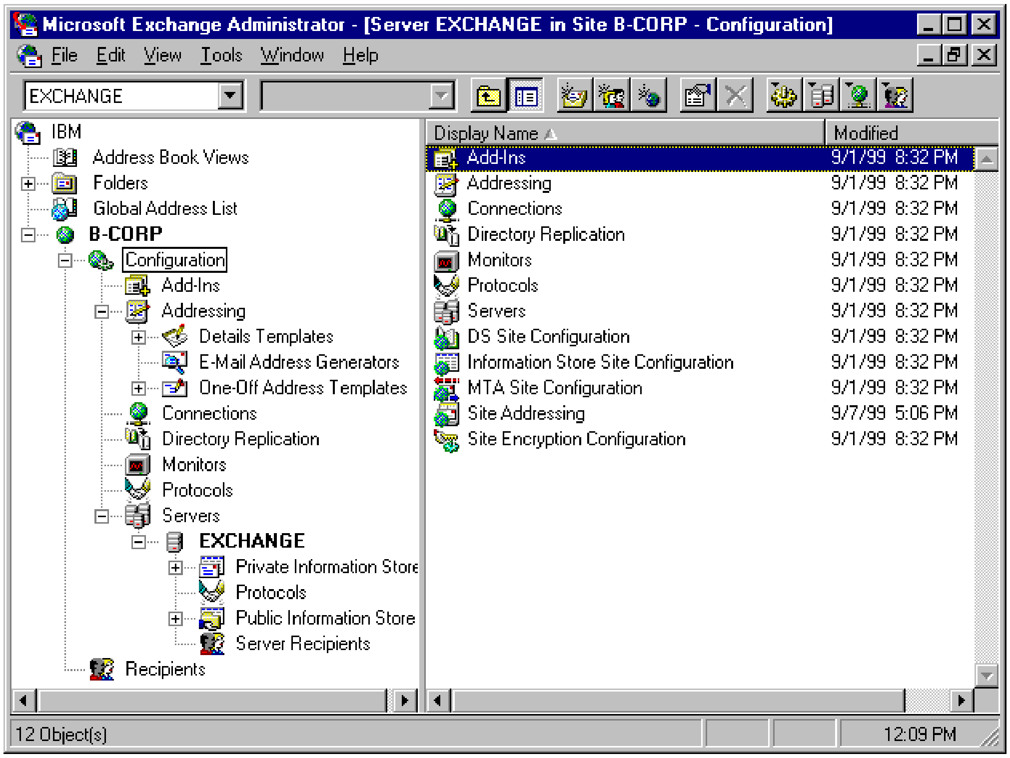

Figure 116. Microsoft Exchange Administrator window

Double-click **Site Addressing** on the panel on the right hand side, and the Site Addressing Properties window is displayed. Select the **Site Addressing** tab, as shown in Figure 117.

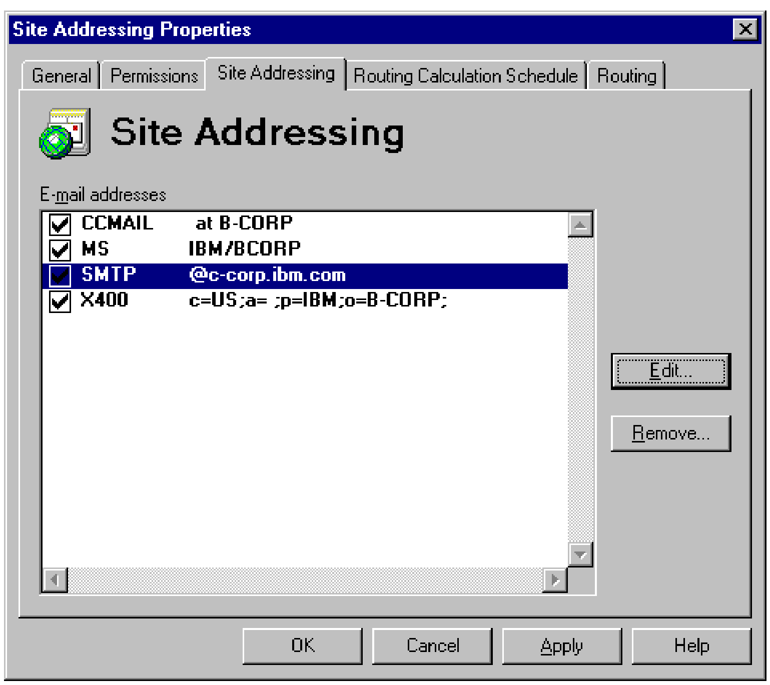

Figure 117. Site Addressing Properties window

The SMTP entry must have the mail address that is considered local for this server. In our example, our local mail address is @c-corp.ibm.com. Click **OK** to close this window and get back to the main administrator window.

Now, double-click **Connections** (see Figure 116), and the options list on this panel will be replaced. Double-click **Internal Mail Service** in the next list, and the Internet Mail Service Properties window is displayed. Select the **Connections** tab (see Figure 118).

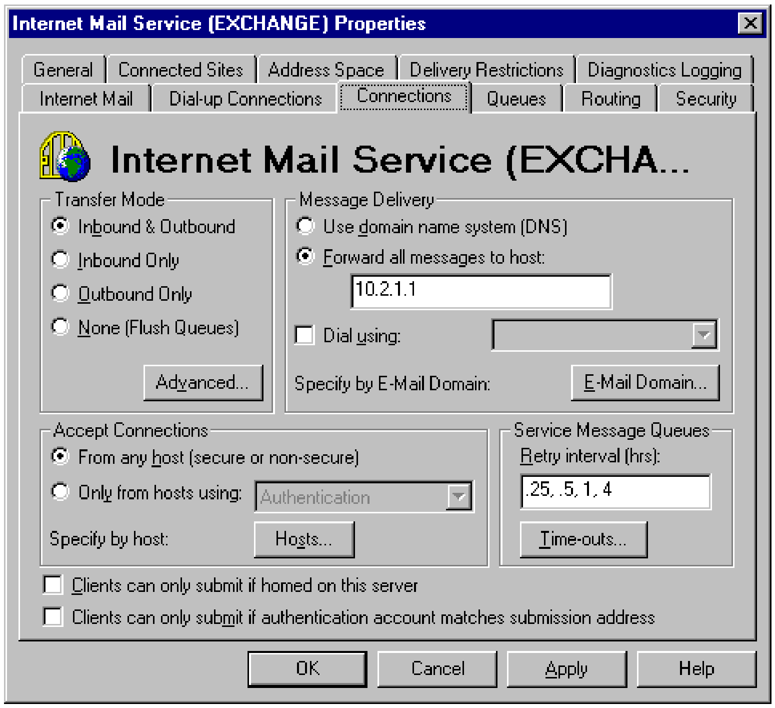

Figure 118. Internal Mail Service properties - connections

In the Message Delivery panel, select **Forward all messages to host**, and fill in the address of the firewall secure interface. In our example, we filled in the IP address `10.2.1.1`. By selecting this option, you are configuring your mail server to forward all messages to the firewall.

We still have to configure this mail server to retain the messages sent to the internal domain. Select the **Routing** tab on this same window (see Figure 119).

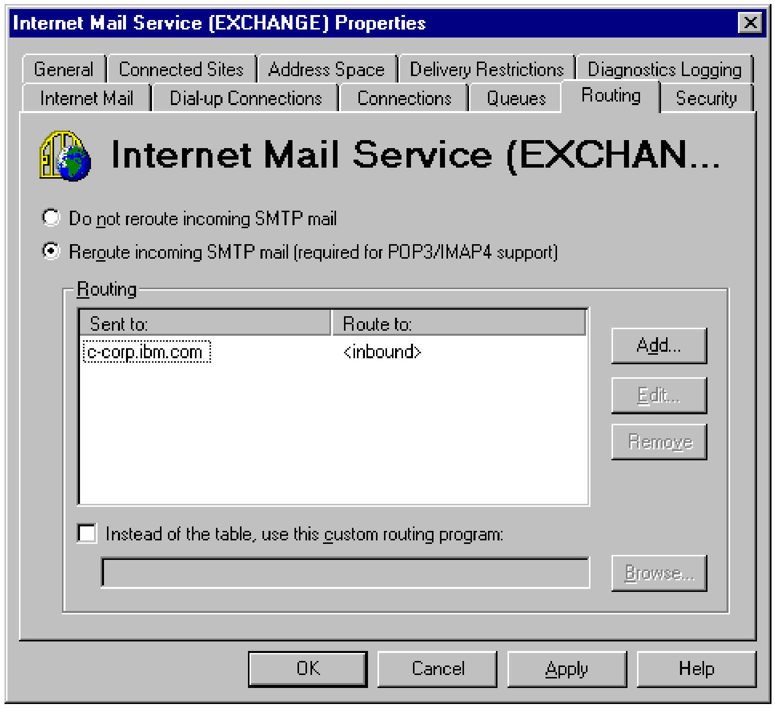

Figure 119. Internal Mail Service properties - routing

Make sure that you have an entry for your mail domain routing to "inbound". This means that all messages sent to <user>@c-corp.ibm.com will be considered as local message and will not be routed to other servers.

After making the changes, restart the mail services.

8.3.3 Client configuration

In our tests we used an SMTP and POP3 client instead of using each product specific client. For information on configuring the mail clients for each product refer to the product documentation.

We used Netscape Communicator V4 mail client. This example shows how we configured it to access the mailbox of the user *cris* using the POP3 server, and to send e-mail using the SMTP server.

Open the Navigator window, and click **Edit** -> **Preferences**. The Preferences window is displayed. Double-click **Mail & Newsgroups**, and click **Identity**, as shown in Figure 120.

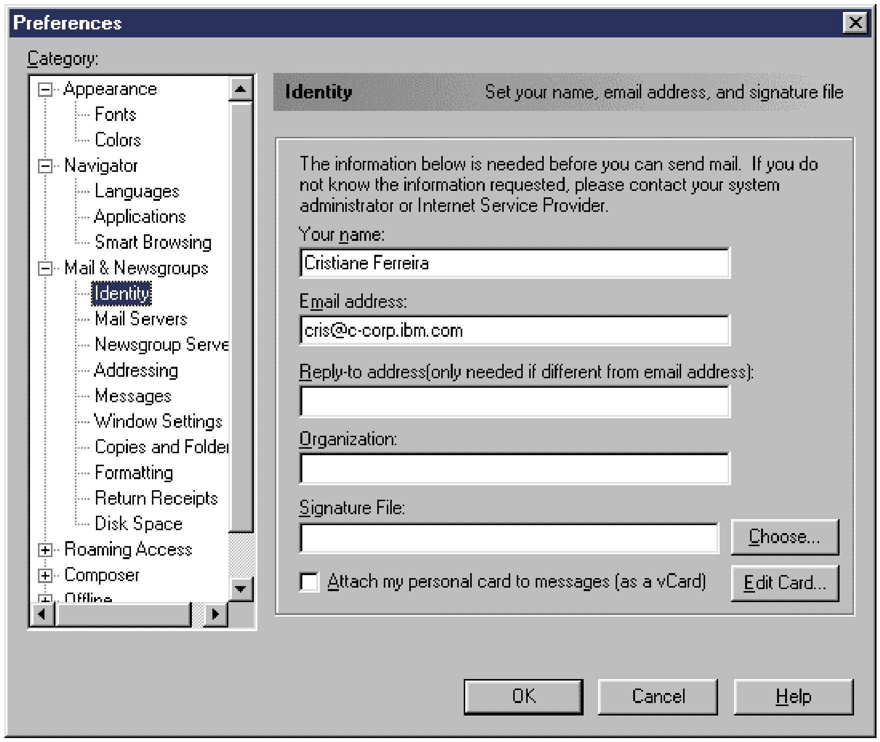

Figure 120. Netscape Communicator: user information

The name and e-mail address you fill in this window will appear in "From" field of the header of your e-mail. So make sure you supply a valid address in the e-mail address field, or the recipients of your messages will not be able to "reply" to this message.

Next, click **Mail Servers** (this option is on the left panel), and the window is updated with new options, as shown in Figure 121.

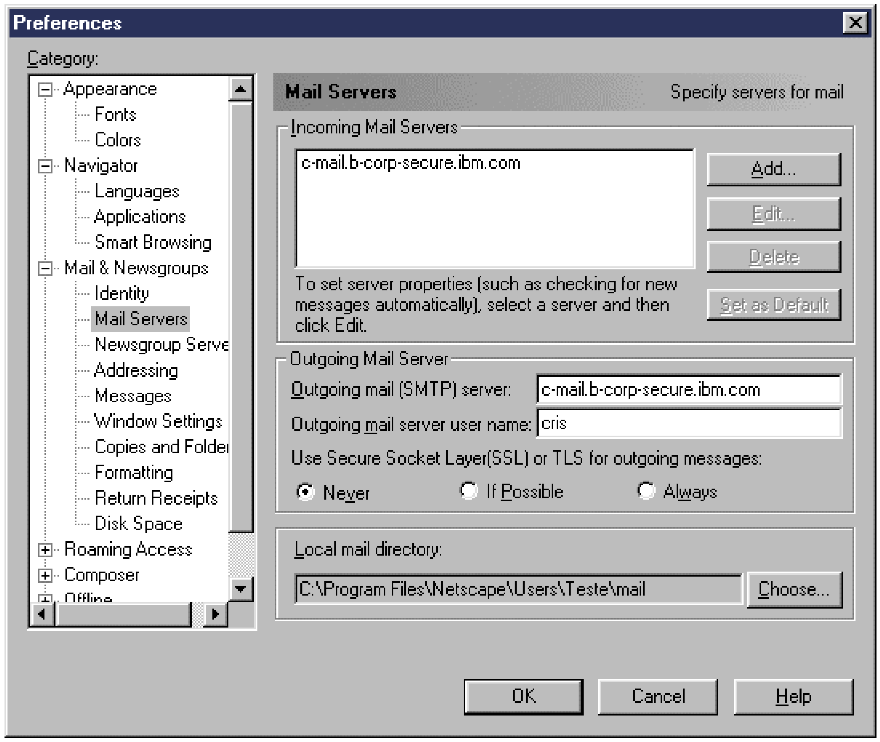

Figure 121. Netscape Communicator: SMTP server configuration

The Outgoing Mail Server is used to input the information about your SMTP server. In our example, our SMTP server is c-mail.b-corp-secure.ibm.com. Fill in the hostname or IP address of your SMTP server in the field Outgoing mail (SMTP) server, as shown in Figure 121. Fill in your username in the field Outgoing server user name (do not include the domain name).

Next, we input the information about the POP3 server. In the same window, locate the panel Incoming Mail Servers. Click **Add**, and the window in Figure 122 is displayed.

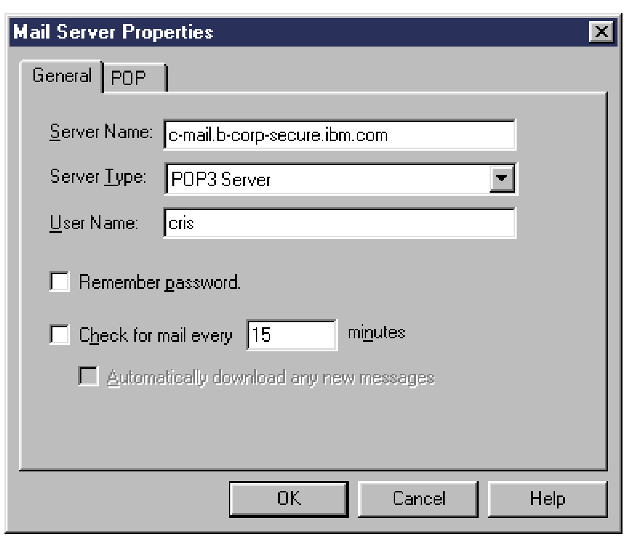

Figure 122. Netscape Communicator: POP3 configuration

On the field Server Name, put the hostname or IP address of your POP3 server. In our example, the POP3 server is c-mail.b-corp-secure.ibm.com. Select **POP3 Server** in the Server Type field, and put the username (no domain) in the User Name field.

For more information on using this mail client refer to the product documentation.

8.3.4 Firewall configuration

In this section we cover the configuration that was made on the firewall. Note that some specific scenarios may need further configuration on other machines or other applications than the firewall, and eventually some changes on the firewall itself. Read carefully all sections before implementing.

8.3.4.1 Filters

We need to add filters to allow the traffic of incoming and outgoing e-mail. Also, we need to keep in mind that we do not want to allow more than what is necessary.

Before creating the connections to allow the mail flow, we created a group (in the Network Objects management window) containing all internal mail servers. This group, which we called "Mail servers", is used in the connections instead of repeating the configuration for each individual mail server (see Figure 123).

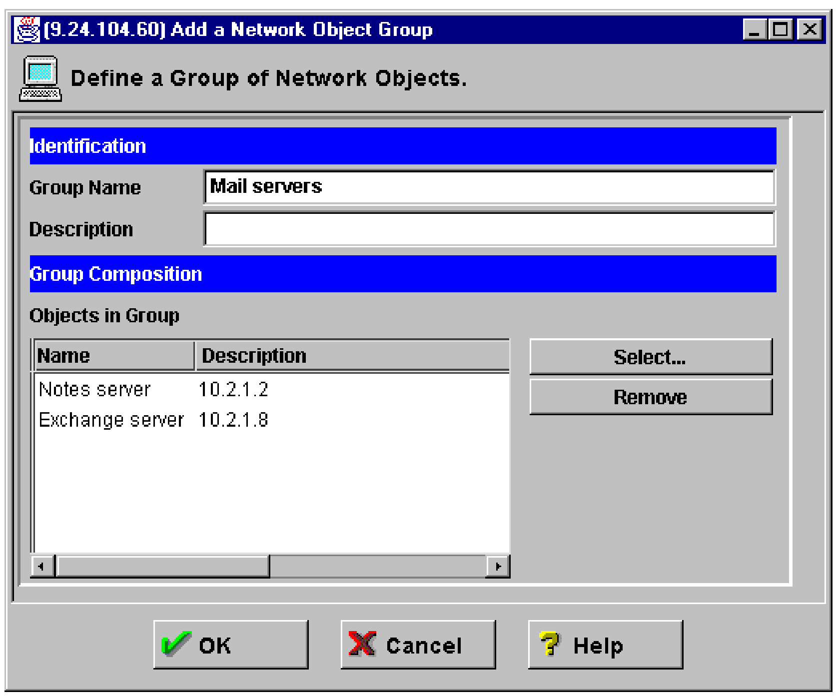

Figure 123. Network group "Mail servers"

The following steps show the connections we need to add in order to be able to send and receive e-mail from the Internet. Note that all the following connections use the same predefined service "Mail". This service allows any mail traffic with local routing (it means that it does not allow routed traffic), so it can be used for all our connections.

1. Allow incoming connections from any server on the Internet to the firewall.

Figure 124 shows the connection we created to allow this traffic.

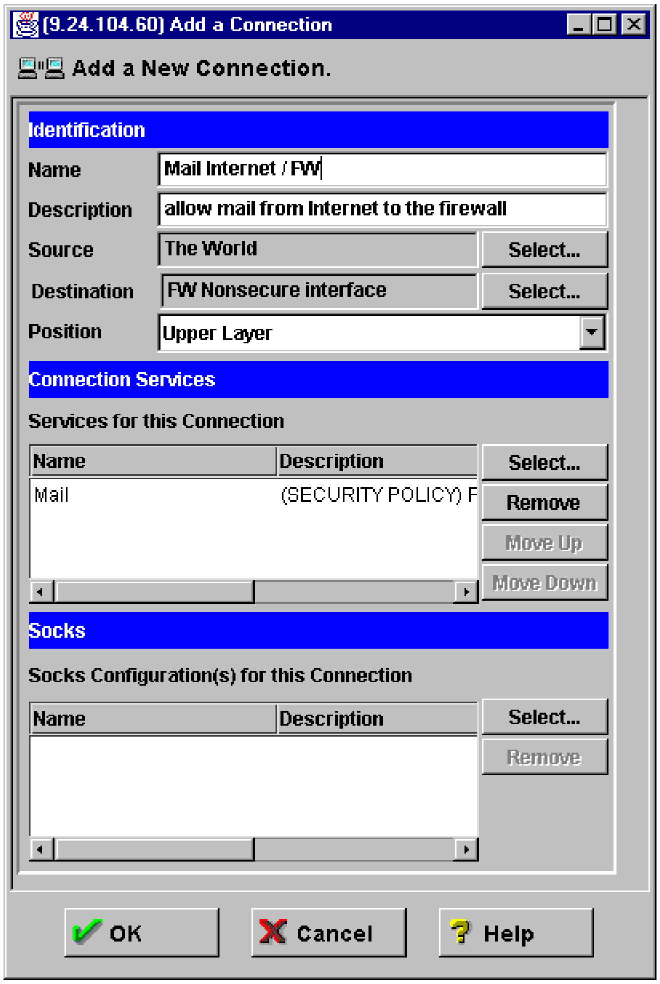

Figure 124. Connection allowing mail from the Internet to the firewall

2. Allow outgoing connections from the firewall to the internal mail server. If there is more than one mail server in the secure network, we recommend that you create an object group and add all mail servers in it, and then create only one connection (as we mentioned in the beginning of this section).

   Figure 125 shows the connection we created to allow this traffic.

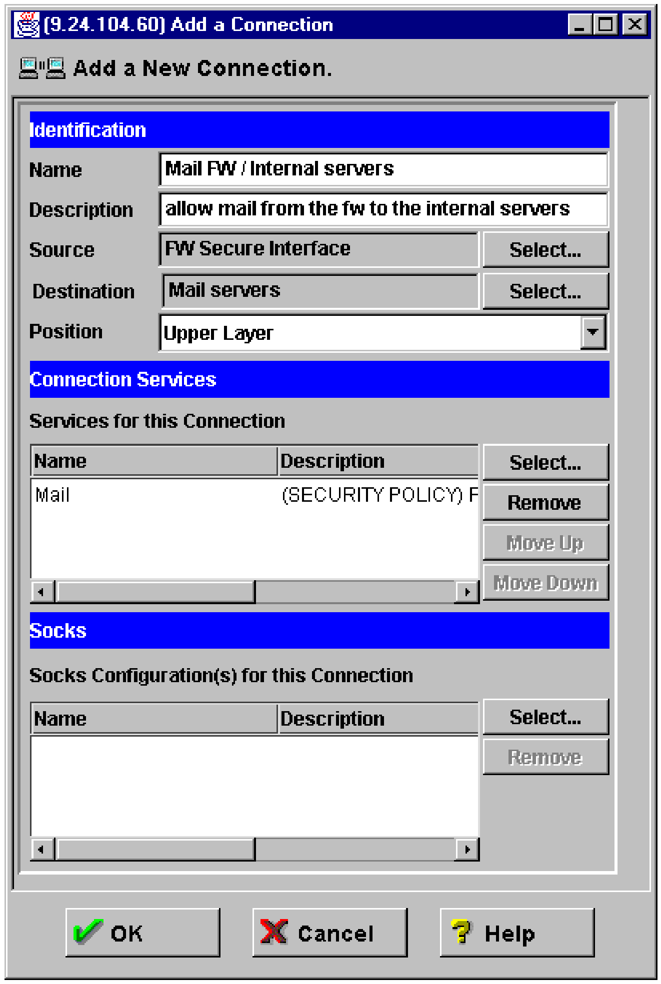

Figure 125. Connection allowing mail from the firewall to the internal servers

3. Allow incoming connections from the internal mail server(s) to the firewall. Figure 126 shows the connection we created to allow this traffic.

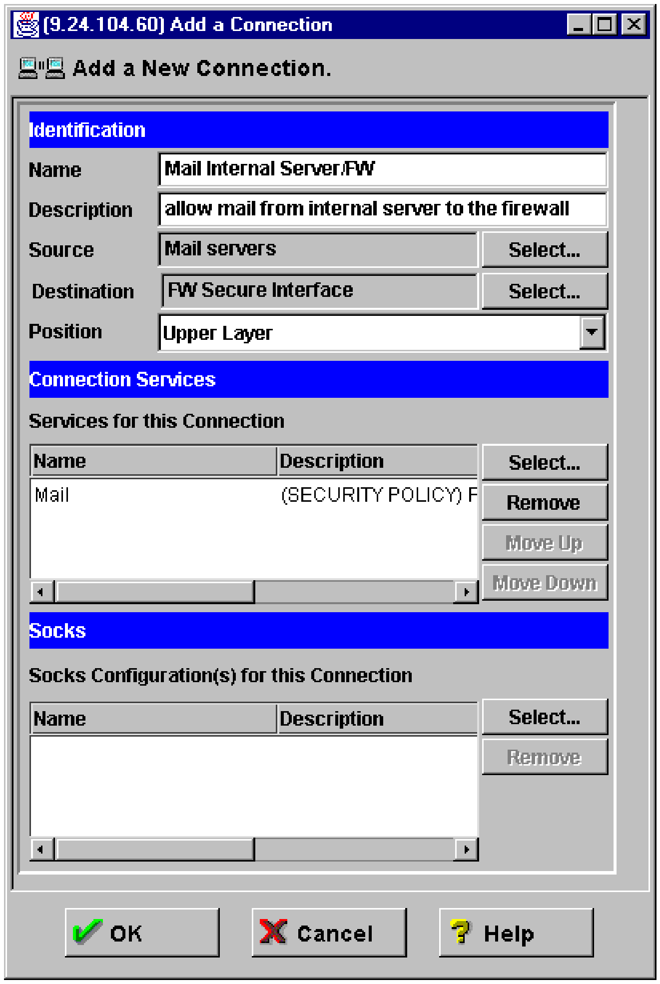

Figure 126. Connection allowing mail from the internal servers to the firewall

4. Allow outgoing connections from the firewall to any server in the Internet.

   Figure 127 shows the connection we created to allow this traffic.

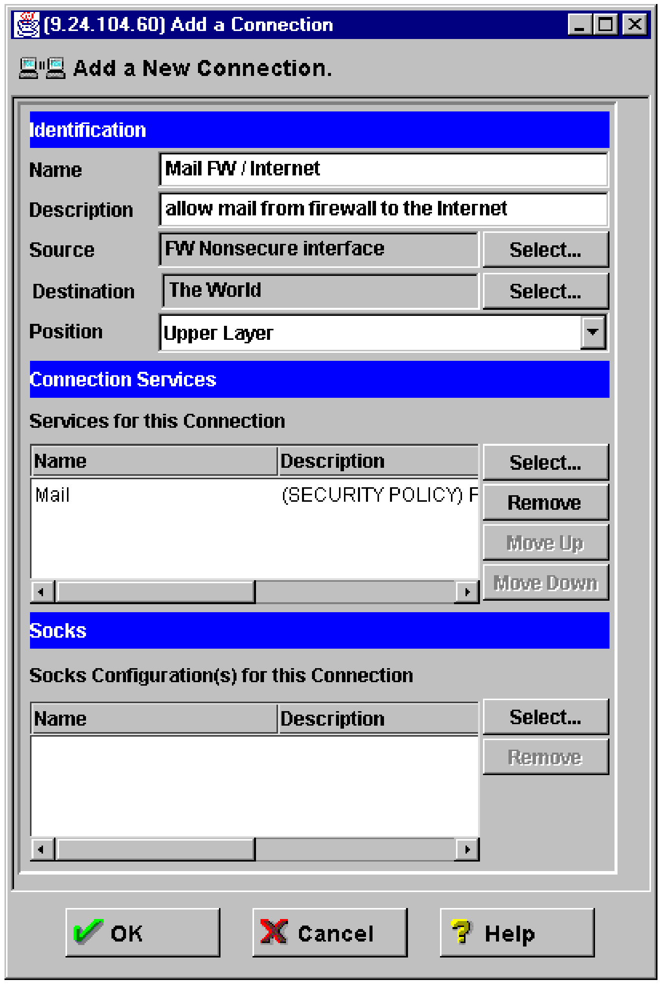

Figure 127. Connection allowing mail from the firewall to server in the Internet

8.3.4.2 Configuring the Secure Mail Proxy

We are using the examples of two companies, and we have set up both firewalls differently, so we can have more possible scenarios to discuss.

In this section we cover the configuration we made for each secure network in the Secure Mail Proxy, and eventual changes are mentioned on the respective scenario.

Company-A firewall

Company-A has one internal mail server, which is a-mail.a-corp-secure.ibm.com. It uses two mail domains: the secure mail domain is a-corp-secure.ibm.com and the nonsecure mail domain is a-corp.ibm.com. To keep this environment transparent to the user, this firewall uses domain name hiding, so the secure domain is translated into the nonsecure domain for outgoing messages and the nonsecure domain is translated into the secure domain for incoming messages.

We also turned off all security options in the Proxy Characteristics window.

In the window Managed Mail Domains, we added the nonsecure domain, as shown in Figure 128:

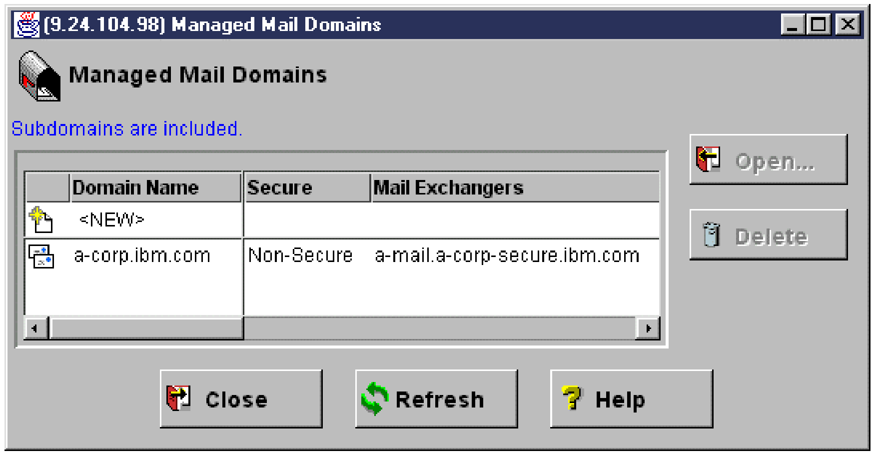

Figure 128. Managed domain for Company-A

Domain name hiding was also configured, and we are explaining about it in“Outgoing mail” on page 206.

Company B firewall

Company B has two internal mail servers: b-mail.b-corp-secure.ibm.com, which receives e-mail for the domain b-corp.ibm.com and c-mail.b-corp-secure.ibm.com, which receives e-mails for the Company C domain, c-corp.ibm.com (remember that this domain is only used for e-mail).

In this network, we decided not to use domain name hiding. So the mail servers are configured to use the nonsecure domain both for e-mails from the internal network and from the Internet.

We turned off all security options in the Proxy Characteristics window for this server too.

In the Managed Mail Domains window, we added both non-secure domains, as shown in Figure 129:

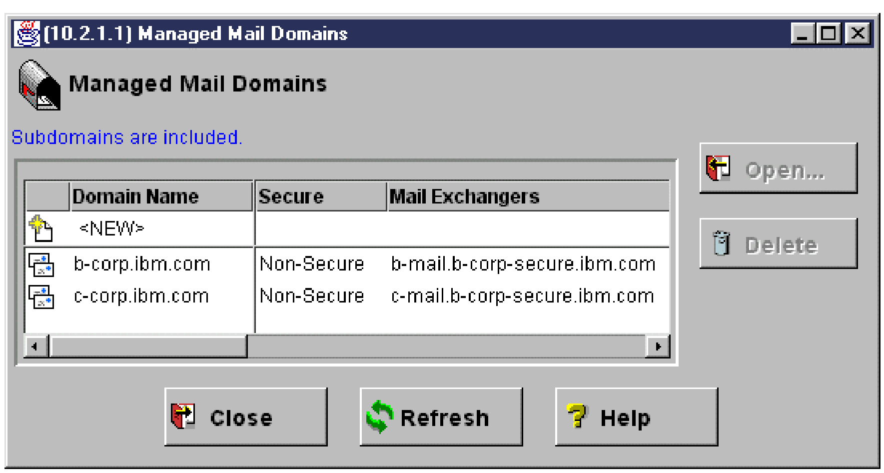

Figure 129. Managed domains for Company B

This configuration will be changed for some scenarios according to the environment we want to simulate.

Logging
In both machines we are using the main log file system (/var/log) to store the logs and for temporary work space for the Secure Mail Proxy. We edited the file /etc/security/fwsecuremail.cfg and changed the following variables:

```
# working directory
componenttoolkit.workdir:              /var/log/sb/work
# Log directory
componenttoolkit.logdir:               /var/log/sb/log
# temporary log storage
SMTPSB.TEMPORARY_STORAGE:      /var/log/sb
# The level of logging desired
SMTPSB.LOG_LEVEL:              IWE
# The directory where the logs will be stored.
SMTPSB.LOGGING_DIR:            /var/log/sb
# The directory where the accounting logs will be stored.
SMTPSB.ACCTING_DIR:            /var/log/sb/sblog
```

As shown in this list, we also selected the level of logging to be I (information), W (warning) and E (errors). On the examples we are discussing in the next section we mention the log entries. They are identified by the first letter of each line, which will be I, W or E according to the level of each entry.

8.3.5 Lab scenarios

Our goal in this section was to set up different environments for several common situations, and try to provide some guidelines for implementing the mail environment for each of these situations.

Note that the scenarios are independent from each other. This means that these scenarios are not a sequence that you should follow. You should rather identify the scenarios that best fit you, and combine the information to come up with a final configuration for your environment.

8.3.5.1 Multiple internal servers using separate mail domains

For this scenario we will use Company B network. We set up two internal mail servers, and each mail server is receiving e-mail for two different domains, as explained in 8.3.4.2, “Configuring the Secure Mail Proxy” on page 199.

The mail server b-mail.b-corp-secure.ibm.com is receiving e-mail for domain b-corp.ibm.com and the mail server c-mail.b-corp-secure.ibm.com is receiving e-mail for the domain c-corp.ibm.com.

The Secure Mail Proxy configuration is shown in Figure 129 on page 200.

When an e-mail arrives at the firewall from the nonsecure network, Secure Mail Proxy checks which is the destination domain. If it is b-corp.ibm.com, it establishes a connection with b-mail.b-corp-secure.ibm.com to receive this e-mail. If the domain is c-corp.ibm.com, then the connection is established with c-mail.c-corp.ibm.com. This is shown in the following log entries:

```
I|19990909 111051 |16232|   smtp_add_alternates()|Applied Map Routes to cris@b-c
orp.ibm.com
I|19990909 111051 |16232|SMTP_Connection::SMTP_Connection|Opened a Connection to
 b-mail.b-corp-secure.ibm.com(25)
I|19990909 111051 |16232|   SMTP_Connection::Rcpt|Sent recipient <cris@b-corp.ib
m.com> to: b-mail.b-corp-secure.ibm.com

I|19990909 111128 |17920|   smtp_add_alternates()|Applied Map Routes to cris@c-c
orp.ibm.com
I|19990909 111128 |17920|SMTP_Connection::SMTP_Connection|Opened a Connection to
 c-mail.b-corp-secure.ibm.com(25)
I|19990909 111129 |17920|   SMTP_Connection::Rcpt|Sent recipient <cris@c-corp.ib
m.com> to: c-mail.b-corp-secure.ibm.com
```

This log shows on the first three entries that an e-mail was received for domain b-corp.ibm.com and it was forwarded to b-mail.b-corp-secure.ibm.com.

The last three entries show another e-mail coming in. This time the recipient domain is c-corp.ibm.com, so the e-mail is forwarded to c-mail.b-corp-secure.ibm.com.

8.3.5.2 Multiple internal servers using load balancing

In this scenario, we are configuring the second mail server, c-mail.b-corp-secure.ibm.com, to receive e-mail also for the domain b-corp.ibm.com. Now we have two machines in our secure network that are mail exchangers for the domain b-corp.ibm.com.

We also changed the Secure Mail Proxy configuration. We opened the Managed Mail Domains window and added the machine c-mail.b-corp-secure.ibm.com as a second server for the domain b-corp.ibm.com, as shown in Figure 130:

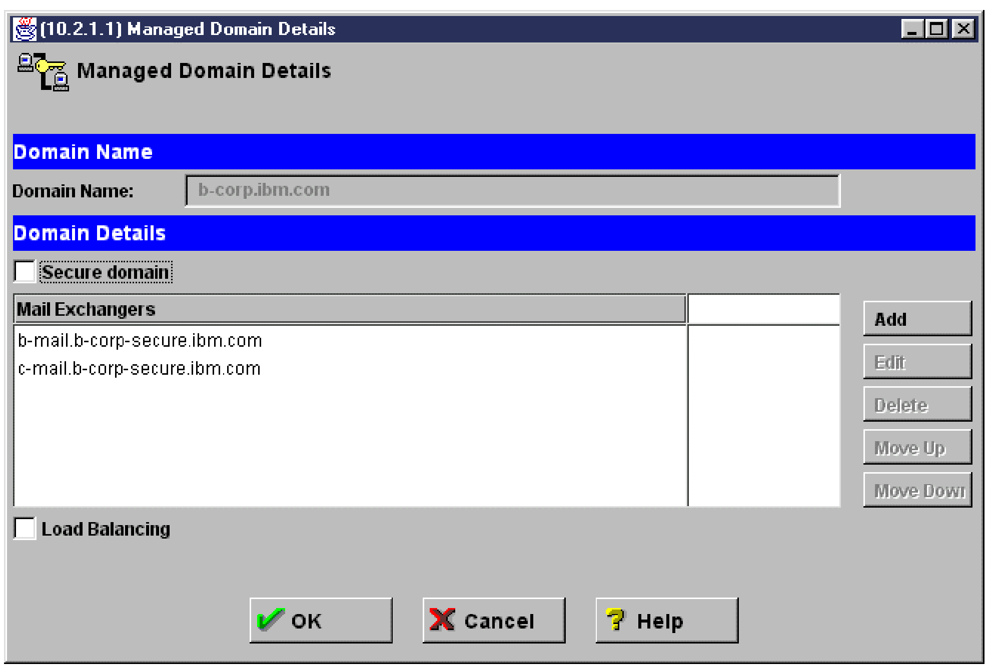

Figure 130. Including new mail server for the domain b-corp.ibm.com

In our first test, we did not use load balancing.

The user in Company A domain sent several messages to Company B. We can see the messages coming in the log (the e-mail is being sent to cris@b-corp.ibm.com):

```
I|19990908 124838 |08616|   smtp_add_alternates()|Applied Map Routes to cris@b-c
orp.ibm.com
I|19990908 124838 |08616|SMTP_Connection::SMTP_Connection|Opened a Connection to
 b-mail.b-corp-secure.ibm.com(25)
I|19990908 124838 |08616|   SMTP_Connection::Rcpt|Sent recipient <cris@b-corp.ib
m.com> to: b-mail.b-corp-secure.ibm.com
I|19990908 124919 |07104|   smtp_add_alternates()|Applied Map Routes to cris@b-c
orp.ibm.com
I|19990908 124919 |07104|SMTP_Connection::SMTP_Connection|Opened a Connection to
 b-mail.b-corp-secure.ibm.com(25)
I|19990908 124919 |07104|   SMTP_Connection::Rcpt|Sent recipient <cris@b-corp.ib
m.com> to: b-mail.b-corp-secure.ibm.com
I|19990908 124926 |09360|   smtp_add_alternates()|Applied Map Routes to cris@b-c
orp.ibm.com
I|19990908 124926 |09360|SMTP_Connection::SMTP_Connection|Opened a Connection to
 b-mail.b-corp-secure.ibm.com(25)
I|19990908 124926 |09360|   SMTP_Connection::Rcpt|Sent recipient <cris@b-corp.ib
m.com> to: b-mail.b-corp-secure.ibm.com
```

The log shows that three messages were received, and all of them were sent to the server b-mail.b-corp-secure.ibm.com.

To simulate a failure in this server, we stopped the mail service in b-mail.b-corp-secure.ibm.com, and tried to send e-mail from Company-A again. This time, it tries to connect to b-mail.b-corp-secure.ibm.com and the connection is refused (because the server is not listening to SMTP port). Then, it connects directly to c-mail.b-corp-secure.ibm.com and delivers the e-mail successfully (see the following entries from the log).

```
I|19990908 142200 |16152|   smtp_add_alternates()|Applied Map Routes to cris@b-c
orp.ibm.com
I|19990908 142200 |16152|SMTP_Connection::SMTP_Connection|Opened a Connection to
 b-mail.b-corp-secure.ibm.com(25)
I|19990908 142200 |16152|   SMTP_Connection::Rcpt|Sent recipient <cris@b-corp.ib
m.com> to: b-mail.b-corp-secure.ibm.com
I|19990908 142247 |14118|   smtp_add_alternates()|Applied Map Routes to cris@b-c
orp.ibm.com
W|19990908 142247 |14118|      Channel::Channel()|Could not connect to b-mail.b-
corp-secure.ibm.com, (79), A remote host refused an attempted connect operation.
W|19990908 142247 |14118|SMTP_Connection::SMTP_Connection|Could not connect to h
ost b-mail.b-corp-secure.ibm.com (MX:b-mail.b-corp-secure.ibm.com)
W|19990908 142247 |14118|SMTP_Connection::SMTP_Connection|Could not connect to a
ny MX host for b-mail.b-corp-secure.ibm.com
I|19990908 142247 |14118|     smtp_assign_recip()|Could not assign mail to any o
pen channel
I|19990908 142247 |14118|SMTP_Connection::SMTP_Connection|Opened a Connection to
 c-mail.b-corp-secure.ibm.com(25)
I|19990908 142248 |14118|   SMTP_Connection::Rcpt|Sent recipient <cris@b-corp.ib
m.com> to: c-mail.b-corp-secure.ibm.com
```

If load balancing is not being used, the Secure Mail Proxy uses the list of managed domains as a preference list. It always tries to connect to the first one on the list. If that server does not respond, it tries the next one, and so on.

For the next test, we turned on the load balancing on the managed domain list (see Figure 130 on page 202). We did the same tests again, sending a few messages to cris@b-corp.ibm.com with both mail servers running and accepting connections.

This is an excerpt from the log that shows the results of this test:

```
I|19990908 143425 |16154|   smtp_add_alternates()|Applied Map Routes to cris@b-c
orp.ibm.com
I|19990908 143425 |16154|SMTP_Connection::SMTP_Connection|Opened a Connection to
 c-mail.b-corp-secure.ibm.com(25)
I|19990908 143426 |16154|   SMTP_Connection::Rcpt|Sent recipient <cris@b-corp.ib
m.com> to: c-mail.b-corp-secure.ibm.com
I|19990908 143434 |14120|   smtp_add_alternates()|Applied Map Routes to cris@b-c
orp.ibm.com
I|19990908 143434 |14120|SMTP_Connection::SMTP_Connection|Opened a Connection to
 b-mail.b-corp-secure.ibm.com(25)
I|19990908 143434 |14120|   SMTP_Connection::Rcpt|Sent recipient <cris@b-corp.ib
m.com> to: b-mail.b-corp-secure.ibm.com
I|19990908 143439 |05984|   smtp_add_alternates()|Applied Map Routes to cris@b-c
orp.ibm.com
I|19990908 143439 |05984|SMTP_Connection::SMTP_Connection|Opened a Connection to
 c-mail.b-corp-secure.ibm.com(25)
I|19990908 143440 |05984|   SMTP_Connection::Rcpt|Sent recipient <cris@b-corp.ib
m.com> to: c-mail.b-corp-secure.ibm.com
I|19990908 143448 |08634|   smtp_add_alternates()|Applied Map Routes to cris@b-c
orp.ibm.com
I|19990908 143448 |08634|SMTP_Connection::SMTP_Connection|Opened a Connection to
 b-mail.b-corp-secure.ibm.com(25)
I|19990908 143448 |08634|   SMTP_Connection::Rcpt|Sent recipient <cris@b-corp.ib
m.com> to: b-mail.b-corp-secure.ibm.com
```

The log shows that the load balancing is done using round-robin algorithm.

Note

In the managed domain configuration, you can use up to mail exchangers for each domain.

8.3.5.3 Multiple domains in the same server

In this scenario, we show an example where one single internal mail server can handle more than one domain.

We made a change in the configuration of the Company B environment: now our server b-mail.b-corp-secure.ibm.com is going to handle both mail domains. That means that for every incoming e-mail with recipients in the domain b-corp.ibm.com or c-corp.ibm.com, the e-mail will be forwarded to the server b-mail.b-corp-secure.ibm.com.

The configuration on Secure Mail Proxy is simple. You just have to add one entry for each domain in managed mail domains list, and these entries will use the same mail exchanger. In our lab, we already had an entry for c-corp.ibm.com using the mail exchanger c-mail.b-corp-secure.ibm.com, so we changed it to use b-mail.b-corp-secure.ibm.com, as shown in Figure 131.

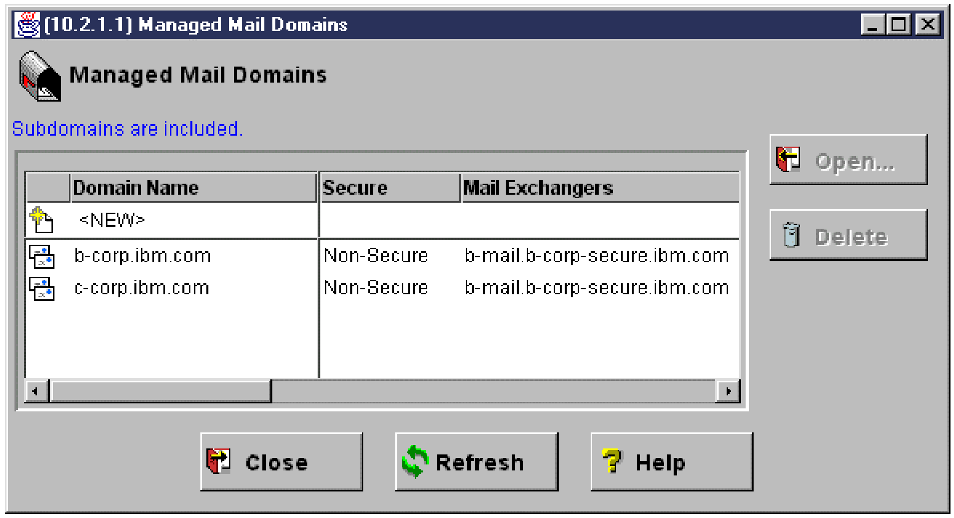

Figure 131. Mail domains using the same mail exchanger

We did not have to change the external DNS configuration because we already have the c-corp.ibm.com domain using the firewall as mail exchanger. If you are adding a new domain, you must add an MX record for this domain in the non-secure DNS server, and the mail exchanger is also the firewall.

You can check the DNS configuration with the command `nslookup` to make sure it is correct (run this command in the nonsecure network):

```
# nslookup
Default Server:  dns.gw.itso.ral.ibm.com
Address:  9.24.104.62

> set type=mx
> c-corp.ibm.com
Server:  dns.gw.itso.ral.ibm.com
Address:  9.24.104.62

c-corp.ibm.com  preference = 10, mail exchanger = b-gateway.gw.itso.ral.ibm.com
b-gateway.gw.itso.ral.ibm.com   internet address = 9.24.104.60
> b-corp.ibm.com
Server:  dns.gw.itso.ral.ibm.com
Address:  9.24.104.62

b-corp.ibm.com  preference = 10, mail exchanger = b-gateway.gw.itso.ral.ibm.com
b-gateway.gw.itso.ral.ibm.com   internet address = 9.24.104.60
```

In this example, the mail exchanger for both domains is the server b-gateway.gw.itso.ral.ibm.com (IP address 9.24.104.60, which is the nonsecure interface of the firewall).

The following log entries show two e-mail messages being received, one for b-corp.ibm.com and one for c-corp.ibm.com, and both being forwarded to b-mail.b-corp-secure.ibm.com.

```
I|19990909 111051 |16232|   smtp_add_alternates()|Applied Map Routes to cris@b-c
orp.ibm.com
I|19990909 111051 |16232|SMTP_Connection::SMTP_Connection|Opened a Connection to
 b-mail.b-corp-secure.ibm.com(25)
I|19990909 111051 |16232|   SMTP_Connection::Rcpt|Sent recipient <cris@b-corp.ib
m.com> to: b-mail.b-corp-secure.ibm.com

I|19990909 111128 |17920|   smtp_add_alternates()|Applied Map Routes to cris@c-c
orp.ibm.com
I|19990909 111128 |17920|SMTP_Connection::SMTP_Connection|Opened a Connection to
b-mail.b-corp-secure.ibm.com(25)
I|19990909 111129 |17920|   SMTP_Connection::Rcpt|Sent recipient <cris@c-corp.ib
m.com> to: b-mail.b-corp-secure.ibm.com
```

8.3.5.4 Domain name hiding

This first scenario shows e-mail being sent between Company-A and Company B, and the difference when using or not the domain name hiding.

In the examples, we just pasted the part of the header added by Company-A servers (the internal mail server and the firewall) to avoid confusions with the other environment.

Outgoing mail

In this test, we are sending e-mail out from Company-A to Company B, which simulates a situation where one company sends e-mail using its Secure Mail Proxy to another company that also has the firewall installed, and it will receive this e-mail by the Secure Mail Proxy as well.

The following header was pasted from e-mail that was sent by the user cris@a-corp.ibm.com to the user cris@b-corp.ibm.com.

```
Received: by a-gateway.gw.itso.ral.ibm.com from a-mail.a-corp-secure.ibm.com
([10.1.1.2]); Wed, 08 Sep 1999 20:26:33 GMT
Received: (from cris@localhost) by a-mail.a-corp-secure.ibm.com (1.0) id PAA13422
for cris@b-corp.ibm.com; Wed, 8 Sep 1999 15:31:04 -0500
Date: Wed, 8 Sep 1999 15:31:04 -0500
From: cris@a-corp.ibm.com
Message-ID: <199909082031.PAA13422@a-mail.a-corp-secure.ibm.com>
To: cris@b-corp.ibm.com
Subject: test - no domain name hiding
```

The first "Received" line was added by the Secure Mail Proxy and the second one was added by the internal mail server (a-mail.a-corp-secure.ibm.com). Note that we can see information about the hostname of the machines and the IP address of the internal mail server (10.1.1.2).

For the next test, we configured domain name hiding for a-corp-secure.ibm.com and 10.2.1.2 (we have to do that if we do not want to reveal the IP address of internal machines). The configuration is shown in Figure 132:

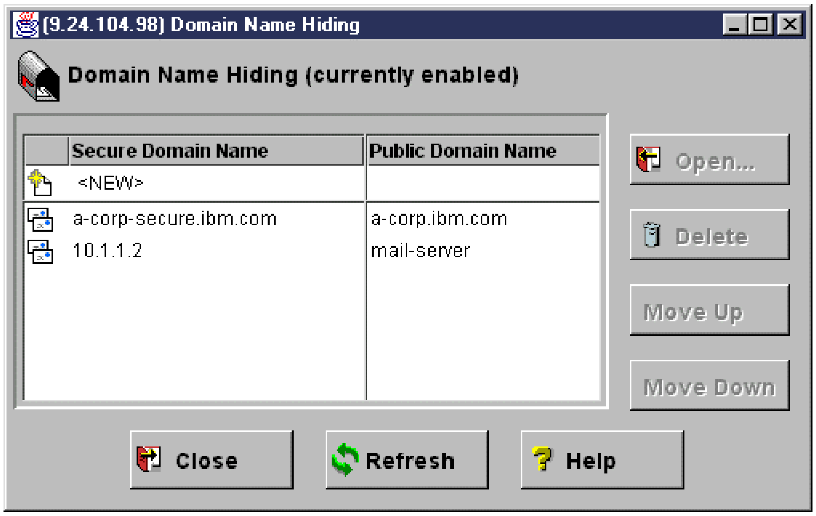

Figure 132. Domain name hiding configuration

The following listing shows how the message arrive at the final mail server (note that some fields of the header were changed).

```
Received: by a-gateway.gw.itso.ral.ibm.com from a-corp.ibm.com mail-server]);
Wed, 08 Sep 1999 20:27:09 GMT
Received: (from cris@localhost) by a-corp.ibm.com (1.0) id PAA15266 for
cris@b-corp.ibm.com; Wed, 8 Sep 1999 15:31:40 -0500
Date: Wed, 8 Sep 1999 15:31:40 -0500
From: cris@a-corp.ibm.com
Message-ID: <199909082031.PAA15266@a-corp.ibm.com>
To: cris@b-corp.ibm.com
Subject: test using domain name hiding
```

The Secure Mail Proxy adds the following entries in the log for this message:

```
I|19990908 154116 |05990|   smtp_add_alternates()|Applied Map Routes to cris@b-c
orp.ibm.com
I|19990908 154117 |05990|SMTP_Connection::SMTP_Connection|Opened a Connection to
b-mail.b-corp-secure.ibm.com(25)
I|19990908 154117 |05990|   SMTP_Connection::Rcpt|Sent recipient <cris@b-corp.ib
m.com> to: b-mail.b-corp-secure.ibm.com
```

Incoming mail

Now we are sending e-mail from Company B to Company -A, replying to the message that was sent in the previous section.

The following e-mail was sent by the user cris@b-corp.ibm.com to the user cris@a-corp.ibm.com.

```
From cris@b-corp.ibm.com Wed Sep  8 18:18:48 1999
Received: from a-gateway.gw.itso.ral.ibm.com (gw.a-corp-secure.ibm.com [10.1.1.1
])
        by a-mail.a-corp-secure.ibm.com (1.0) with ESMTP id SAA14210
        for <cris@a-corp.ibm.com>; Wed, 8 Sep 1999 18:18:47 -0500
Received: by a-gateway.gw.itso.ral.ibm.com from b-gateway.gw.itso.ral.ibm.com ([
9.24.104.60]); Wed, 08 Sep 1999 23:14:14 GMT
[...]
Message-ID: <37D6DF6C.9192643F@b-corp.ibm.com>
Date: Wed, 08 Sep 1999 18:13:00 -0400
From: Cristiane <cris@b-corp.ibm.com>
To: cris@a-corp.ibm.com
Subject: test - no domain name hiding
```

The symbol [...] indicates that we removed the Received lines added by the sending mail server to avoid confusion.

The header fields do not have anything unusual, except for the Received lines added by the Secure Mail Proxy (the first one) and the internal mail server (the second one), that show the hostname of internal servers, but this is not a problem since this is an incoming message (this information is not going outside).

We repeated the test using domain name hiding in the Company-A firewall. Look at the following header:

```
From cris@b-corp.ibm.com Wed Sep  8 18:10:10 1999
Received: from a-gateway.gw.itso.ral.ibm.com (gw.a-corp-secure.ibm.com [10.1.1.1
])
        by a-mail.a-corp-secure.ibm.com (1.0) with ESMTP id SAA19560
        for <cris@a-corp.ibm.com>; Wed, 8 Sep 1999 18:10:09 -0500
Received: by a-gateway.gw.itso.ral.ibm.com from b-gateway.gw.itso.ral.ibm.com ([
9.24.104.60]); Wed, 08 Sep 1999 23:05:37 GMT
[...]
Message-ID: <37D6DD67.D14BE19D@b-corp.ibm.com>
Date: Wed, 08 Sep 1999 18:04:23 -0400
From: Cristiane <cris@b-corp.ibm.com>
To: cris@a-corp-secure.ibm.com
Subject: test - using domain name hiding
```

The fields are the same as when we do not use domain name hiding because this is an incoming message. The only difference here is the field "To:", which was translated into the secure domain. That is automatically done when you are using domain name hiding; it changes the nonsecure domain into a secure domain when the message is coming in.

8.3.5.5 Overflow server on the firewall

A simple configuration for an overflow server is to use sendmail on the same box as the firewall. When the Secure Mail Proxy needs to send any message to the overflow server, it just connects to `localhost` using one port specified by the administrator (remember that port 25 is already in use by the Secure Mail Proxy).

In the configuration we made, sendmail is using DNS to find the mail exchanger for each domain, including the internal domains. We chose this approach because it is simple to configure and troubleshoot, and you can easily handle more than one internal mail exchanger and different internal mail domains.

Note that the configuration we show here was done using sendmail V8.8.8. Make sure you are using the same version or check if the configuration is valid for your version of sendmail.

Configuring DNS

First, we configure the internal DNS server. We add the mail domains to the boot file. See in the following listing the lines that were added to /etc/named.boot:

```
primary        b-corp.ibm.com  b-corp.data
primary        c-corp.ibm.com  c-corp.data
```

The mail exchanger (MX record) for the domain b-corp.ibm.com is the machine b-mail.b-corp-secure.ibm.com. And this machine is also the nameserver for this domain. So the data file (b-corp.data) was configured as follows:

```
; NAME          TTL     CLASS   TYPE    RDATA
;
; setting default domain to "b-corp-secure.ibm.com"
;
@               9999999 IN      SOA     b-corp.ibm.com. root.b-corp.ibm.com. (
                                        1.1             ; Serial
                                        3600            ; Refresh
                                        300             ; Retry
                                        3600000         ; Expire
                                        86400 )         ; Minimum
                9999999 IN      NS      b-mail.b-corp-secure.ibm.com.
b-corp.ibm.com. 9999999 IN      MX      10      b-mail.b-corp-secure.ibm.com.
```

Note that the data file for this domain only contains an MX record. The hostname b-mail.b-corp-secure.ibm.com is configured in the data file of the b-corp-secure.ibm.com domain (see Chapter 5, “Domain Name System (DNS) Service” on page 77 for more information about this configuration).

Similar to the b-corp.ibm.com domain, c-corp.ibm.com uses the machine c-mail.b-corp-secure.ibm.com as the mail exchanger and b-mail.b-corp-secure.ibm.com is the name server. The data file for the domain c-corp.ibm.com is:

```
; NAME          TTL     CLASS   TYPE    RDATA
;
@               9999999 IN      SOA     c-corp.ibm.com. root.c-corp.ibm.com. (
                                        1.1             ; Serial
                                        3600            ; Refresh
                                        300             ; Retry
                                        3600000         ; Expire
                                        86400 )         ; Minimum
                9999999 IN      NS      b-mail.b-corp-secure.ibm.com.
c-corp.ibm.com. 9999999 IN      MX      10      c-mail.b-corp-secure.ibm.com.
```

This data file also contains only an MX record. After making these changes, it is necessary to refresh the named daemon.

Now, when sendmail receives an e-mail from the Secure Mail Proxy, it queries the DNS for the mail exchanger for any domain. When it is one of the internal domains, it receives the MX record as we configured above. Any other domain will be queried on the Internet, and the connection will be done directly to the corresponding external server.

If you also use a separate secure mail domain, you should add the MX record for this domain too (only in the internal DNS server).

Configuring sendmail

Sendmail configuration is done by editing the file /etc/sendmail.cf.

Add the following lines to make sendmail bind to the address 127.0.0.1 and port 2500 (we chose this port because it is the default port for the overflow server on the GUI).

```
# SMTP daemon options
O DaemonPortOptions=Port=2500
O DaemonPortOptions=Address=127.0.0.1
```

Forcing sendmail to bind to the localhost address is recommended for two main reasons:

1. It does not need any filter to allow this connection. Even if you have problems with the filters, the Secure Mail Proxy can still connect to the overflow server.
2. With this configuration, sendmail will not listen either on the secure or on the nonsecure interfaces, so it is not possible to connect to this server from outside.

We also recommend you change the information that is inserted in the header of the message. First, locate the following lines in the configuration file:

```
#########################
#   Format of headers   #
#########################

H?P?Return-Path: <$g>
HReceived: $?sfrom $s $.$?_($?s$|from $.$_) $.by $j ($v/$Z)$?r with $r$. id $i$?
u for $u$.; $b
```

Change the HReceived line to avoid including information about this overflow server (type it as a single line; do not press Enter at the end):

```
HReceived: $?sfrom $s $.$?_ $.by b-gateway.gw.itso.ral.ibm.com with $r$. id $i$?
u for $u$.; $b
```

In the preceding example, type the nonsecure name of your firewall in the same place where you read "b-gateway.gw.itso.ral.ibm.com".

Start sendmail and check which address and port it is listening to:

```
# startsrc -s sendmail -a "-bd -q30m"
0513-059 The sendmail Subsystem has been started. Subsystem PID is ...
# netstat -an | grep 2500
tcp4       0      0  127.0.0.1.2500         *.*                LISTEN
```

Uncomment the following line from /etc/rc.tcpip, so sendmail is automatically restarted on every reboot.

```
start /usr/lib/sendmail "$src_running" "-bd -q${qpi}"
```

Sendmail logging

You can set up separate logging for this overflow server. Sendmail sends its log entries to AIX syslog daemon. All you have to do is edit the /etc/syslog.conf file, and add the following line (do not change anything else on this file):

```
mail.info       /var/log/sendmail.info
```

Run the following commands to create the log file and refresh the syslogd daemon:

```
# touch /var/log/sendmail.info
# refresh -s syslogd
0513-095 The request for subsystem refresh was completed successfully.
```

Sendmail queue

Use the command `mailq` to keep track of the sendmail queue.

We recommend you create a separate file system for this mail queue, so it will not be using /var space.

Testing the overflow server

We tested the overflow server by forcing an error situation on the internal server. We stopped the SMTP service in the c-mail.b-corp-secure.ibm.com server, so it does not listen to port 25. While this server is out, the user cris@a-corp.ibm.com sends e-mail to the user cris@c-corp.ibm.com.

Since the internal mail server was down, Secure Mail Proxy delivered this e-mail to the overflow server. The following listing is an excerpt from the Secure Mail Proxy log file.

```
I|19990907 184519 |08588|   smtp_add_alternates()|Applied Map Routes to cris@
c-corp.ibm.com
W|19990907 184519 |08588|      Channel::Channel()|Could not connect to c-mail.
b-corp-secure.ibm.com, (79), A remote host refused an attempted connect operatio
n.
W|19990907 184519 |08588|SMTP_Connection::SMTP_Connection|Could not connect to h
ost c-mail.b-corp-secure.ibm.com (MX:c-mail.b-corp-secure.ibm.com)
W|19990907 184519 |08588|SMTP_Connection::SMTP_Connection|Could not connect to a
ny MX host for c-mail.b-corp-secure.ibm.com
I|19990907 184519 |08588|     smtp_assign_recip()|Could not assign mail to any o
pen channel
I|19990907 184519 |08588|SMTP_Connection::SMTP_Connection|Opened a Connection to
 localhost(2500)
I|19990907 184519 |08588|   SMTP_Connection::Rcpt|Sent recipient <cris@c-corp
.ibm.com> to: localhost
```

The sendmail log we set up in the previous section (the file /var/log/sendmail.info) contains the following entries:

```
Sep  7 18:45:20 b-gateway sendmail[21416]: SAA21416: from=<cris@.a-corp-secur
e.ibm.com>, size=708, class=0, pri=30708, nrcpts=1, msgid=<Pine.A41.4.05.9909
071846260.15772-100000@a-mail.a-corp.ibm.com>, proto=ESMTP, relay=localhost
[127.0.0.1]
Sep  7 18:45:20 b-gateway sendmail[21930]: SAA21416: to=<cris@c-corp.ibm.c
om>, delay=00:00:01, xdelay=00:00:00, mailer=esmtp, relay=c-mail.b-corp-sec
ure.ibm.com. [::ffff:10.2.1.9], stat=Deferred: Connection refused by c-mail
.b-corp-secure.ibm.com.
```

The first entry in this log shows the message being received by sendmail. It tries to send the message right away, but it also cannot establish a session with the mail exchanger, so it puts this message in the queue (stat=Deferred).

Running the command mailq, we can confirm that the message is queued:

```
# mailq
There is 1 request in the mail queue
---QID---- --Size-- -----Q-Time----- --------------Sender/Recipient-------------

SAA21416       10 Tue Sep  7 18:45 <cris@a-corp.ibm.com>
                 (Deferred: Connection refused by c-mail.b-corp-secure.ibm.c)
                                    <cris@c-corp.ibm.com>
```

This message is going to be handled from now on by sendmail, so to keep track of it we have to check the sendmail queue and log.

At this moment, we started the mail service in the internal mail server, so it is accepting connections at port 25. We forced sendmail to process the queue by running the command sendmail -q, then we checked the queue again and the message was gone:

```
# sendmail -q

# mailq
The mail queue is empty.
```

We checked the log, and confirmed the message was successfully delivered to the internal mail server (stat=Sent):

```
Sep  7 18:48:21 b-gateway sendmail[21658]: SAA21416: to=<cris@c-corp.ibm.c
om>, delay=00:03:02, xdelay=00:00:01, mailer=esmtp, relay=c-mail.b-corp-sec
ure.ibm.com. [::ffff:10.2.1.9], stat=Sent (OK)
```

We also checked the header of the message that was received by the user cris@c-corp.ibm.com. It has an extra Received entry that was added by sendmail, but it does not reveal any information about this server, since we changed this line in sendmail.cf file. The information it shows is:

```
Received: from b-gateway.gw.itso.ral.ibm.com by b-gateway.gw.itso.ral.ibm.com with
ESMTP id SAA21416 for <cris@c-corp.ibm.com>; Tue, 7 Sep 1999 18:45:19 -0400
```

This line may seem meaningless for the final user (it shows that the message was received by b-gateway from itself), but it is useful for the administrator to track the path of the message. This entry can be easily identified as being added by the overflow server (among all other "Received" entries), since it was sent and received by the same machine, so the administrator knows that the overflow server handled this message.

8.3.5.6 Excluding mail domains

Using the Company B network, we added the domain spamdomain.com to the excluded domain list, as shown in Figure 133:

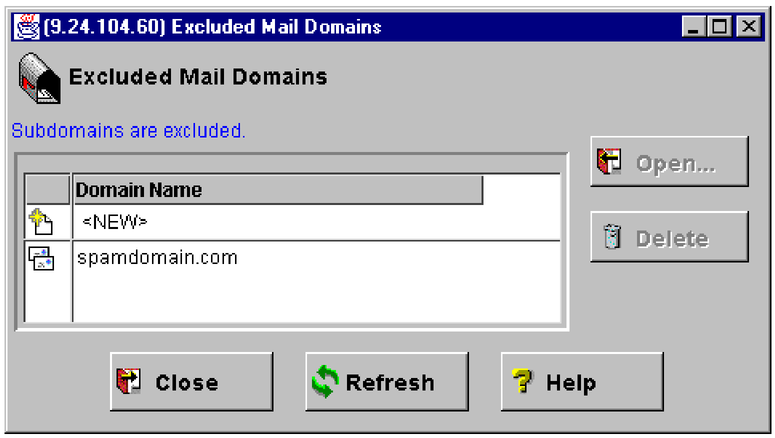

Figure 133. Excluded mail domain configuration

We configured an external mail server using sendmail, and we set up a masquerade to make all outgoing messages from this server use the domain spamdomain.com.

Figure 134 shows the output from `mail -v` when trying to send e-mail to b-corp.ibm.com from the nonsecure network using the excluded domain spamdomain.com.

```
cris@b-corp.ibm.com... Connecting to b-corp.ibm.com. via smtp...
220-b-gateway.gw.itso.ral.ibm.com Connection Established.
220 ESMTP
>>> EHLO test.gw.itso.ral.ibm.com
250-Hello, [9.24.104.70]
250-SIZE=20971520
250-8BITMIME
250-DSN
250 Okay
>>> MAIL From:<cris@spamdomain.com>
551 b-gateway.gw.itso.ral.ibm.com will not accept mail from
spamdomain.com.

cris@b-corp.ibm.com... Service unavailable
/home/cris/dead.letter... Saved message in /home/crismari/dead.letter
/home/cris/dead.letter... Closing connection to b-corp.ibm.com.
>>> QUIT
```

Figure 134. Sending e-mail using an excluded mail domain

8.4 Advanced configuration

Configuration of the Secure Mail Proxy is achieved via the administrative GUI.

Outside of the GUI environment there are several configuration tasks we have found useful. These tasks are related to the one-time setup of various scenarios that may arise in your configuration.

Configuration Files

Always keep a backup copy of any configuration file you edit. Hint: Use the copy command, for example: cp thisfile thisfile.orig.

Always be prepared to compare your changes against the original configuration. Hint: Use the diff command, for example: diff thisfile thisfile.orig.

Always know how to abort the editor in case you corrupt your file. Hint: Using the vi editor, this is achieved via the key sequence <Esc>q!.

8.4.1 Refresh

The following command needs to be executed after any manual changes to the configuration file.

```
# fwsecuremail cmd=refresh
```

8.4.2 Disable

The Secure Mail Proxy can be permanently disabled. This may be required for the following reasons:

- Your security policy does not require the use of SMTP on the firewall and you choose not to have this process running.
- You will be using an alternate SMTP process, such as sendmail.

Edit the file /etc/rc.tcpip and comment out the line that starts the smtpsb process.

The following except shows how the line will appear commented out in the file:

```
#/usr/sbin/smtpsb &                                         #FW#
```

The currently running smtpsb process can be stopped as described in the next section, or if you prefer you can reboot the firewall, so the smtpsb process will not start again.

8.4.3 Stop

If the Secure Mail proxy needs to be terminated immediately, run the following command:

```
# fwsecuremail cmd=shutdown
```

8.4.3.1 Has is really stopped?

The following checks can be performed to confirm the process has been terminated and no residual mail sockets are in use.

Confirm the process is not currently running via the the following command:

```
# ps -ef | grep smtpsb
```

The command will not return any output if the process has stopped.

The following command will look at remote or local TCP connections using the default SMTP port (25). No output indicates SMTP activity has stopped.

```
# netstat -an | grep '.25 '
```

8.4.4 Start

After the Secure Mail Proxy has been stopped (as shown in 8.4.3, "Stop" on page 215) it may be necessary to start it manually.

The following command will manually start the process:

```
# /usr/sbin/smtpsb &
```

8.4.4.1 Has it really started?

The following checks can be performed to confirm the process is currently running and ready to accept SMTP connections.

Confirm the process is currently running via the following command and output:

```
# ps -ef | grep smtpsb
    root  7546 19164   0 14:46:55 pts/10  0:00 smtpsb
    root  7804 19164   0 14:46:55 pts/10  0:00 smtpsb
    root  8062 19164   0 14:46:55 pts/10  0:00 smtpsb
    root  8320 19164   0 14:46:55 pts/10  0:00 smtpsb
    root  8586 19164   0 14:46:55 pts/10  0:00 smtpsb
    root  8836 19164   0 14:46:55 pts/10  0:00 smtpsb
    root  9352 19164   0 14:46:55 pts/10  0:00 smtpsb
    root  9610 19164   0 14:46:55 pts/10  0:00 smtpsb
    root  9868 19164   0 14:46:55 pts/10  0:00 smtpsb
    root 10126 19164   0 14:46:55 pts/10  0:00 smtpsb
    root 10384 19164   0 14:46:55 pts/10  0:00 smtpsb
```

The output will display many smtpsb processes. The default configuration will create 64 processes that can process SMTP transactions in parallel.

The following command will look at remote or local sockets using the default SMTP port (25). At least one entry will be present in the LISTEN state to indicate the SMTP process is ready to accept SMTP connections:

```
# netstat -an | grep '.25 '
tcp4       0      0*.25 *.* LISTEN
```

8.4.5 Logging

No logs of the Secure Mail Proxy are collected by default.

Logs may be required for the following reasons:

- Your security policy requires auditing of SMTP transactions.
- Tracing the flow of SMTP transactions during problem determination.
- Other reasons as directed by your support organization.

Logging Filesystem

Did you create the dedicated logging filesystem as recommended in Chapter 2, “Installation” on page 7? This is especially important in the case of a Denial of Service attachment trying to fill your filesystem. The dedicated logging filesystem will minimize possible interference with other subsystems.

8.4.5.1 Configure a logging filesystem

Create a directory to store the logs. We will create a directory in the dedicated logging filesystem you had created during the installation process. The following command will create the logging directory:

```
# mkdir /var/log/mail
```

8.4.5.2 Configure Secure Mail Proxy to use logging filesystem

The Secure Mail Proxy needs to point to this directory. Edit the configuration file /etc/security/fwsecuremail.cfg and update the LOGGING_DIR attribute as follows:

```
SMTPSB.LOGGING_DIR:             /var/log/mail
```

8.4.5.3 Select logging level

You will need to select what level of logging is required from the following table:

Table 24. Logging severity levels

Code	Name	Description
G	General	A basic status message, such as a report on configuration settings.
W	Warning	A minor error has occurred, but SMTPSB has automatically recovered.
E	Error	A serious error has occurred, which has resulted in a failed connection or e-mail.
S	Security	SMTPSB has rejected a mail item or connection because of a security check.

Code	Name	Description
I	Informational	Detailed information about the operation of SMTPSB, such as a report about the success of an outbound connection.
D	Debug	Very detailed information about the operation of SMTPSB, useful primarily to support personnel.
T	Trace	Messages reporting on the actual step-by-step execution of SMTPSB. Useful only to support personnel.
i	i/o	This level records all SMTP read and write operations. Useful primarily to support personnel.

Only information you are interested in should be logged to keep the manageability of the log files reasonable.

Code selection

In practical terms only the following code groups would be used:

Table 25. Practical logging levels

Codes	Description
WES	Default recommended level. Log only warnings, errors and security-related events.
I	Include Informational logging to capture successful transactions for audit purposes.
GiTD	This verbose trace would present internal information readable only by your support organization.

Code configuration

Edit the configuration file /etc/security/fwsecuremail.cfg and update the LOG_LEVEL attribute as per your selection:

```
SMTPSB.LOG_LEVEL:              WES
```

8.4.5.4 Refresh server

You will need to refresh the Secure Mail Proxy after changing the configuration file. See 8.4.1, “Refresh” on page 214.

8.4.5.5 Viewing Logs

The log files you had configured above will be in the directory /var/log/mail.

The file names will be in the following format, representing the timestamp of creation:

```
YYYYMMDDHHMMSS.LOG
```

The following excerpt shows some warning messages when a remote SMTP server could not be contacted:

```
W|19990903 181527 |10938|      Channel::Channel()|Could not connect to
rtpmail01.raleigh.ibm.com, (79), A remote host refused an attempted connect
operation.
W|19990903 181527 |10938|SMTP_Connection::SMTP_Connection|Could not connect to
host
rtpmail01.raleigh.ibm.com (MX:raleigh.ibm.com)
```

8.4.5.6 Log file management

Log file management is automatically performed by the Secure Mail Proxy.

Log data exceeding 512 KB will cause a new log file to be created.

When the fifth log file is filled, the oldest one is automatically deleted.

The following output shows the five log files:

```
# /var/log/mail > ls -al
-rw-------   1 root sys        9776 Sep 03 18:13 19990903181027.LOG
-rw-------   1 root sys        5165 Sep 03 18:14 19990903181331.LOG
-rw-------   1 root sys         289 Sep 03 18:15 19990903181507.LOG
-rw-------   1 root sys       10317 Sep 03 18:36 19990903181938.LOG
-rw-------   1 root sys           0 Sep 03 18:36 19990903183614.LOG
```

8.4.6 Temporary storage

The Secure Mail Proxy is unlike the store-and-forward style of SMTP mailers, such as sendmail regarding storage of SMTP items. Sendmail will store the body of SMTP items in its spool directory for as long as required to deliver the item, up to a specified limited. The default on some systems is as long as five days. The Secure Mail Proxy keeps a temporary file on disk only during the lifetime of the actual TCP socket connection. After the connection is closed, the file is erased.

The default storage location for intermediate files used by the Secure Mail Proxy is the /tmp filesystem. The SMTP message body is temporarily stored here during the relaying of the message.

8.4.6.1 Why dedicated storage?

We recommend that you create a dedicated temporary storage volume for the following reasons:

- Filling the /tmp filesystem to 100% can be detrimental to the operation of AIX.
- Allows you to accept larger message sizes.
- Larger capacity disks can be taken advantage of by allowing more parallel Secure Mail Proxy threads to run.
- Increased performance by off-loading SMTP temporary storage from the default root volume group (rootvg) defined by AIX.

8.4.6.2 Configure additional disk

We have installed an addition physical volume on our firewall (hdisk1) that we will use purely as a temporary spool area for the Secure Mail Proxy. By default, hdisk0 is used by the operating system.

The following command shows the two disks we have installed in our machine:

```
# lscfg | grep disk
+ hdisk0            04-C0-00-4,0      16 Bit SCSI Disk Drive (4500 MB)
+ hdisk1            04-C0-00-5,0      16 Bit SCSI Disk Drive (4500 MB)
```

A new volume group (VG) is created using the additional disk (hdisk1) that we have installed. The following command will create the volume group:

```
# mkvg -s 8 -y mailvg hdisk1
0516-014 lcreatevg: The physical volume appears to belong to another
        volume group.
00735006bc78e311
0516-631 mkvg: Warning, all data belonging to physical
        volume hdisk1 will be destroyed.
mkvg: Do you wish to continue? y(es) n(o)? y
mailvg
```

The whole disk will be used in our case solely for the temporary spool area. The following command will create a large 4GB filesystem called /var/spool/smtpsb:

```
# crfs -v jfs -a bf=true -g mail -m /var/spool/smtpsb -Ayes -asize=8000000

Based on the parameters chosen, the new /var/spool/smtpsb JFS file system
is limited to a maximum size of 134217728 (512 byte blocks)

New File System size is 8011776
```

The following commands will mount the new filesystem, then verify the amount of storage it has available:

```
# mount /var/spool/smtpsb

# df -k /var/spool/smtpsb
Filesystem 1024-blocks Free %Used Iused %Iused Mounted on
/dev/lv03 4005888   3880100    4% 17     1% /var/spool/smtpsb
```

8.4.6.3 Configure Secure Mail proxy to use new disk

The Secure Mail Proxy needs to point to this new filesystem /var/spool/smtpsb. Edit the configuration file /etc/security/fwsecuremail.cfg to update the TEMPORARY_STORAGE attribute as follows:

```
SMTPSB.TEMPORARY_STORAGE:       /var/spool/smtpsb
```

8.4.6.4 Refresh server

You will need to refresh the Secure Mail Proxy after changing the configuration file. This step is described in 8.4.1, “Refresh” on page 214.

Chapter 9. Network Address Translation

In order to be assured of any-to-any communication across the Internet, all IP addresses have to be officially assigned by the Internet Assigned Numbers Authority (IANA). This is becoming increasingly difficult to achieve because the number of available address ranges is now severely limited. Many organizations use locally assigned IP addresses, basically used from the three blocks as described in RFC 1918 to avoid colliding with officially assigned IP addresses. These addresses, namely 10/8, 172.16/12 and 192.168/16, will not be routed on the Internet, but you can do it with NAT.

On the convention for subnets used above, 192.168/16 is the same as 192.168.0.0 mask 255.255.0.0. The number 16 refers to the number of bits ON from the leftmost bit of the mask.

NAT takes the source IP address of an outgoing packet and translates it to an official address. For incoming packets it translates the official destination address back to an internal address. The basic idea of NAT is to transparently translate the internal IP addresses of a network to official IP addresses so that they can be routed on the Internet or to hide them for privacy reasons. We now can use NAT as a routing solution for networks that have private address ranges or illegal addresses but want to communicate with hosts on the Internet.

In fact, by implementing the firewall we have already circumvented part of the problem. Clients that communicate with the Internet by using a proxy or SOCKS server do not expose their addresses to the Internet, so their addresses do not have to be translated. However, when we do not want, for whatever reason, to use a proxy or SOCKS server, or when proxy and SOCKS are not possible, we can use NAT.

The NAT solution implemented in the IBM SecureWay Firewall V4.1 for AIX makes sense if you have many hosts initiating sessions from the secure network to the Internet, but not vice versa. If you want to let users reach a server in your intranet from the Internet, you will have to use the MAP function of NAT to map the IP address of every internal server to a unique registered IP address.

You should also remember that using NAT means you have to allow routed traffic. Whenever possible, you should use proxy or SOCKS instead of routed traffic for higher security; for example, you could use the logon functions provided by the proxy to authenticate non-secure users and the Internet would only know about your non-secure interface IP address. See “Proxy” on page 87 for details.

9.1 Translation mechanism

The original so-called traditional NAT as described in RFC 1631 has been implemented in the previous versions of IBM Firewall for AIX. The new implementation is called Network Address Port Translation (NAPT). This type of NAT is also called a many-to-one translation because you need to have just one officially assigned IP address to communicate to the Internet. For more information see the following Web page:

`http://www.ietf.org/html.charters/nat-charter.html`

Traditional NAT translates the source IP address of outgoing packets to an official IP address taken out of a pool of official IP addresses. It records which official IP address was allocated to a certain internal IP address. This association can be considered a NAT session. So for incoming packets it is able to retranslate the official destination IP address to the internal IP address. The drawback to this method is that the number of concurrently active NAT sessions is limited to the number of registered IP addresses available in the pool. Once this limit has been reached, no more NAT sessions can be established until some of the existing sessions time out due to lack of activity. If NAT is configured to translate a particular new connection and it can't due to no more IP addresses in the pool, packets from the new session will be dropped.

Figure 135 shows the implementation of NAT:

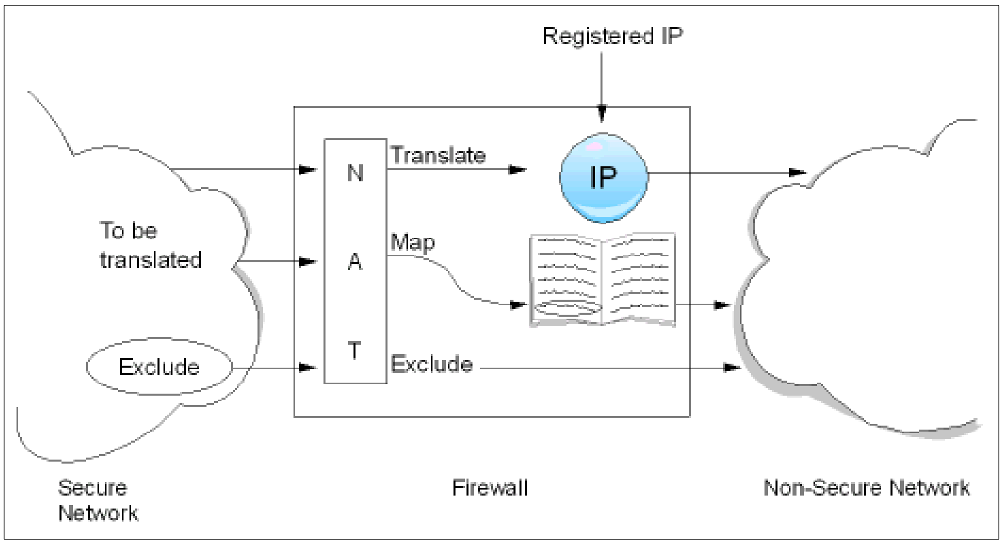

Figure 135. NAT implementation

So, how does NAPT work? For each outgoing packet the source IP address is checked by the NAT configuration rules. If the source IP address matches a translation rule, the address and port translation is performed. NAPT not only translates the internal IP address to an official IP address, but also translates the source port into a free port selected by NAPT. The firewall builds up a table with these two pieces of information to register the outgoing requests for translation, like this:

```
<src IP:src port>        -> <new src IP:new src port>
<10.1.1.2:1378>          -> <9.24.104.117:2484>
```

The destination IP address and port remain the same.

When a packet from the Internet arrives at the firewall, NAT looks up the table to find a matching record for this packet and retranslates it to the original sender's IP address and port.

```
<dst IP:dst port>        -> <original src IP:original src port> 1
<9.24.104.117:2484>      -> <10.1.1.2:1378>
```

For incoming packets, the source IP address and port remain the same.

[1] The original source IP is now the destination IP address of the incoming packet.

This translation includes adjusting of the checksum field(s) of the packet.

There are 65536 possible ports; 1024 are reserved for well-known protocols. This means that we still have more than 64000 ports left for communicating to the Internet with just one registered IP address.

This translation works transparently on most TCP, UDP and ICMP packets. But for certain FTP packets the task is even more difficult because the packets can contain IP addresses in their payload.

For example, the `FTP PORT` command contains an IP address and a port number in ASCII. These numbers are also translated correctly by the firewall NAPT including checksum updates and even TCP sequence and acknowledgment numbers. FTP is the only protocol that needs such a correction; that is handled by NAPT in the IBM SecureWay Firewall V4.1 for AIX. See 9.6, “Inside the packets” on page 236 for details.

In the above, we were talking about normal FTP. There is another way of using FTP, called passive. In the FTP passive mode, the client always initiates the communication, even for data transfer. In FTP passive mode, NAT will work fine if the client is in the secure network, but it cannot work if the client is in the non-secure network trying to reach an FTP server in the secure network.

For other protocols that have IP addresses in the payload of the packet, you should use traditional proxies or SOCKS if possible.

Regarding the packet filters, basically you need to create filter rules that will allow packets to flow from a secure network to the Internet and back (routed traffic). NAT will take care of the address translation of the secure addresses. “NAT configuration” on page 224 shows how packets flow through the firewall. You will notice that NAT translation will occur for the outgoing packet after the packet has gone through packet filtering for both interfaces (secure and non-secure adapters). This means that all packet filtering is performed using the real addresses, *not* the translated ones.

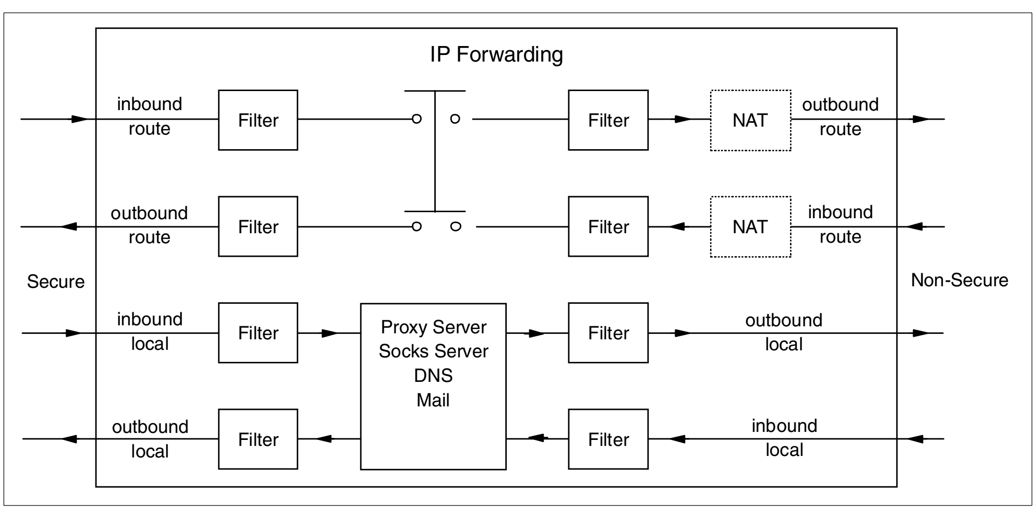

Figure 136. Flow of the packets inside the firewall

9.2 NAT configuration

The address translation is performed according to NAT rules. The NAT keywords are shown in Figure 137.

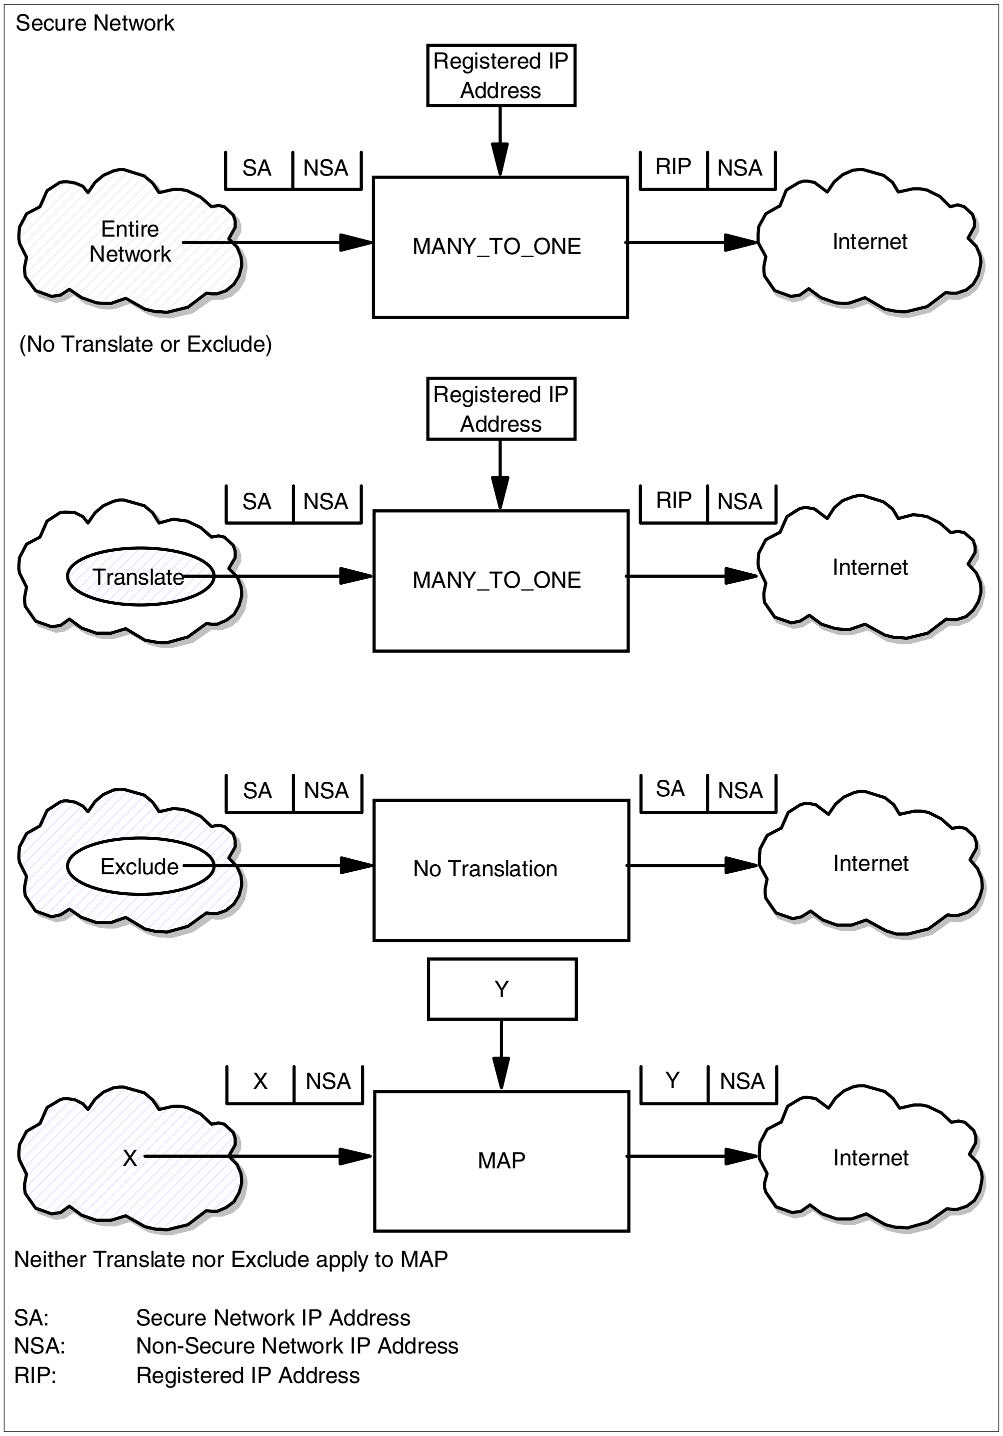

Figure 137. NAT keywords

1. Many-to-one will translate the source IP address of all outgoing routed packets, unless we use the translate or exclude NAT options.

First of all NAT has to know which registered IP address it may use for the translation. The entry consists of a registered address and a timeout value. The timeout value is the number of minutes before NAT deletes an idle address/port pair from its table. The default value is 15 minutes; the minimum value allowed is 5 minutes. An example is given in Figure 138.

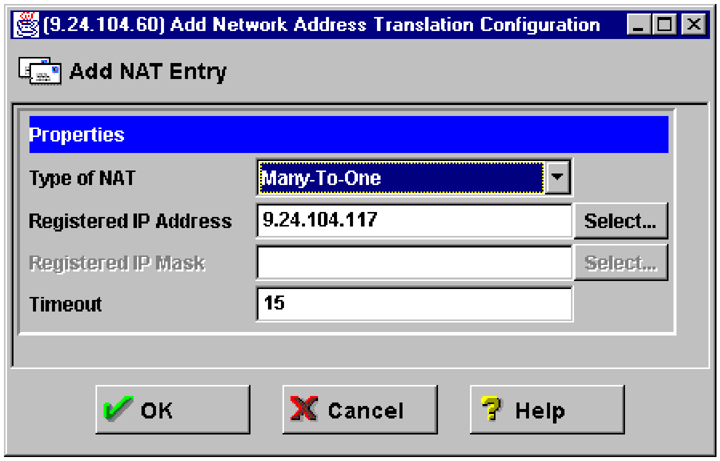

Figure 138. NAT many-to-one

Figure 138 registers the IP address 9.24.104.117 to be used for NAT.

2. Define addresses to be translated.

   By default all addresses in the secure network are translated by NAT. If you want to restrict the translation, you may specify one or more ranges of addresses that must be translated.

   For example, if you want the class A network 10.0.0.0 to be translated, you define it as shown in Figure 139.

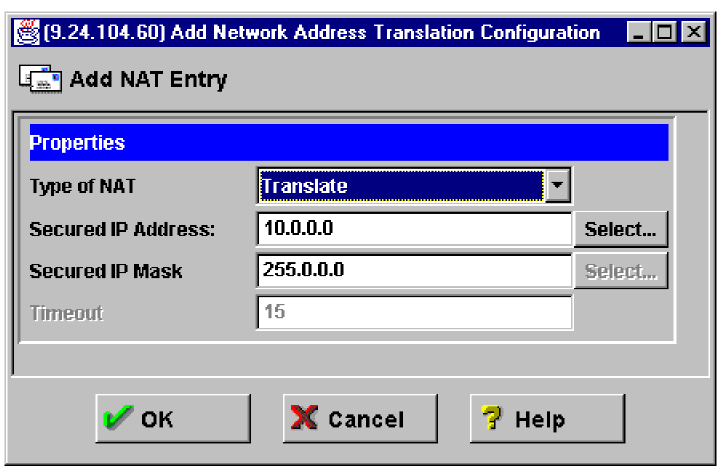

Figure 139. NAT translate

   Once the addresses to be translated by NAT have been set as shown in the preceding figure, the secure addresses not in this set will not be translated by NAT.

3. Define addresses to be excluded from translation.

Use this step if you want to exclude some addresses from the range of addresses that are allowed for translation.

For example, if you do not want to translate the host 10.1.1.3, you define it as shown in Figure 140.

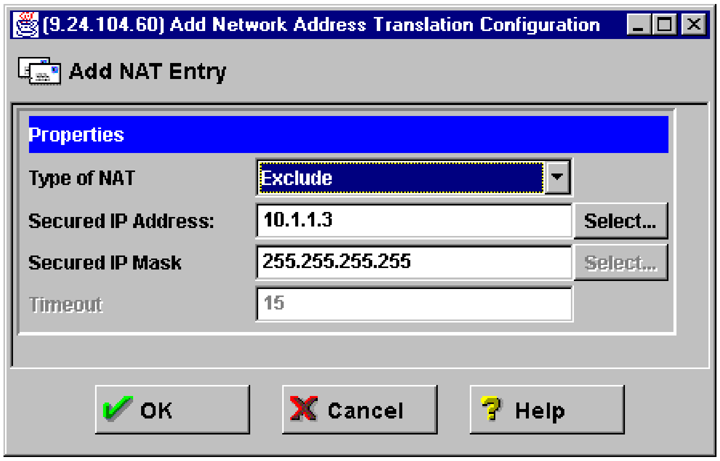

Figure 140. NAT exclude

4. Define address mappings.

   An address mapping allows you to map a secure address to a specific registered IP address. One reason for using mapping instead of many-to-one is when you want to allow users from the Internet to initiate connections to hosts in the secure network, routing through the firewall. In this case, many-to-one cannot be used.

   For example, if you want the secure address 10.1.1.2 to be translated into 9.24.104.27 then define this mapping as shown in Figure 141.

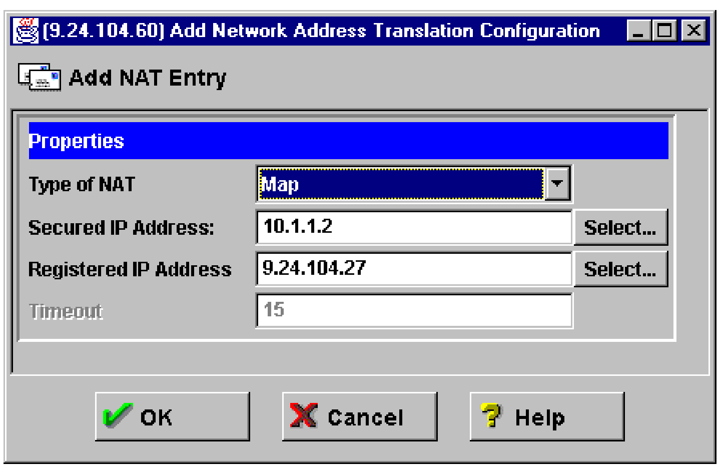

Figure 141. NAT map

A packet that arrives at the non-secure interface of the firewall with the destination IP address 9.24.104.27 will be translated into the destination IP address 10.1.1.2. Of course, a similar translation will be done with a packet that arrives at the secure side of the firewall with the source IP address 10.1.1.2.

When you add a map entry, you do not need to add a translate entry. The translate entries should be added only for many-to-one translations.

5. Activate or update the NAT configuration and logging.

   After the initial configuration and after every change you make, you have to activate/update the NAT configuration. You may also decide whether or not to activate the NAT logging facility. See Figure 142for the available options on the NAT Control Activation Status Panel.

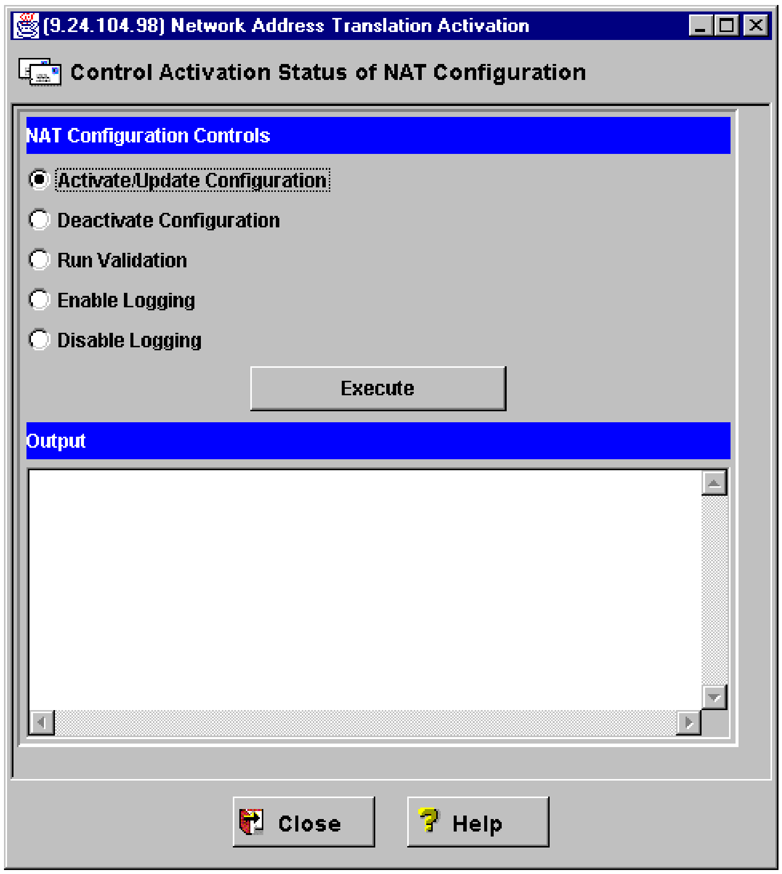

Figure 142. Control activation status window

The definitions mentioned earlier would appear in the NAT Setup window, as shown in Figure 143.

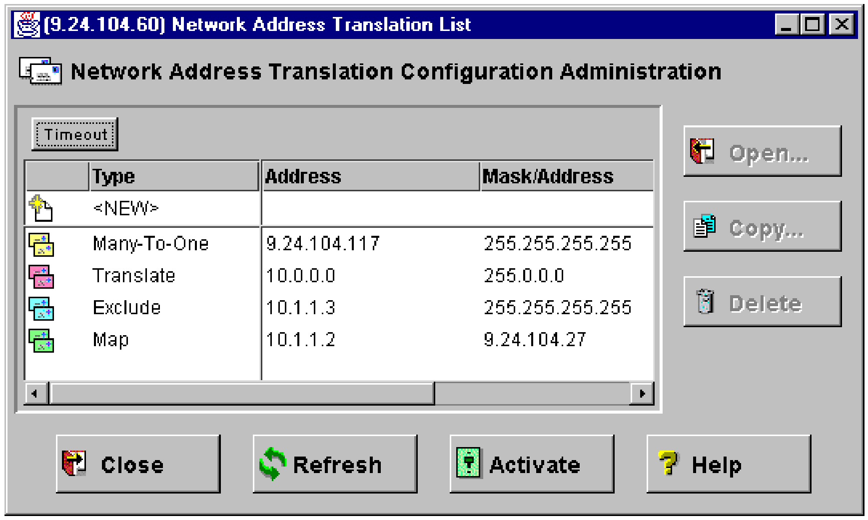

Figure 143. Example configuration

The use of NAT is independent of the filter rules, but you will have to create the filter rules that will allow the packets translated by NAT. Remember, in the filter rules you have to use your secure network addresses. The NAT address must *not* be used in the filter rules.

For an overview of the possible combinations of active NAT configuration entry and behavior, see the following table:

Table 26. Possible NAT rule combinations

Active NAT Entry	What NAT Code Does
None	When there are no active NAT configuration entries, NAT is not active and no secure addresses are translated.
Many-to-one only	All secure source addresses are translated.
Translate only	Outbound packets with source addresses matching the translate entry are discarded because there is no many-to-one entry specifying an available registered address. Outbound packets that do not match the translate entry are allowed through without translation.
Exclude only	Outbound packets with secure source addresses matching the exclude entry are allowed through without translation. Outbound packets that do not match the exclude entry are discarded.
Map only	Packets with secure source or destination addresses matching the map entry are translated. Outbound packets that do not match the map entry are allowed through without translation.
Many-to-one and translate	Outbound packets with source secure addresses matching the translate entry are translated because there is many-to-one entry specifying an available registered address. Outbound packets that do not match the translate entry are allowed through without translation.
Many-to-one and exclude	Outbound packets with secure source addresses matching the exclude entry are allowed through without translation. Outbound packets that do not match the exclude entry are translated.

Active NAT Entry	What NAT Code Does
Many-to-one and map	Packets with secure source or destination addresses matching the map entry are translated. Outbound packets that do not match the map entry are translated with the many-to-one address.
Exclude and translate	Outbound packets with secure source addresses matching the exclude entry are allowed through without translation. Outbound packets with secure source addresses matching the translate entry are discarded because there is no many-to-one entry specifying an available registered address. Outbound packets that do not match either entry are allowed through without translation even though they do not match the exclude entry.
Exclude and map	Outbound packets with secure source addresses matching the exclude entry are allowed through without translation. Packets with secure source or destination matching the map entry are translated. Outbound packets that do not match either entry are allowed through without translation.
Map and translate	Packets with secure source or destination matching the map entry are translated. Outbound packets with secure source addresses matching the translate entry are discarded because there is no many-to-one entry specifying available external addresses. Outbound packets that do not match either entry are allowed through without translation.
Many-to-one, translate and exclude	Outbound packets with secure source addresses matching the translate entry are translated because there is a many-to-one entry specifying an available registered address. Outbound packets with secure source addresses matching the exclude entry are allowed through without translation. Outbound packets that do not match either entry are allowed through without translation even though they do not match the exclude entry.
Many-to-one, translate and map	Outbound packets with secure source addresses matching the translate entry are translated because there is a many-to-one entry specifying an available registered address. Packets with secure source or destination addresses matching the map entry are translated. Outbound packets that do not match either entry are allowed through without translation.
Exclude, translate and map	Outbound packets with secure source addresses matching the exclude entry are allowed through without translation. Outbound packets with secure source addresses matching the translate entry are discarded because there is no many-to-one entry specifying an available registered address. Packets with secure source or destination addresses matching the map entry are translated. Outbound packets that do not match any entry are allowed through without translation.
Many-to-one, exclude and map	Outbound packets with secure source addresses matching the exclude entry are allowed through without translation. Packets with secure source or destination addresses matching the map are translated. Outbound packets that do not match either entry are translated.

Active NAT Entry	What NAT Code Does
Many-to-one, translate, exclude and map	Outbound packets with secure source addresses matching the translate entry are translated because there is a many-to-one entry specifying an available registered address. Outbound packets with secure source addresses matching the exclude entry are allowed through without translation. Packets with secure source or destination addresses matching the map entry are translated. Outbound packets that do not match any entry are allowed through without translation.

9.3 How to configure routing when using NAT

When we use NAT to translate internal IP addresses into registered IP addresses, we need to associate these addresses with the non-secure adapter of the firewall.

Since these registered IP addresses do not exist physically, it is necessary that somehow the packets are actually sent to the firewall.

We can consider two possible network layouts: first, your registered IP address is in the same subnet as the non-secure adapter of the firewall; second, your registered IP address is in a different subnet. Depending on the layout you have, you must perform specific steps to guarantee that all packets sent to registered addresses will be routed to the non-secure adapter of the firewall.

We will consider two situations:

- The registered NAT IP address is in the same subnet as the non-secure interface
- The registered NAT IP address is in a different subnet from the non-secure interface

After discussing each situation, we will discuss MTU sizing, which is also an important issue when working with NAT.

9.3.1 The registered NAT IP address is in the same subnet

If your NAT IP addresses are in the same subnet as your non-secure adapter IP address, you need to get your firewall host to respond to ARP requests for the NAT IP addresses with the MAC address of the non-secure adapter. Basically you will add a permanent entry for this IP address in the ARP table of the firewall, pointing to its own non-secure adapter.

Follow these steps:

1. Determine the MAC address of the non-secure adapter of the firewall.

   You can use the command `netstat` to get this information:

```
# netstat -in

Name  Mtu   Network     Address            Ipkts Ierrs    Opkts
lo0   16896 link#1                         2393863     0  2393863
lo0   16896 127         127.0.0.1          2393863     0  2393863
lo0   16896 ::1                            2393863     0  2393863
tr0   1492  link#2      0.4.ac.63.31.b9    2616320     0    12509
tr0   1492  9.24.104    9.24.104.60        2616320     0    12509
tr1   1492  link#3      0.20.35.44.e1.b8   2602066     0      130
tr1   1492  10.1.1      10.1.1.1           2602066     0      130
```

The command `netstat -in` shows you the IP address for each interface. It also shows you the MAC address of the respective adapter. In this example, our non-secure interface is tr0, with IP address 9.24.104.60, and the MAC address of the adapter is 0.4.ac.63.31.b9.

2. Add an entry in the ARP table for each NAT IP address.

   Use the `arp` command with the following syntax:

   ```
   arp -s <adapter type> <NAT IP address> <nonsecure adapter MAC address> pub
   ```

   In our example, we are going to add an entry for the NAT IP address 9.24.104.117. We are using token-ring, so the `arp` command will look like:

   ```
   # arp -s 802.5 9.24.104.117 00:04:ac:63:31:b9 pub
   ```

   We will review each parameter used in this command.

   The first parameter right after `arp -s` stands for the type of adapter you are using. In this case, 802.5 stands for token-ring. The supported keywords for this parameter are listed in Table 27:

Table 27. Types of adapters available in arp command

Type of adapter	Keyword
Ethernet	ether
IEEE 802.3	802.3
FDDI	fddi
Token-Ring	802.5

The second parameter is the NAT IP address (the registered IP address that belongs to your non-secure network).

The third parameter is the physical address of your non-secure adapter. Note that the syntax of the MAC address is not the same as shown in the output of `netstat -in`. In the `arp` command, we have to use colons (:) instead of periods (.). This means that we have to write the address `0.4.ac.63.31.b9` as `0:4:ac:63:31:b9` (or `00:04:ac:63:31:b9`).

The last parameter is the word "pub", which means that the new IP address will be *published*: ARP requests for it will be responded to by this host. This new ARP entry will not be removed from the ARP table; it is a permanent entry. The only way to remove it is by running the `arp` command with the option `-d` or rebooting the machine. Now, you can list your ARP table and check this permanent entry:

```
# arp -an
? (9.24.104.123) at 10:0:5a:b1:b5:1a [token ring]
? (9.24.104.1) at 40:0:22:16:aa:0 [token ring]
? (9.24.104.108) at 10:0:5a:b1:d7:31 [token ring]
? (9.24.104.117) at 0:4:ac:63:31:b9 [token ring] permanent published
```

3. Add the `arp` command in /etc/rc.net

   The ARP table is cleaned when the machine is rebooted (including the permanent entries), so we need to include the arp command(s) in the /etc/rc.net file. Now, when the machine reboots, the ARP table will be automatically updated.

   We added the following lines to the end of the /etc/rc.net file:

   ```
   # FW - update to ARP table - used by NAT
   /usr/sbin/arp -s 802.5 9.24.104.117 00:04:ac:63:31:b9 pub
   ```

   Remember that you have to run the `arp` command for each NAT address you are using, and you must also add all of them to /etc/rc.net.

   In case you need to remove the NAT configuration for this address, you also need to remove its entry from the ARP table and from the /etc/rc.net file.

9.3.2 The registered NAT IP address is in a separate subnet

If your NAT registered IP addresses are in a separate subnet from your non-secure adapter IP address, you must add static routes on the non-secure router (the one that connects your non-secure network to your ISP network).

The destination address of these static routes are the NAT addresses (or range), and the gateway is the firewall non-secure address.

This will work as if this range of NAT addresses were part of a virtual LAN behind the firewall.

The external router needs to advertise this route through the Internet to the backbone routers. Remember to configure your router to advertise its static routes.

9.3.3 Routing inside the secure network

When we are using NAT, we allow a connection between machines in our secure network and the Internet. To be able to send the packets to the Internet, our internal machines must have a route to go out of our network into the Internet. In other words, their default gateway must be the firewall box.

If your secure network has one or more routers (to communicate with subnets, remote locations, and so forth), you should first identify your main router. Then add a static route on this router to use the firewall as its default gateway.

If you have only one router, add the firewall as its default gateway. If you have only one subnet and have no routers inside your secure network, then your clients and servers will use the firewall itself as default gateway.

Also remember to add routes for your remote internal networks on the firewall, since we recommend that you do not use dynamic routing on the firewall box.

9.3.4 NAT and ICMP

An improvement in NAT in IBM SecureWay Firewall V4.1 for AIX is that ICMP Error packets are now fully translated. This means that if you choose to allow ICMP Error packets to flow through your firewall filters, NAT will no longer prevent them from being delivered to the original packet sender, regardless of whether that sender is on the secure or non-secure side of your firewall. Mechanisms that relay on ICMP error packets, such as Path MTU Discovery, will now work through NAT.

If you decide not to allow ICMP packets through your firewall, see 4.5, “ICMP traffic and MTUs” on page 73 for how to set MTU in order to avoid ICMP traffic.

9.4 Timeout value

The default timeout value for NAT dynamically translated connections (many-to-one) is 15 minutes.

For example, if your telnet connection through NAT is idle for more than 15 minutes, your NAT address/port will be released from the table. Effectively it means that your telnet connection is lost. NAT will log these timeout IP address releases to the syslog; the log message number is ICA9047.

The timeout value does not apply to map connections.

You may increase the timeout value if you expect the clients to have idle periods longer than 15 minutes. Bear in mind that even when the connection is closed, the NAT address/port pair will remain allocated for an amount of time equivalent to the timeout value. If you run out of address/port pairs, an ICA9046 message will be logged.

9.5 Example configurations for using NAT

Now we will present some useful scenarios with NAT.

9.5.1 Using NAT to the Internet

In this example, we use a setup as shown in Figure 144:

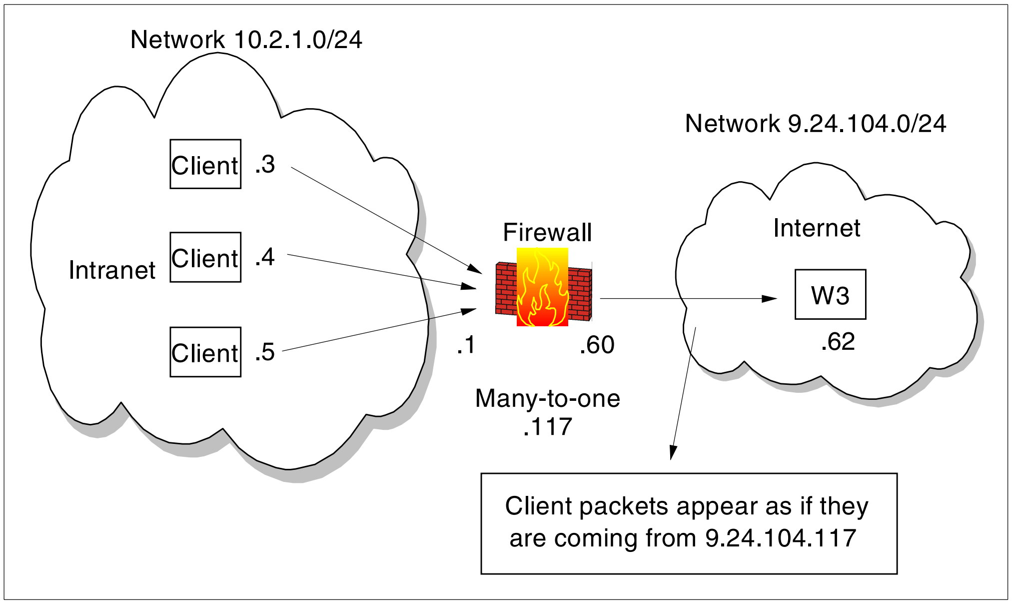

Figure 144. Basic NAT configuration

For this scenario IP forwarding has to be activated in the firewall. The client's default gateway is the secure interface of the firewall, since we do not have a router in our test configuration (in the secure network).

In your case, if you have a router in your secure network (or more than one), you do not need to change your client's routing table. You just have to add a default route on your main internal router, so if any client tries to reach an IP address outside your network, this router will direct this packet to the firewall.

There are no proxies or SOCKS configured for the Web browsers on the client. We added the many-to-one registered IP address 9.24.104.117 in the NAT configuration, as seen in Figure 144. Remember to activate this configuration in the NAT Activation Panel.

From now o, all secure IP addresses will be translated into the address 9.24.104.117.

After that, we added this address to the ARP table of the firewall, using the following command:

```
arp -s 802.5 9.24.104.117 00:04:ac:63:31:b9 pub
```

In the next step, we created filter rules that allow direct outgoing HTTP traffic from the secure network to the world. Predefined services for "HTTP direct out" are shipped with the firewall. You only need to create a connection and use it.

The filter rules for this connection are:

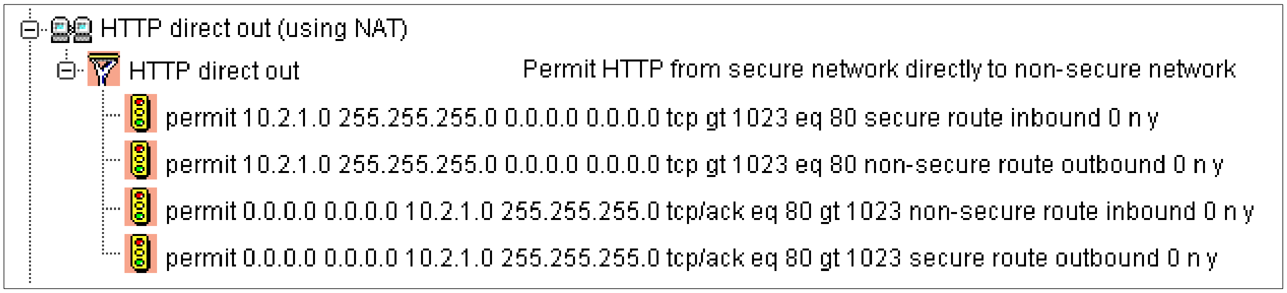

Figure 145. Filter rules for routed outgoing HTTP traffic

Notice that we did not specify our registered IP address 9.24.104.117 in our filter rule.

When we tried to access a non-secure Web server, it worked fine. The access log of the Web server showed only connections from 9.24.104.117, which means that the real IP address of the client is not seen outside.

9.5.2 Mapping a server

In this example we are using a setup like the one shown in Figure 146. We want to allow users in the Internet to access your internal Web server that resides in your secure network.

Since the HTTP sessions will start from the non-secure network and not from the secure network, we cannot use many-to-one translation; we need to use a static mapping.

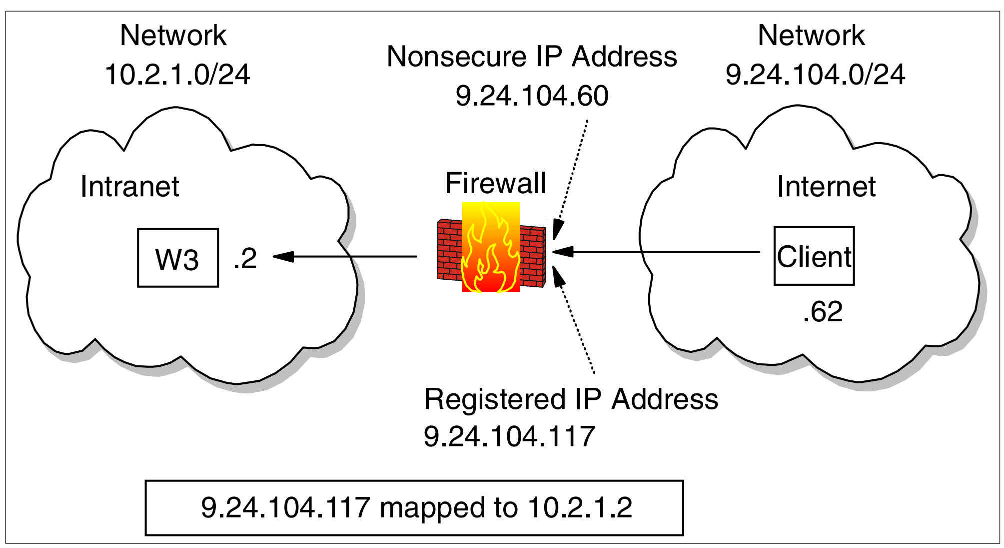

Figure 146. Mapping IP addresses with NAT

Even if you plan to set up a DMZ it might be a good idea to use NAT anyhow, since you protect your servers a little more by hiding their real addresses, and you can use IP filtering to allow only the necessary connections from outside. [2]

[2] We are using the same registered address, 9.24.104.117, as in the previous sample, where we used many-to-one. This is due to a limitation in the number of IP addresses we were able to use; you should not use the same registered IP address for many-to-one and for mapping.

We copied the "HTTP direct out" templates and created a "HTTP direct in" where we just exchanged the secure and non-secure interfaces on the four rules. With this we defined a new connection from "The World" to the secure IP address of the Web server. Remember not to specify any NAT IP addresses in the rules.

Checking the filter rules we see the following:

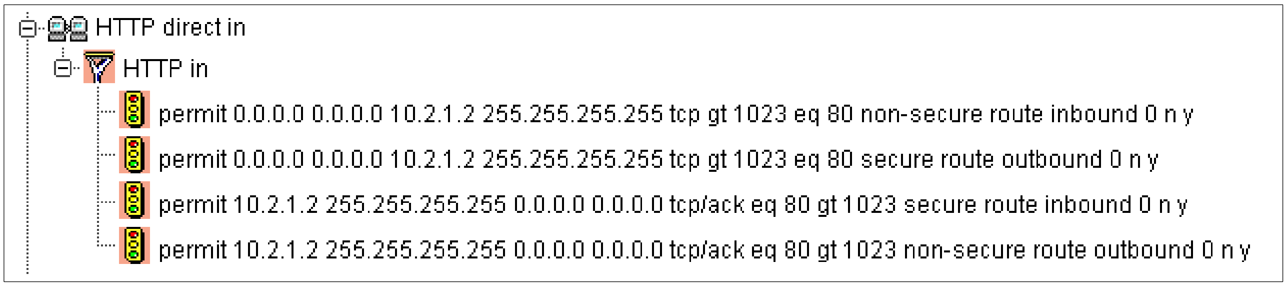

Figure 147. Filter rules for routed outgoing HTTP traffic

We also need to add the ARP table entry for this address. The command used is:

```
arp -s 802.5 9.24.104.117 00:04:ac:63:31:b9 pub
```

The last step is to activate the new NAT configuration and the new rules. Then we used the client's Web browser, which does not have any proxy or SOCKS configured. Accessing the IP address 9.24.104.117 worked as expected.

9.6 Inside the packets

In this section we show some traces to illustrate the translation of addresses done by NAT.

We begin with a simple example. Using the setup shown in Figure 146, we added filters to allow the traffic of ICMP packets also. That means that now it is possible to ping any machine on the non-secure network using NAT.

Now, we start a ping from the machine 10.2.1.3 in the secure network (not shown in the figure) to the machine 9.24.104.62, which is our external DNS server.

When the "echo request" is sent, we see the following packets in an iptrace:

```
Packet Number 1
TOK: ====( 82 bytes received on interface tr0 )==== 17:43:54.968982394
TOK: 802.5 packet
TOK: 802.5 MAC header:
TOK: access control field = 10, frame control field = 40
TOK: [ src = 00:06:29:f3:65:78, dst = 00:20:35:44:e1:b8]
TOK: 802.2 LLC header:
TOK: dsap aa, ssap aa, ctrl 3, proto 0:0:0, type 800 (IP)
IP:     < SRC =        10.2.1.3 >
IP:     < DST =     9.24.104.62 >
IP:     ip_v=4, ip_hl=20, ip_tos=0, ip_len=60, ip_id=45363, ip_off=0
IP:     ip_ttl=32, ip_sum=6d33, ip_p = 1 (ICMP)
ICMP:   icmp_type=8 (ECHO_REQUEST)  icmp_id=512  icmp_seq=20224

Packet Number 2
TOK: ====( 82 bytes transmitted on interface tr1 )==== 17:43:54.969000025
TOK: 802.5 packet
TOK: 802.5 MAC header:
TOK: access control field = 0, frame control field = 40
TOK: [ src = 00:04:ac:63:31:b9, dst = 40:00:52:00:51:94]
TOK: 802.2 LLC header:
TOK: dsap aa, ssap aa, ctrl 3, proto 0:0:0, type 800 (IP)
IP:     < SRC =    9.24.104.117 >
IP:     < DST =     9.24.104.62 >
IP:     ip_v=4, ip_hl=20, ip_tos=0, ip_len=60, ip_id=45363, ip_off=0
IP:     ip_ttl=31, ip_sum=7ab, ip_p = 1 (ICMP)
ICMP:   icmp_type=8 (ECHO_REQUEST)  icmp_id=512  icmp_seq=20224
```

In this firewall, tr0 is the secure interface and tr1 is the non-secure interface.

In the preceding iptrace, you can see that packet number 1 arrives at the secure interface (see the first line of this packet), coming from the machine 10.2.1.3 (see the field SRC in the first line of the IP header) and its destination is the machine 9.24.104.62 (see the field DST in the second line of the IP header).

The packet number 2 is going out to the non-secure network (it is being transmitted by the non-secure interface), and note that the SRC field on the IP header has changed. Now it is 9.24.104.117, which is the NAT address.

So if you look at the packet number 2, which is the one that is sent to the Internet, it does not have any information about the internal IP address of this machine.

See the field icmp_id; because we used mapping, this field is not changed. If we would have used many-to-one, it would have been changed by NAT.

See also field ip_sum; it represents the checksum. NAT must update it to reflect the changes it has done on some fields of the packet.

Now, let’s take a look at the response from the machine 9.24.104.62:

```
Packet Number 3
TOK: ====( 82 bytes received on interface tr1 )==== 17:43:54.970365754
TOK: 802.5 packet
TOK: 802.5 MAC header:
TOK: access control field = 18, frame control field = 40
TOK: [ src = 40:00:52:00:51:94, dst = 00:04:ac:63:31:b9]
TOK: 802.2 LLC header:
TOK: dsap aa, ssap aa, ctrl 3, proto 0:0:0, type 800 (IP)
IP:     < SRC =     9.24.104.62 >
IP:     < DST =    9.24.104.117 >
IP:     ip_v=4, ip_hl=20, ip_tos=0, ip_len=60, ip_id=20040, ip_off=0
IP:     ip_ttl=128, ip_sum=996, ip_p = 1 (ICMP)
ICMP:   icmp_type=0 (ECHO_REPLY)  icmp_id=512  icmp_seq=20224

Packet Number 4
TOK: ====( 82 bytes transmitted on interface tr0 )==== 17:43:54.970412149
TOK: 802.5 packet
TOK: 802.5 MAC header:
TOK: access control field = 0, frame control field = 40
TOK: [ src = 00:20:35:44:e1:b8, dst = 00:06:29:f3:65:78]
TOK: 802.2 LLC header:
TOK: dsap aa, ssap aa, ctrl 3, proto 0:0:0, type 800 (IP)
IP:     < SRC =     9.24.104.62 >
IP:     < DST =        10.2.1.3 >
IP:     ip_v=4, ip_hl=20, ip_tos=0, ip_len=60, ip_id=20040, ip_off=0
IP:     ip_ttl=127, ip_sum=711e, ip_p = 1 (ICMP)
ICMP:   icmp_type=0 (ECHO_REPLY)  icmp_id=512  icmp_seq=20224
```

We can see these packets are the ping response ("echo reply") sent by the machine 9.24.104.62.

In packet number 3, we can see the SRC field in IP header is 9.24.104.62, and the DST (destination) field is 9.24.104.117, which is the NAT address. This packet was received by the non-secure interface (tr1).

When the firewall sends this packet to the secure network (see packet number 4), note that it automatically changes the field DST of the IP header, which is the real address of the internal machine (10.2.1.3).

In the preceding example, the only fields that the firewall changed were in the header of the packet. Now, we will see an example where it also changes the information inside the body of the packet.

We used the same environment for this next test, but we added more filters to allow FTP from the client machine 10.2.1.3 to the server 9.24.104.62.

Now, we open an FTP connection from 10.2.1.3 to 9.24.104.62, using NAT. See the following iptrace:

```
Packet Number 39
TOK: ====( 83 bytes received on interface tr0 )==== 17:44:18.139690664
TOK: 802.5 packet
TOK: 802.5 MAC header:
TOK: access control field = 10, frame control field = 40
TOK: [ src = 00:06:29:f3:65:78, dst = 00:20:35:44:e1:b8]
TOK: 802.2 LLC header:
TOK: dsap aa, ssap aa, ctrl 3, proto 0:0:0, type 800 (IP)
IP:     < SRC =       10.2.1.3 >
IP:     < DST =    9.24.104.62 >
IP:     ip_v=4, ip_hl=20, ip_tos=21, ip_len=61, ip_id=50739, ip_off=0DF
IP:     ip_ttl=128, ip_sum=b817, ip_p = 6 (TCP)
TCP:    <source port=1275, destination port=21(ftp) >
TCP:    th_seq=1ada4e4, th_ack=487396aa
TCP:    th_off=5, flags<PUSH | ACK>
TCP:    th_win=5691, th_sum=656e, th_urp=0
TCP: 00000000     504f5254 2031302c 322c312c 332c342c     |PORT 10,2,1,3,4,|
TCP: 00000010     3235330d 0a                             |253..           |
```

This packet was sent right after issuing the `dir` command (it arrived at the secure interface of the firewall). The client (10.2.1.3) sends a `PORT` command to the server containing its own IP address (it uses commas instead of periods). If this data is not translated also, this connection may fail. Now, let’s see the next packet:

```
Packet Number 40
TOK: ====( 87 bytes transmitted on interface tr1 )==== 17:44:18.139712327
TOK: 802.5 packet
TOK: 802.5 MAC header:
TOK: access control field = 0, frame control field = 40
TOK: [ src = 00:04:ac:63:31:b9, dst = 40:00:52:00:51:94]
TOK: 802.2 LLC header:
TOK: dsap aa, ssap aa, ctrl 3, proto 0:0:0, type 800 (IP)
IP:     < SRC =   9.24.104.117 >
IP:     < DST =    9.24.104.62 >
IP:     ip_v=4, ip_hl=20, ip_tos=21, ip_len=65, ip_id=50739, ip_off=0DF
IP:     ip_ttl=127, ip_sum=528b, ip_p = 6 (TCP)
TCP:    <source port=1275, destination port=21(ftp) >
TCP:    th_seq=1ada4e4, th_ack=487396aa
TCP:    th_off=5, flags<PUSH | ACK>
TCP:    th_win=5691, th_sum=9772, th_urp=0
TCP: 00000000     504f5254 20392c32 342c3130 342c3131 |PORT 9,24,104,11|
TCP: 00000010     372c342c 3235330d 0a                |7,4,253..       |
```

This packet was sent by the firewall non-secure interface (tr1) to the Internet. Note that the SRC field of the IP header was changed; it is now 9.24.104.117. And it also changes the parameters of the `PORT` command inside the data of the packet. It is now showing 9,24,104,117.

The source port number (4,253) was not changed because this translation was done with a static map.

9.7 NAT and Virtual Private Networks (VPNs)

One basic thing to remember with NAT is that the registered IP addresses will not be used in your local filter rules and there is no exception when using NAT together with VPN. Figure 148 shows where NAT and VPN take place in IBM SecureWay Firewall V4.1 for AIX.

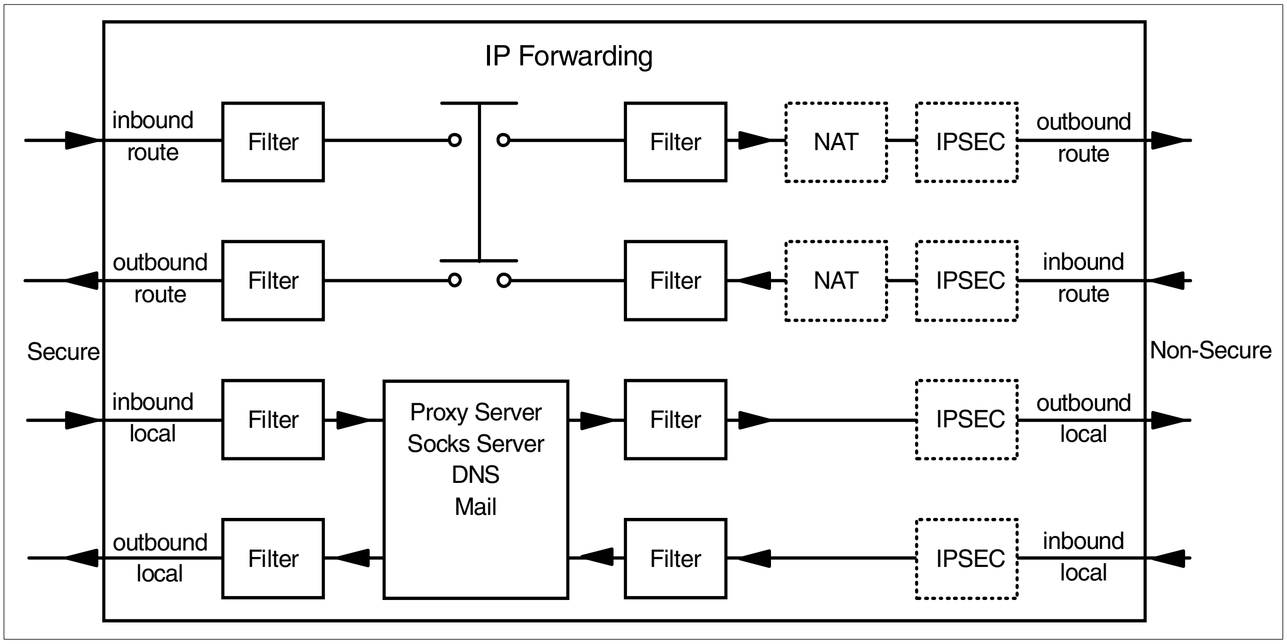

Figure 148. Where NAT and IPSEC functions are performed

But your registered (translated) IP address will be visible in the Internet and, of course, in the other firewall where the tunnel terminates.

Remember that NAT registered IP addresses of a given firewall should never appear in its own filter rules.

See 10.5.1.1, “Using NAT with IPSec” on page 273 for more details.

9.8 NAT and multiple adapters

The NAT implementation in IBM SecureWay Firewall V4.1 for AIX supports multiple secure and non-secure adapters. NAT is able to route the translated packets to the proper interface.

While planning NAT for a firewall with multiple adapters you should keep in mind that the NAT configuration is on a firewall basis, not on an individual adapter basis. You cannot have different NAT configurations for different adapters.

Chapter 10. Virtual Private Network

In this chapter we discuss secure IP tunnels. They are a mechanism provided by IBM SecureWay Firewall V4.1 for AIX to allow secure communications across a non-secure intervening network like the Internet. It constructs a virtual private network (VPN) between two different sites providing authentication and encryption. A redbook with more detailed information about virtual private networks is A *Comprehensive Guide to Virtual Private Networks, Volume I,* SG24-5201. Volume II of this redbook, *Comprehensive Guide to Virtual Private Networks, Vol II,* SG24-5234*,* is also recommended reading.

10.1 Secure IP tunnel standards - inter operability

Secure IP tunnel products have existed in the market for years. However, due to the lack of an IP security standard, these products were proprietary in nature (that is, they can establish secure tunnels only with their own product). IBM and other organizations were actively involved in the development of standards that would allow firewalls from different manufacturers to establish tunnels between them. The basis for these standards is a group of RFCs of the Internet Engineering Task Force (IETF) IP Security Protocol (IPSec) working group. The charter for the IPSec group, plus links to the IPSec RFCs, can be found at the following URL:

`http://www.ietf.org/html.charters/ipsec-charter.html`

Basically, IPSec is a network layer security protocol that will provide authentication, integrity checking, and encryption to IP datagrams. IPSec is the IETF-chosen security framework for both IPv4 and IPv6 environments, and is recommended as the security for PPTP, L2F and L2TP tunnels. IPSec defines two mechanisms to achieve these security objectives. They are the IP Authentication Header (AH) and the IP Encapsulating Security Payload (ESP). The details of AH and ESP can be found in RFC 2402 and RFC 2406 respectively, and the overall architecture of IPSec is described in RFC 2401. The supported data integrity algorithms are Hashed Message Authentication Code (HMAC) using Message Digest 5 (MD5) or Secure Hash Algorithm (SHA); these are described in RFC 2403 and RFC 2404. The supported encryption algorithm Data Encryption Standard (DES) is described in RFC 2405. Triple DES is described in RFC 2451.

IBM SecureWay Firewall V4.1 for AIX provides two kinds of IPSec tunnels to cater to different situations.

We can establish manual tunnels with dynamic or static filter rules, both using the IPSec standard. This has to be done manually by exporting and importing the tunnel configurations. They do not support Key Exchange; this means that the keys are set at tunnel definition and do not change throughout the life of the tunnel.

To establish an IPSec tunnel between the IBM SecureWay Firewall V4.1 for AIX and another host, the partner node must support the new IPSec headers, described in the RFCs mentioned above. IPSec tunnels can be established between an IBM SecureWay Firewall V4.1 for AIX and:

- IBM SecureWay Firewall V4.1 for AIX

- An IBM eNetwork Firewall for Windows NT V3.3
- OS/390 Firewall Technologies V2R7 or higher
- AIX 4.3 operating system (supports the old and new IPSec headers)
- AS/400 operating system V4R4 (supports new IPSec headers)
- 2210/2212/2216 routers running MRS/AIS/MAS software V3.1 or higher
- Other vendors' software that conform to the RFCs mentioned above

Incompatibility

The IPSec implementation of IBM SecureWay Firewall V4.1 for AIX is not compatible with previous versions of Firewall for AIX and the Windows95 IPSec client shipped with these versions.

The current IPSec standards are described in RFCs 2401, 2402, 2403, 2404 and 2405; the old ones are described in RFC 1825, 1826 and 1827.

10.2 Operation of the secure tunnel

The IPSec tunnel relies on symmetric-key cryptography to enforce data security. The secure tunnels established with the Firewall have shared secrets, i.e., keys for authentication and/or encryption that are known to both ends of the tunnel before data is passed through the tunnel. The IPSec tunnel provides two different types of security:

1. Authentication, in which the sending firewall appends a message authentication code (MAC) to the messages it sends through the tunnel. The MAC is constructed from the message contents and the authentication key using a one-way hash function. The receiving firewall performs the same operation and, if the MAC matches, it knows that the message is authentic and has not been altered while being transported over the non-secure network.
2. Encryption, in which the data within the message is encrypted using the secure key, so that it cannot be viewed in transit.

Authentication and encryption can be used independently. In fact, you can enable or disable the two features for each tunnel. A typical scenario will have multiple secure networks (for example, branches of a company that are in different cities) with tunnels between them in order to protect the information. There may be more than one tunnel between a single pair of nodes, which might be useful for different encryption and authentication choices.

For example, your computer department may wish to monitor machines in the finance department using SNMP. In this case, the information itself is not sensitive, but you want to be sure that it is accurate, so you could use a tunnel that provides authentication only.

However, you also want the computer department to send mail to the finance department and you would like to protect this mail from being read in the non-secure network. This would require a second tunnel providing both authentication and encryption.

When a packet has to go from one secure network to another secure network through the IPSec tunnel, the whole IP packet will be encrypted and authentication data will be created at the first end and sent in a new IP packet to the second end of the tunnel. Note that the packet is not sent using the normal IP protocols (TCP or UDP), but using a special security protocol (AH or ESP). In Figure 149we show a black border around the original IP packet to show that it is being protected in the non-secure network by the IPSec tunnel.

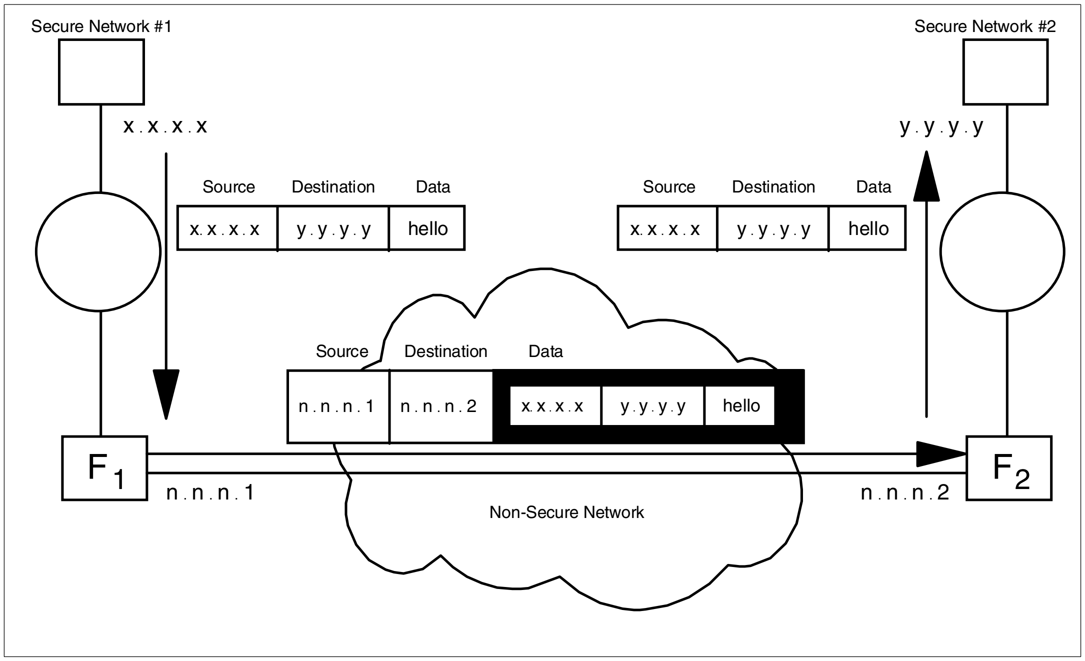

Figure 149. Operation of the IPSec tunnel

When a packet leaves a host in secure network #1 the source IP address is x.x.x.x with destination IP address y.y.y.y. In the firewall F1 the packet is encapsulated and the new IP addresses are those of the non-secure adapters of both firewalls. When the packets reaches the firewall F2, the original packet is restored.

10.3 Implementing the IPSec tunnel

In order to configure a tunnel with dynamic or static filters with the IBM SecureWay Firewall V4.1 for AIX you will have to follow these steps:

1. Add the tunnel definition in one node (tunnel owner).
2. Export the tunnel definition to a file.
3. Transfer the tunnel definition file to the partner node.
4. Import the tunnel definition in the partner node.
5. Activate the tunnel at both ends.
6. Specify which protocols you want to tunnel using filtering rules and activate the rules (for tunnels with static filter rules only).
7. Reactivate the tunnel when the lifetime has expired.

You also have to consider that you need the following prerequisites:

1. IP forwarding enabled in both firewalls for routed traffic.

2. Coherent IP addresses in both secure networks (for example, you cannot use the same private IP addresses). If this is not the case, we will have to use Network Address Translation (NAT); see “NAT and Virtual Private Networks (VPNs)” on page 239 for details.
3. Proper routes in the clients (they point to the firewall for addresses in the other secure network).
4. Name resolution for the remote networks (this is important if you want to pass hidden DNS information through the tunnel).

We will describe each implementation step in turn.

10.3.1 Adding the tunnel definition in one node

The tunnel has to be created manually; we can create a tunnel with dynamic or static filter rules. Tunnels with dynamic filters are easy to configure; they have fixed filter rules that are generated automatically when the tunnel is used. Tunnels with static filter rules have user-defined filter rules and provide fine-grained control.

10.3.1.1 When to use each type of tunnels

Based on the descriptions above, you may want to use an IPSec tunnel with dynamic filter rules in the following cases:

- Trusted location; we have to remember that once the tunnel is established, the remote hosts have almost the same access to our secure network resources as any host in our network, but limited to the end points of the tunnel.
- Testing purposes, in order to avoid any possible filter rules errors.

You may want to use an IPSec tunnel with static filter rules in the following cases:

- You want to allow only specific protocols to flow through the tunnel.
- You have more than one tunnel between two locations and you want to distribute the traffic on each tunnel based on protocols.
- You have more than one tunnel between two locations and you want to send encrypted traffic through one tunnel and authenticated traffic through the other.
- There is a Network Address Translation (NAT) at one or both ends of the tunnel.
- You want to limit the use of the tunnel to certain hours of the day.

We will now configure a tunnel with dynamic rules, using the configuration shown in Figure 174 on page 267. A tunnel with static rules will be configured in 10.5, “Virtual Private Network scenarios” on page 266.

When we select **Virtual Private Network** from the main menu, the window shown in Figure 150 is displayed. Using this window we can open, copy, delete, import, export, activate, and deactivate tunnels.

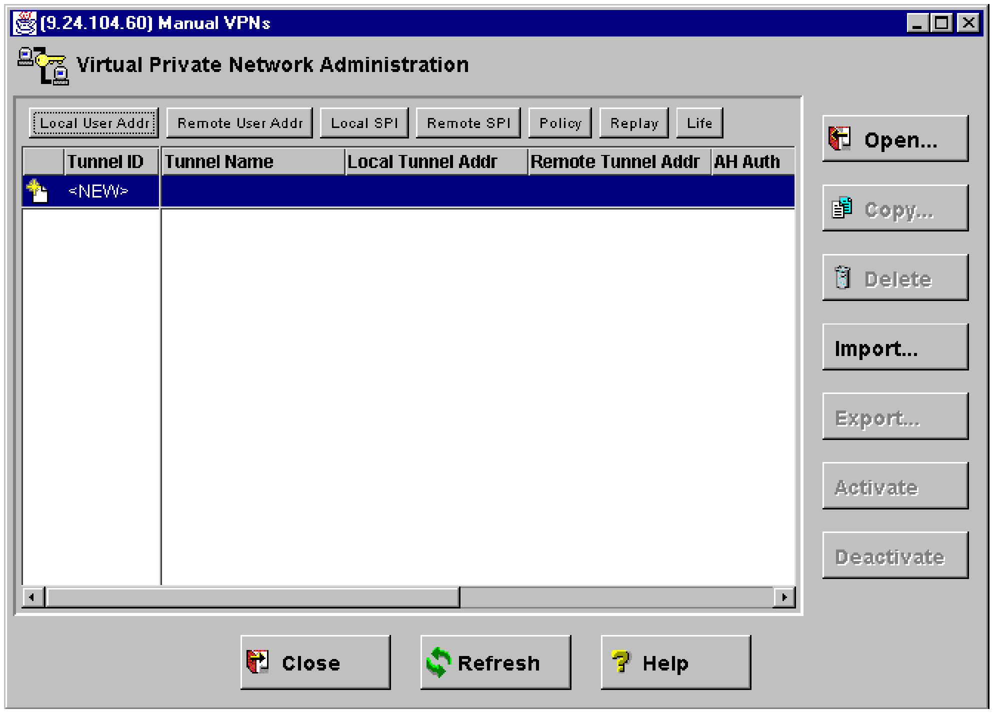

Figure 150. The Virtual Private Network administration window

When we want to create a tunnel for the first time we can use only Open to create a new tunnel, and Import if we want to import the tunnel definitions from another tunnel owner.

When we create the tunnel we must define, using the Configuration Client GUI, the characteristics of the tunnel: the tunnel type, the addresses of both ends of the tunnel, the authentication/encryption desired, and the parameters for the session key. When creating a tunnel with dynamic rules, we also have to define the local and remote addresses. Figure 151 shows the Configuration Client GUI screen.

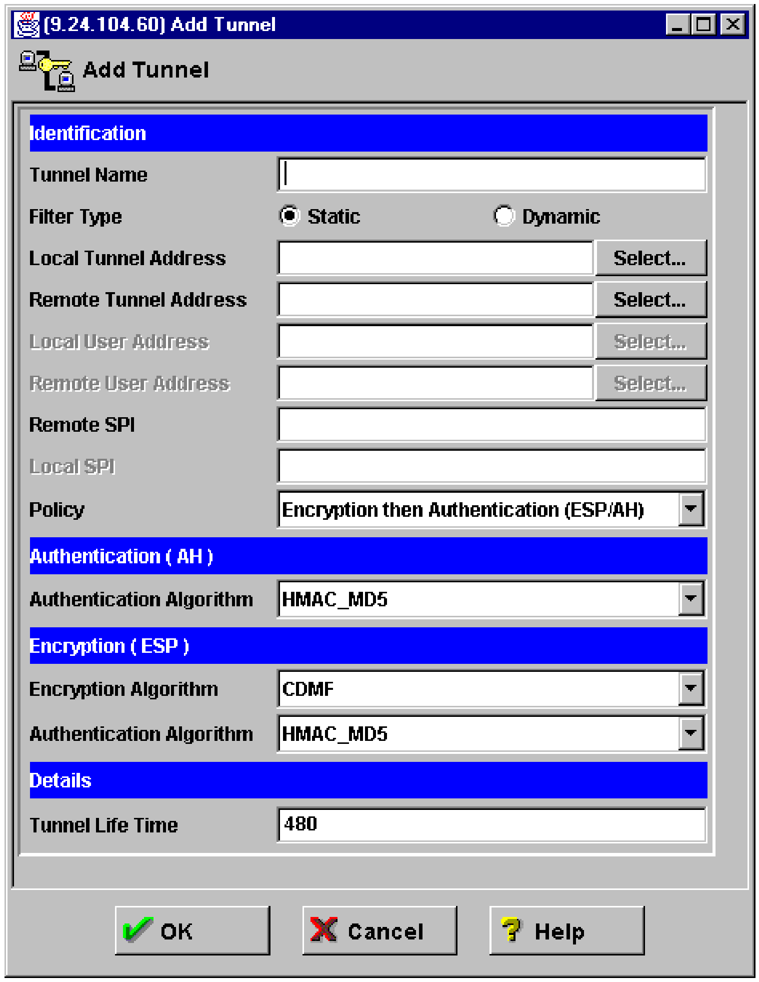

Figure 151. Adding a tunnel configuration

At this point it is very important to double-check the addresses of the tunnel. When we later import them at the partner end, there is no validation against the addresses.

In the tunnel's definition, the following fields have special importance:

- Tunnel Name: Enter the name of the tunnel.
- Filter Type: This can be either static or dynamic.

  When we select **Dynamic**, the Firewall will generate dynamic filter rules each time the tunnel is activated. This means that all the traffic between the specified networks will be accepted in the tunnel.

  When we select **Static**, we must create the filter rules for the tunnel. We have the possibility to create multiple tunnels; we can decide which protocol will be used for each tunnel. These filter rules may be more selective; for example, they may allow specific protocol traffic through the tunnel.
- Local Tunnel Address: IP address of the non-secure interface of the local firewall. Clicking **Select** gives us the list with the interfaces.
- Remote Tunnel Address: IP address of the remote partner's non-secure interface. Clicking **Select** gives us the list of the network objects.

- Local User Address: IP address of the secure network or secure host who will use the tunnel. Clicking **Select** gives us the list with the network objects (used for Dynamic Filter type only).
- Remote User Address: IP address of the remote network or host to which we will connect through the tunnel. Clicking **Select** gives us the list with the network objects (used for Dynamic Filter type only).
- Remote SPI: Specifies the security parameter index (SPI) value the tunnel partner will use. The value entered must be greater than 255. The definition of SPI is described in RFC 2401. Basically, the SPI in conjunction with the target address will uniquely identify the set of security information (such as encryption key(s), key lifetime, etc.) for your tunnel partner. You should check with the tunnel partner and obtain an unassigned SPI from it.
- Local SPI: Specifies the security parameter index (SPI) value the tunnel owner will use. The value is entered automatically.
- Policy: Define which policy we will use; we can select Authentication (AH), Encryption (ESP) or both (ESP/AH).
- Authentication Algorithm (AH): Enter the type of authentication algorithm we will use; the types available are: HMAC_MD5 and HMAC_SHA.
- Encryption Algorithm (ESP): Enter the type of authentication algorithm and encryption algorithm we will use. The encryption types are CDMF, DES_CBC, 3DES_CBC or none, depending on the country version of the firewall. For authentication we can select HMAC-MD5, HMAC-SHA or none.

  We cannot select none for both authentication and encryption with ESP.
- Tunnel Lifetime: specifies the time in minutes that the tunnel will be operational. Put a value in the entry field. The default is 480 (eight hours) and the maximum time allowed is 99999. If you specify 0, it means that the tunnel will not time out - the lifetime is unlimited.

After entering all the parameters for our tunnel and confirming this by clicking **OK**, we return to the Virtual Private Network Administration window, which shows us all the tunnels we have created. We can select a tunnel and open to modify, copy, delete, import, export, activate, or deactivate it, as shown in Figure 152.

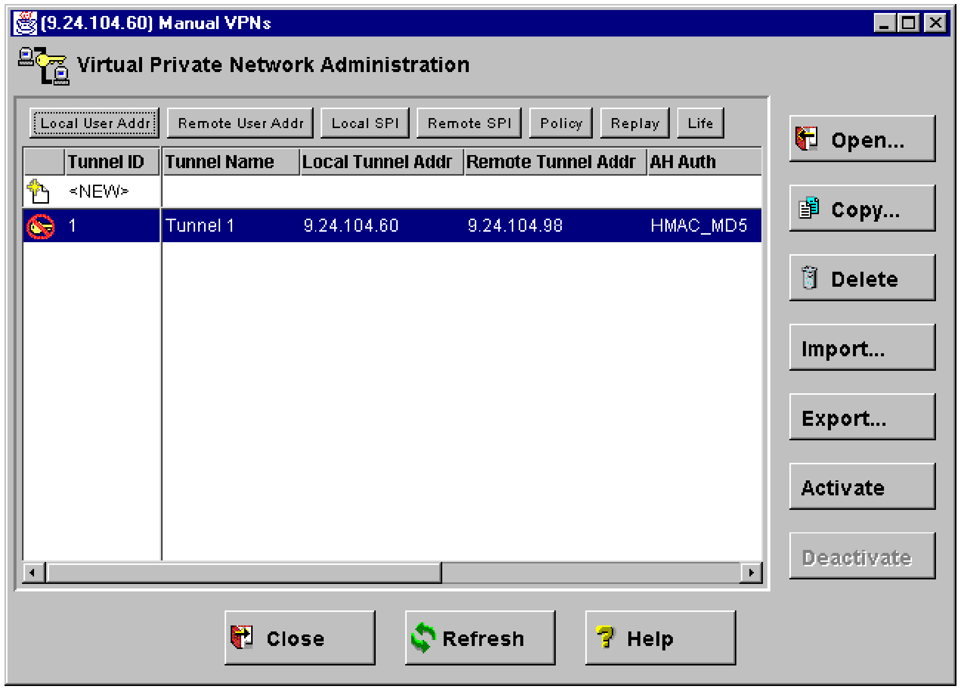

Figure 152. The tunnel list

10.3.2 Export the tunnel definition to a file

We will export the tunnel definition to a file using the Configuration Client GUI. In the tunnel list (see Figure 152, we select the tunnel(s) we want to export. An ipsec_tun_man.exp file will be generated in the specified directory.

Note that since a fixed filename is used for the export file, only one set of tunnel definitions may exist in a given directory at a time. If we are going to have tunnels between different pairs of nodes, we should create different directories for each pair of nodes.

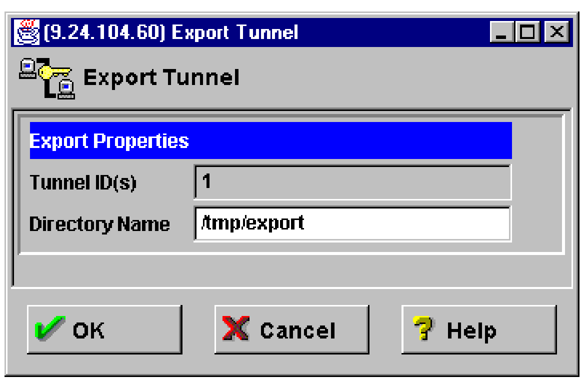

Figure 153. The export tunnel window

When the tunnel configuration export file is created correctly, we should see Figure 154.

Figure 154. Successfully exported tunnel

The following listing is an example of an export file generated by the Configuration Client GUI.

```
Line #------------------------------------
 1.  4
 2.  9.24.104.60
 3.  9.24.104.98
 4.  1
 5.  257
 6.  257
 7.  257
 8.  257
 9.  DES_CBC
10.  8
11.  0E8ED0F9A89718A5
12.  DES_CBC
13.  8
14.  CDE9CF366302F2C0
15.  HMAC_MD5
16.  16
17.  477D919F24BE56BE0E5DE2DC1DFF1DC9
18.  HMAC_MD5
19.  16
20.  81E334FED8BDFC993AABECB39E32F4D7
21.  0
22.  28800
23.  tunnel
24.  tunnel
25.  eaea
26.  0
27.  1
28.  HMAC_MD5
29.  16
30.  682BB0D42D2FE24FAE53233655B31227
31.  HMAC_MD5
32.  16
33.  1AB8AC0E86CF9F3A7BF5A5E5F7FA7342
34.  0
35.  -
36.  -
37.  Tunnel 1
38.  0
39.  0.0.0.0
40.  0.0.0.0
41.  0.0.0.0
42.  0.0.0.0
```

Figure 155. Export file

Table 28 shows the layout of the export file. When we compare the export file in Figure 155 with the import file layout in Table 28 on page 250, we see on line 38 the value 0, which means a manual tunnel with static filter rules and that authentication and encryption are used.

Table 28. Tunnel export - import file layout

Line	Field in export file	Corresponding field in tunnel structure to use when exporting	Corresponding field in tunnel structure to use when importing
1.	IP version number	IP version	IP version
2.	source address	source IP address	destination IP address
3.	destination address	destination IP address	source IP address
4.	tunnel ID	tunnel ID	tunnel ID
5.	dest encr spi > 255	remote esp spi	local esp spi
6.	dest auth spi > 255	remote ah spi	local ah spi
7.	src encr spi > 255	local esp spi	remote esp spi
8.	src auth spi > 255	local ah spi	remote ah spi
9.	receiving encr algorithm	remote esp alg	local esp alg
10.	receiving encr key length	remote esp alg length	local esp alg length
11.	receiving encr key	remote esp key	local esp key
12.	sending encr algorithm	local esp alg	remote esp alg
13.	sending encr key length	local esp alg length	remote esp alg length
14.	sending encr key	local esp key	remote esp key
15.	receiving mac algorithm	remote ah alg	local ah alg
16.	receiving mac key length	remote ah alg length	local ah alg length
17.	receiving mac key	remote ah key	local ah key
18.	sending mac algorithm	local ah alg	remote ah alg
19.	sending mac key length	local ah alg length	remote ah alg length
20.	sending mac key	local ah key	remote ah key
21.	start - defaults to 0	n/a	n/a
22.	time in seconds that the tunnel will be operational	lifetime	lifetime
23.	esp mode - must be tunnel mode	n/a	n/a
24.	ah mode - must be tunnel mode	n/a	n/a
25.	policy	local policy, remote policy	remote policy, local policy
26.	replay (=1), no replay (=0)	replay	replay
27.	new header - must be 1	n/a	n/a
28.	receiving encr mac algorithm	remote enc mac alg	local enc mac alg
29.	receiving encr mac key length	remote enc mac key length	local enc mac key length

Line	Field in export file	Corresponding field in tunnel structure to use when exporting	Corresponding field in tunnel structure to use when importing
30.	receiving encr mac key	remote enc mac key	local enc mac key
31.	sending encr mac algorithm	local enc mac alg	remote enc mac alg
32.	sending encr mac key length	local enc mac key length	remote enc mac key length
33.	sending encr mac key	local enc mac key	remote enc mac key
34.	through FW	n/a	n/a
35.	FW address	n/a	n/a
36.	destination mask	n/a	n/a
37.	tunnel_name	tunnel name	tunnel name
38.	filter type, static (=0) dynamic (=1)	filter type	filter type
39.	source user address (only dynamic)	local user IP address	remote user IP address
40.	source user mask (only dynamic)	local user mask	remote user mask
41.	dest user address (only dynamic)	remote user IP address	local user IP address
42.	dest user mask (only dynamic)	remote user mask	local user mask

When exporting and importing the tunnel configuration, the owner and partner authentication and encryption keys are switched. The sending key on the owner firewall is the receiving key on the partner.

Note: The SPI authentication and encryption keys are switched when exporting the tunnel configuration. The source and destination addresses are switched when importing the tunnel configuration on the partner.

If the tunnel ID is already defined in the receiving firewall, it will be changed automatically; if the tunnel has dynamic filter rules, they will also be changed. If the tunnel has static filter rules, the user will have to point to the new tunnel ID in the filter rules.

10.3.3 Import the tunnel definition in the partner node

After exporting we have to take the file from the local firewall to the partner firewall. Currently, IBM SecureWay Firewall V4.1 for AIX does not provide any mechanism to do this transfer. The file contains the encryption key for the secure tunnel, so we should devise a secure way to transmit them. We can copy the file on a diskette and go to the partner firewall or we can send the file within an encrypted mail message.

In the partner node, once we have received the files we click **Import** (see Figure 152 on page 248), and we can import the definitions as shown in Figure 156.

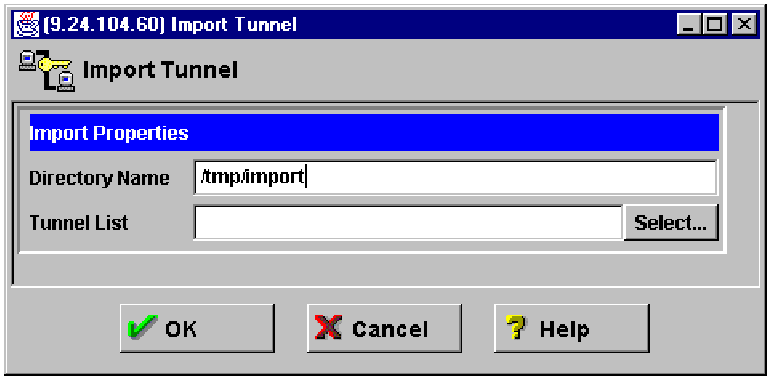

Figure 156. Import tunnel window

We type the directory name and then click **Select** to get the list of tunnels we can import as shown in Figure 157.

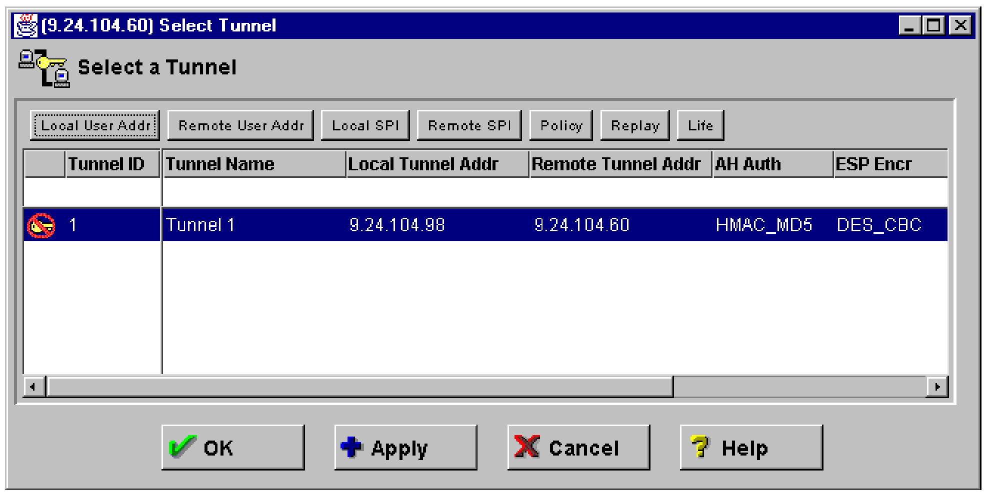

Figure 157. Tunnel import selection window

Once we have selected the tunnels we want to import, we click **OK** and return to the import tunnel window (Figure 156), where we again click on **OK** to import our tunnels. When imported successfully we should see the window in Figure 158.

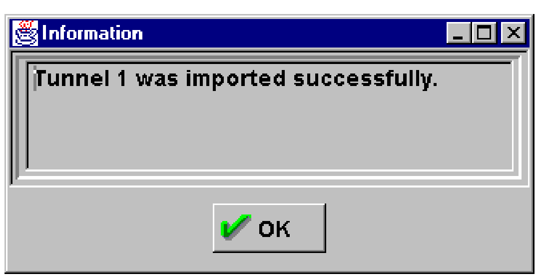

Figure 158. Successfully imported tunnel

The import function swaps the source and destination addresses.

Importing tunnels

The Import function can only be used without modifying the export file if the owner of the tunnel is an IBM SecureWay Firewall V4.1 for AIX or an IBM eNetwork Firewall for Windows NT 3.3 or an IBM SecureWay Firewall for Windows NT V4.1. When the IBM SecureWay Firewall V4.1 for AIX is *not* the owner of the tunnel, the export files on the IBM SecureWay Firewall V4.1 for AIX *must* be updated manually.

10.3.4 Activate/deactivate the tunnel at both ends

Because we are creating a tunnel with dynamic filter rules, we do not need to create any filter rules or activate the filter rules. We can activate the tunnel at both ends using the Configuration Client GUI. Select the tunnel that you would like to activate and click **Activate**.

Figure 152 shows the tunnel list before activating the tunnel, and Figure 159 shows the active tunnel. You can see that the icon at the left of the tunnel ID is different after the tunnel is activated.

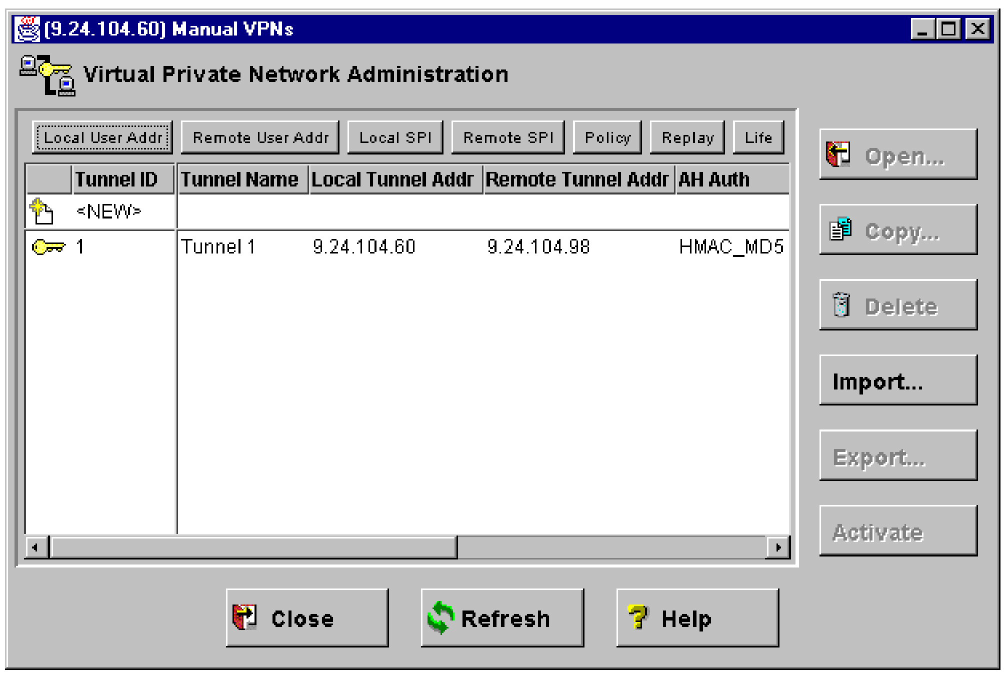

Figure 159. List showing an active tunnel

The tunnel will be marked active even if the other end is not running or connected.

If we want to stop the tunnel, we select it and click the **Deactivate** button. Every time we activate or deactivate a tunnel, a message will be written in the log file (see Figure 160).

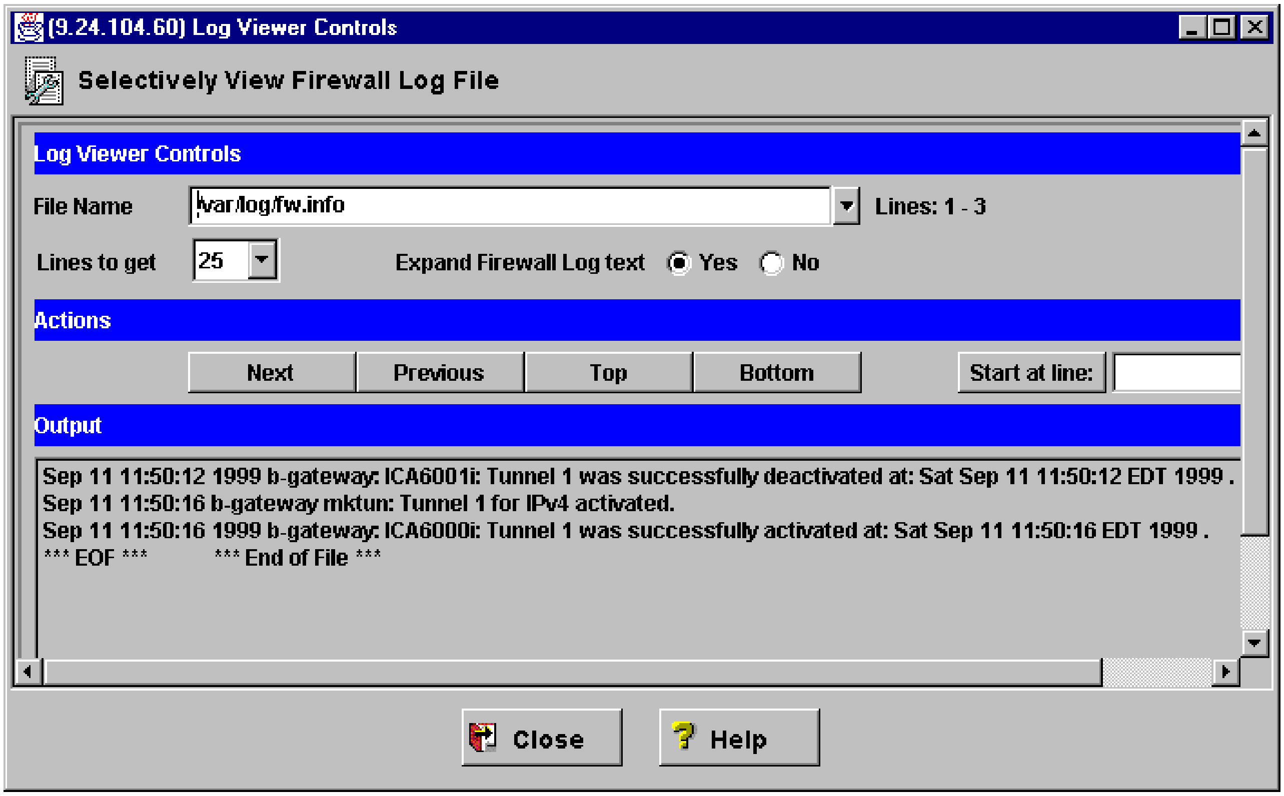

Figure 160. Tunnel logging

When we have executed all the previous steps successfully, we have a tunnel running. When dynamic filters are used, all traffic between the specified users will be transported through the tunnel without regard to the protocol of the traffic. The filter rules are automatically activated.

For example, the filter rules created for the sample configuration in Figure 161are shown in Figure 162.

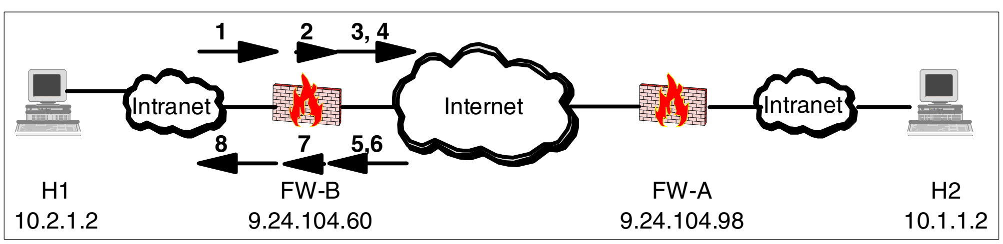

Figure 161. Sample of manual tunnel with dynamic filter rules

These are the corresponding filter rules:

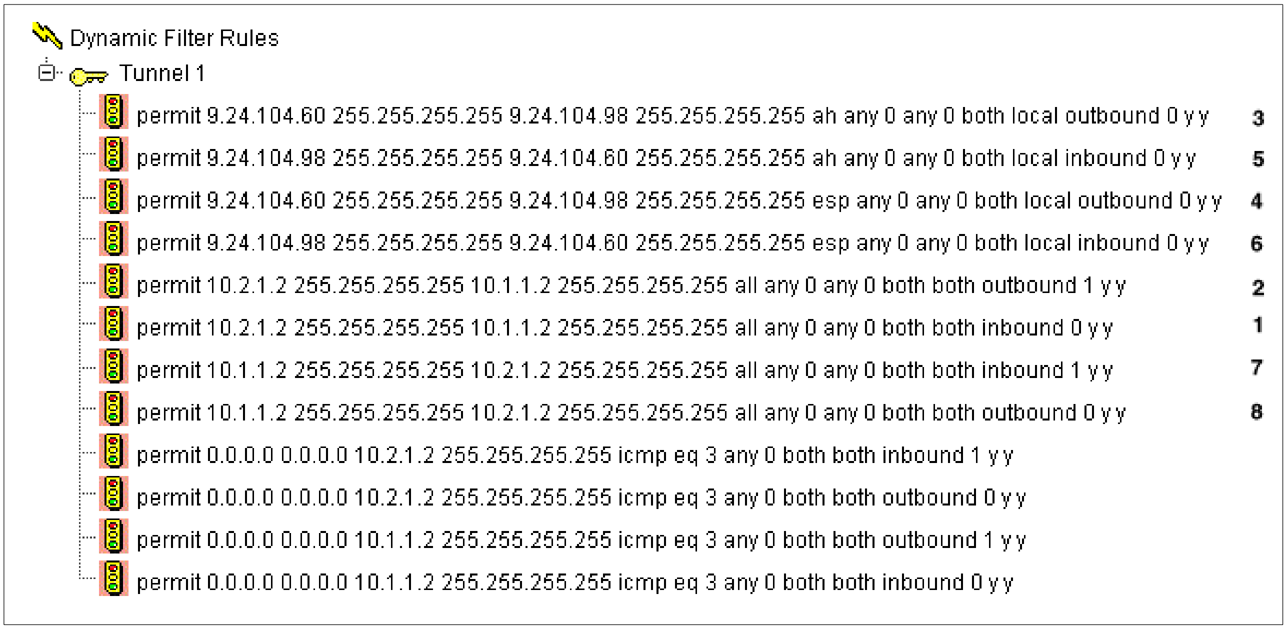

Figure 162. Generated dynamic filter rules

The first eight rules in Figure 162 correspond to the rules in Figure 161; the numbers on the right side of Figure 162 correspond with the numbers of the arrows in the previous figure. These rules are rather generic, so they can be used for local or routed traffic, and for tunnels that go to the Internet or to the secure intranet.

The last four rules are added to the set of dynamic filter rules for the tunnel to allow ICMP type 3, code <any> error packets to be sent back through the tunnel to the packet originator. This way the client can be informed about any "Destination Unreachable" errors for the connection. For example, if the client (originator) was employing Path MTU Discovery, if the packet hit a router in the remote secure net that needed to fragment the packet but couldn't because the "don't fragment" (DF) bit was set, the ICMP Error type 3, code 4 packet that would be generated would be permitted by the last four rules to travel back through the tunnel to the originating client.

10.3.5 Using static filter rules

By using static filter rules, we can be more specific about which traffic is going to flow through each tunnel. These rules will be like normal rules (with source, target, protocol, ports and port operations), but some of them will also have a tunnel ID. So when a packet must be transferred, the IBM SecureWay Firewall V4.1 for AIX will search the filtering rules. If it matches a rule, and this rule has a specific tunnel ID, the packet will be sent according to the authentication/encryption rules specified in this specific tunnel.

The IBM SecureWay Firewall V4.1 for AIX has provided three services that we can use to define our own rules as shown in Figure 163.

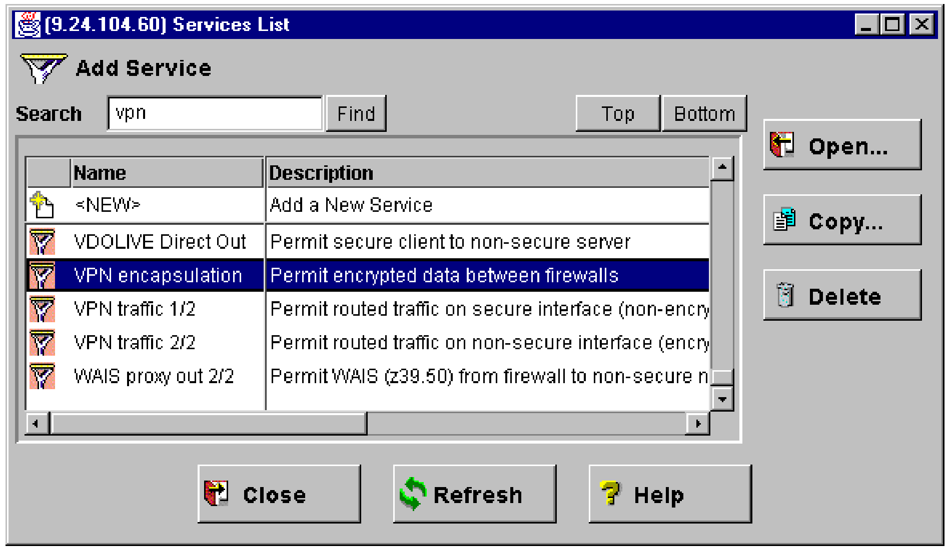

Figure 163. VPN services

We will create a tunnel with static filter rules using the predefined services. We have to create two connections, one for the encapsulation between the two non-secure sides of the firewalls and one for the data traffic between the two secure clients as shown in Figure 164.

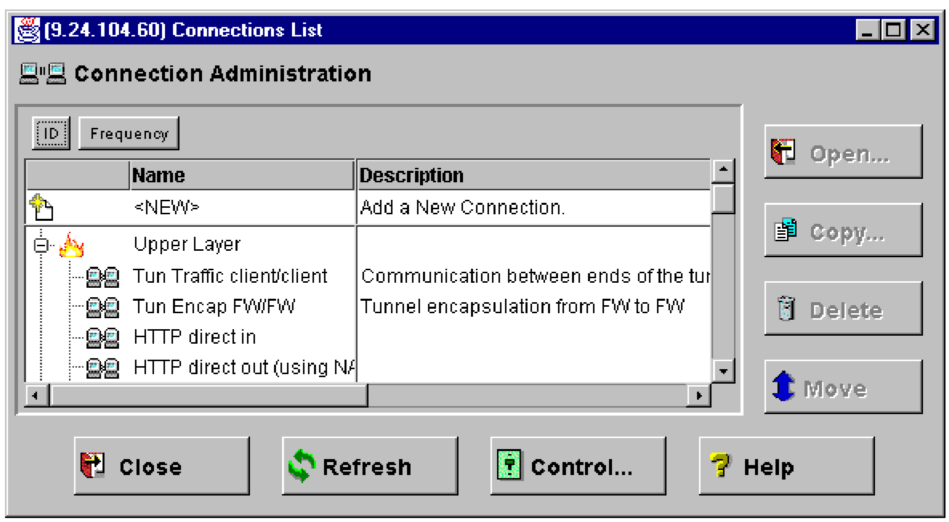

Figure 164. Connection list

The encapsulation takes place between the two non-secure interfaces of the tunnel partners. For this connection we can always use the predefined service. This will be the same for tunnels with static and dynamic filter rules. We can see the connection setup in Figure 165. In Figure 167 on page 259 we can see the four filter rules that are created.

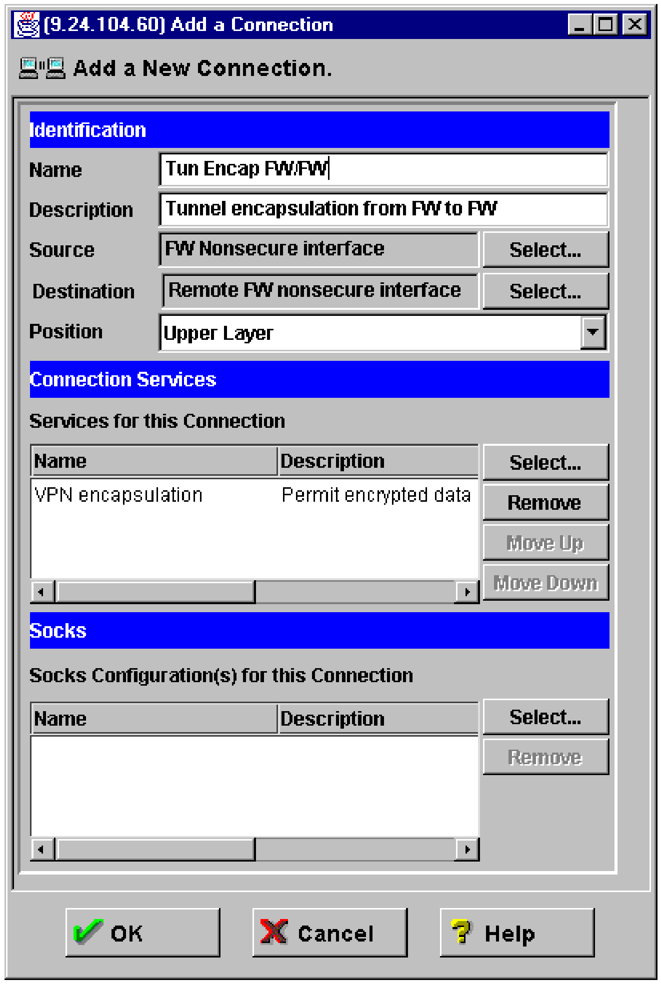

Figure 165. Tunnel connection encapsulation

The data traffic connection is established between the two clients on each secure side. Here we have two services: one to transport the data from the secure client to the secure interface of the firewall, and one from the non-secure interface to the tunnel. To get the traffic into the tunnel we need to define the tunnel ID in the second part of the traffic connection. For that reason we cannot use the default service; we *must* copy the VPN traffic 2/2 service into a new service VPN traffic 2/2 Tunnel 3 in our example, and define the tunnel ID we want to use.

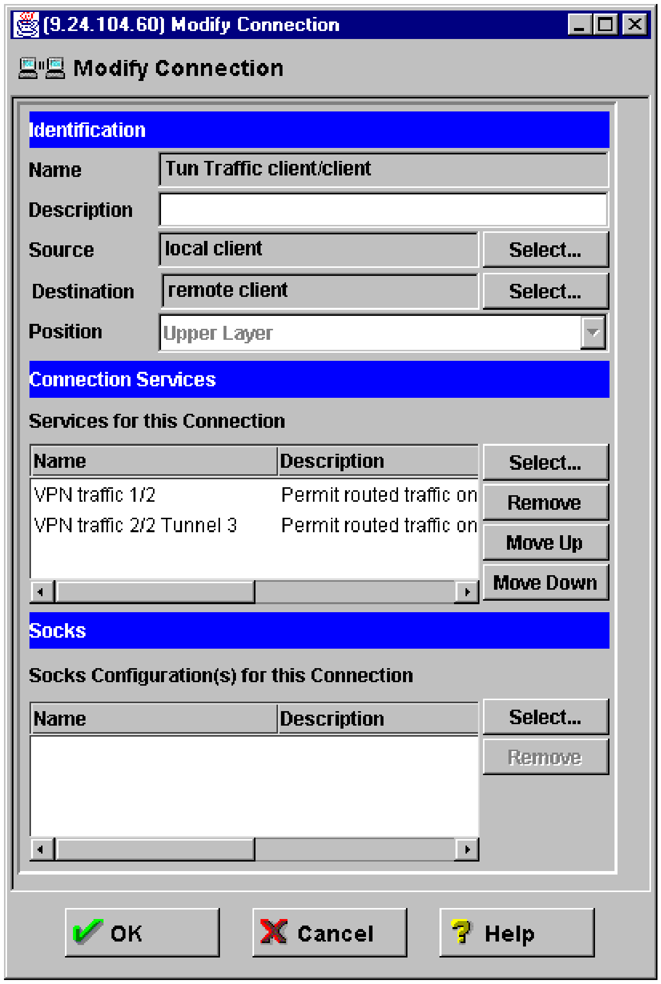

Figure 166. Tunnel connection traffic

When we look at Figure 167we see the eight rules that are created. The first two provide the traffic from the secure client to the secure interface of the firewall. The next two are for the transport from the non-secure interface to the tunnel. At the end of the rules we see that all our traffic will be using Tunnel 3 (the number three right after the words outbound and inbound). The last four rules are for the encapsulation.

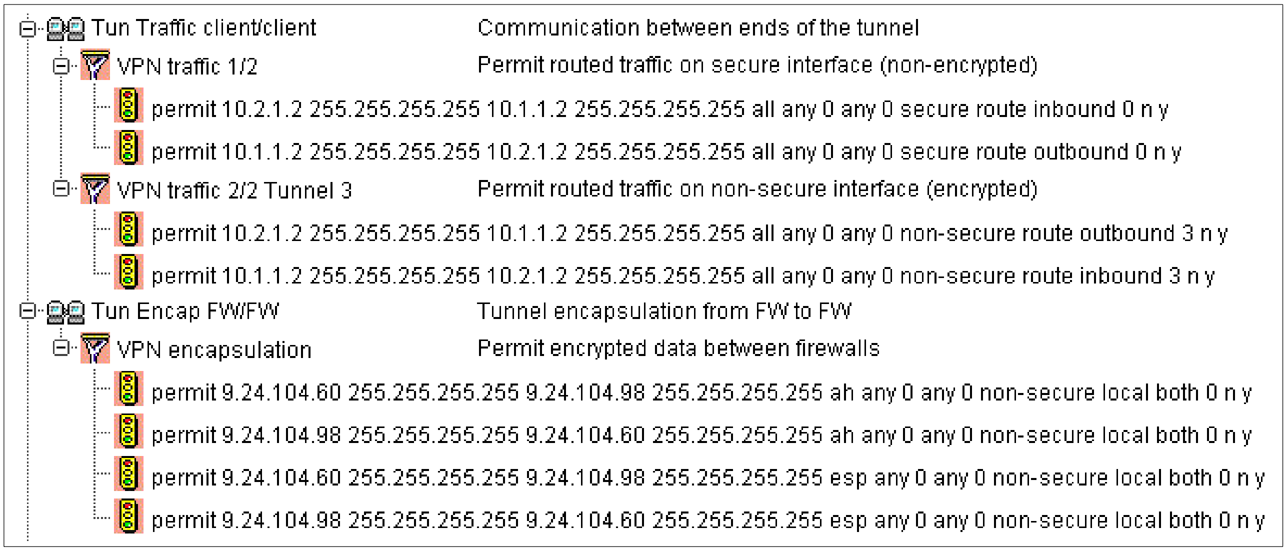

Figure 167. Static rules

This is when we use the default services that come with the IBM SecureWay Firewall V4.1 for AIX. In 10.5, “Virtual Private Network scenarios” on page 266, we see more detailed examples.[1]

Note

When you are working with VPN tunnels, you should not "open" the firewall by selecting the Test IP Routing (debug only) option in the Security Policy window. If you do so, the filter rules that direct the traffic to a specific tunnel ID will not be reached, and no traffic will flow through the tunnels.

10.3.6 Reactivate tunnel when lifetime has expired

When the tunnel lifetime is reached, the tunnel will cease operation until it is reactivated.

To reactivate the tunnel, simply select the tunnel and click the **Activate** button. However, reactivating a tunnel would only reactivate it to an operational state. The keys used in the tunnel remain the same. To re-establish the tunnel with new session keys, you need to delete the tunnel, and then add a new tunnel with the same tunnel ID and characteristics back into the firewall. After that, you export the new tunnel definition to the tunnel partner. New session keys are stored inside the definition file. Your tunnel partner is also required to delete the existing tunnel and then re-import the new definition.

A convenient shortcut to generate new keys for a tunnel would be to modify the algorithms used in the tunnel (e.g., change to HMAC-SHA from HMAC-MD5, or CDMF from DES), save the modified tunnel definition, and finally, modify them back to the original set of algorithms (and save again). This effectively regenerates keys for the tunnel. Remember to modify all used algorithms.

[1] Optionally you may want to add the ICMP filters as shown in Figure 162., “Generated dynamic filter rules” on page 255

10.3.7 Summary

The following list is a summary of the steps to create and activate a tunnel:

1. Create a firewall object for the non-secure interface of the remote firewall.
2. Create a network object for the secure network of the remote firewall (or for the specific hosts with which we want to be connected).
3. Create the tunnel itself; local address=nonsecure interface of the local firewall and remote address=remote firewall object (from item 1).
4. Export definitions, transport to the remote firewall, import definitions (which automatically switches local and remote addresses).
5. Add connection services for both VPN encapsulation; source=nonsecure interface of my firewall, destination=remote firewall object from item 1. This allows firewall-to-firewall communication of encapsulated data.
6. Copy the VPN 2/2 rule to a new rule called "VPN traffic 2/2 tunnel xx" and set the tunnel ID in that rule.
7. Create a connection service "VPN traffic 1/2" and "VPN traffic 2/2 tunnel xx" from item 6. Source=my secure network (or the set of hosts allowed to use VPN), destination=secure network of remote firewall from 2).
8. Ensure ip forwarding is on.
9. Repeat steps 1, 2, 5, 6, 7 and 8 at the remote firewall. (Note that now "remote" and "local" are relative to that firewall.)
10. Activate rule sets.
11. Activate the tunnel at both ends.
12. Try to ping between the networks; it should work!

10.4 Authentication and encryption examples

In order to understand how the tunnels work, we will show two examples, both using the same tunnel, in the network shown in Figure 174 on page 267. Both firewalls are configured to allow routed FTP traffic using the default services VPN Traffic 1/2 and VPN Traffic 2/2. We made an FTP connection from client 10.1.1.2 to FTP server 10.2.1.2 and executed the `DIR` command, which gave us the output in Figure 168. In the first example we use only authentication and in the second authentication and encryption. We compared the data packet that was transported from the FTP server to the client.

Telnet - 10.2.1.2

Connect Edit Terminal Help

```
root@b-mail.b-corp-secure.ibm.com:/ > ftp 10.1.1.2
Connected to 10.1.1.2.
220 a-mail.a-corp-secure.ibm.com FTP server (Version 4.1 Mon Jul 26 19:58:48 CDT
 1999) ready.
Name (10.1.1.2:root): cris
331 Password required for cris.
Password:
230 User cris logged in.
ftp> dir
200 PORT command successful.
150 Opening data connection for /bin/ls.
total 10
-rwxr-----   1 cris     staff        267 Sep 07 15:57 .profile
-rw-------   1 cris     staff       1634 Sep 14 11:24 .sh_history
-rw-r--r--   1 cris     staff       1024 Sep 08 18:10 a
-rw-r--r--   1 cris     staff       1019 Sep 08 18:18 b
drwx------   2 cris     staff        512 Sep 08 11:37 mail
226 Transfer complete.
ftp> bye
221 Goodbye.
root@b-mail.b-corp-secure.ibm.com:/ >
```

Figure 168. FTP download

Before using the tunnel we took a trace of the FTP data packet that was sent to the client. This packet will be compared with the packets in each example. Figure 169 shows us the data packet not using a tunnel (this was acquired using the command `iptrace`).

```
IP:     < SRC =        10.1.1.2 >
IP:     < DST =        10.2.1.2 >
IP:     ip_v=4, ip_hl=20, ip_tos=8, ip_len=355, ip_id=30936, ip_off=0
IP:     ip_ttl=58, ip_sum=f0ae, ip_p = 6 (TCP)
TCP:    <source port=20(ftp-data), destination port=35837 >
TCP:    th_seq=fcb0019d, th_ack=23d5a167
TCP:    th_off=5, flags<PUSH | ACK>
TCP:    th_win=15972, th_sum=9dac, th_urp=0
TCP: 00000000     746f7461 6c203130 0d0a2d72 7778722d     |total 10..-rwxr-|
TCP: 00000010     2d2d2d2d 20202031 20637269 73202020     |----   1 cris   |
TCP: 00000020     20207374 61666620 20202020 20202032     |  staff        2|
TCP: 00000030     36372053 65702030 37203135 3a353720     |67 Sep 07 15:57 |
TCP: 00000040     2e70726f 66696c65 0d0a2d72 772d2d2d     |.profile..-rw---|
TCP: 00000050     2d2d2d2d 20202031 20637269 73202020     |----   1 cris   |
TCP: 00000060     20207374 61666620 20202020 20203136     |  staff       16|
TCP: 00000070     33342053 65702031 34203131 3a323420     |34 Sep 14 11:24 |
TCP: 00000080     2e73685f 68697374 6f72790d 0a2d7277     |.sh_history..-rw|
TCP: 00000090     2d722d2d 722d2d20 20203120 63726973     |-r--r--   1 cris|
TCP: 000000a0     20202020 20737461 66662020 20202020     |     staff      |
TCP: 000000b0     20313032 34205365 70203038 2031383a     | 1024 Sep 08 18:|
TCP: 000000c0     31302061 0d0a2d72 772d722d 2d722d2d     |10 a..-rw-r--r--|
TCP: 000000d0     20202031 20637269 73202020 20207374     |   1 cris     st|
TCP: 000000e0     61666620 20202020 20203130 31392053     |aff       1019 S|
TCP: 000000f0     65702030 38203138 3a313820 620d0a64     |ep 08 18:18 b..d|
TCP: 00000100     7277782d 2d2d2d2d 2d202020 32206372     |rwx------   2 cr|
TCP: 00000110     69732020 20202073 74616666 20202020     |is     staff    |
TCP: 00000120     20202020 35313220 53657020 30382031     |    512 Sep 08 1|
TCP: 00000130     313a3337 206d6169 6c0d0a                |1:37 mail..     |
```

Figure 169. FTP data packet (no tunnel)

In Figure 169 you can see the packet sent by the FTP server (10.1.1.2) containing the output of the `DIR` command, as shown in Figure 168. The destination address is 10.2.1.2, which is our client. You can see the translation of the contents of the packet on the column on the right. Note that the size of this packet is 355 bytes (see `ip_len=355` in the third line). Since we are not using any encryption, we can read the contents of the packet by reading this iptrace.

10.4.1 Authentication example

In this first example we send the packet through the tunnel using authentication only. To get a better view we have traced the packet before (see Figure 170) and after (see Figure 171) it was authenticated.

```
IP:     < SRC =        10.1.1.2 >
IP:     < DST =        10.2.1.2 >
IP:     ip_v=4, ip_hl=20, ip_tos=8, ip_len=355, ip_id=30992, ip_off=0
IP:     ip_ttl=58, ip_sum=f076, ip_p = 6 (TCP)
TCP:    <source port=20(ftp-data), destination port=35842 >
TCP:    th_seq=640250bd, th_ack=2f129567
TCP:    th_off=5, flags<PUSH | ACK>
TCP:    th_win=15972, th_sum=e7f8, th_urp=0
TCP: 00000000     746f7461 6c203130 0d0a2d72 7778722d     |total 10..-rwxr-|
TCP: 00000010     2d2d2d2d 20202031 20637269 73202020     |----   1 cris   |
TCP: 00000020     20207374 61666620 20202020 20202032     |  staff        2|
TCP: 00000030     36372053 65702030 37203135 3a353720     |67 Sep 07 15:57 |
TCP: 00000040     2e70726f 66696c65 0d0a2d72 772d2d2d     |.profile..-rw---|
TCP: 00000050     2d2d2d2d 20202031 20637269 73202020     |----   1 cris   |
TCP: 00000060     20207374 61666620 20202020 20203136     |  staff       16|
TCP: 00000070     33342053 65702031 34203131 3a323420     |34 Sep 14 11:24 |
TCP: 00000080     2e73685f 68697374 6f72790d 0a2d7277     |.sh_history..-rw|
TCP: 00000090     2d722d2d 722d2d20 20203120 63726973     |-r--r--   1 cris|
TCP: 000000a0     20202020 20737461 66662020 20202020     |     staff      |
TCP: 000000b0     20313032 34205365 70203038 2031383a     | 1024 Sep 08 18:|
TCP: 000000c0     31302061 0d0a2d72 772d722d 2d722d2d     |10 a..-rw-r--r--|
TCP: 000000d0     20202031 20637269 73202020 20207374     |   1 cris     st|
TCP: 000000e0     61666620 20202020 20203130 31392053     |aff       1019 S|
TCP: 000000f0     65702030 38203138 3a313820 620d0a64     |ep 08 18:18 b..d|
TCP: 00000100     7277782d 2d2d2d2d 2d202020 32206372     |rwx------   2 cr|
TCP: 00000110     69732020 20202073 74616666 20202020     |is     staff    |
TCP: 00000120     20202020 35313220 53657020 30382031     |    512 Sep 08 1|
TCP: 00000130     313a3337 206d6169 6c0d0a                |1:37 mail..     |
```

Figure 170. FTP data packet before authentication

When we compare the packet in Figure 170 before authentication with the packet in Figure 169, we see there is no difference. The source and destination address are still the same, the data sent is still 355 bytes, and no header has been added.

```
IP:     < SRC =     9.24.104.98 >
IP:     < DST =     9.24.104.60 >
IP:     ip_v=4, ip_hl=20, ip_tos=8, ip_len=399, ip_id=31293, ip_off=0
IP:     ip_ttl=59, ip_sum=2129, ip_p = 51 (unknown internet protocol)
IP:  00000000     04040000 00000134 00000017 ec0c630a     |.......4......c.|
IP:  00000010     bce4d594 0d94e861 45080163 79100000     |.......aE..cy...|
IP:  00000020     3b06ef76 0a010102 0a020102 00148c02     |;..v............|
IP:  00000030     640250bd 2f129567 50183e64 e7f80000     |d.P./..gP.>d....|
IP:  00000040     746f7461 6c203130 0d0a2d72 7778722d     |total 10..-rwxr-|
IP:  00000050     2d2d2d2d 20202031 20637269 73202020     |----   1 cris   |
IP:  00000060     20207374 61666620 20202020 20202032     |  staff        2|
IP:  00000070     36372053 65702030 37203135 3a353720     |67 Sep 07 15:57 |
IP:  00000080     2e70726f 66696c65 0d0a2d72 772d2d2d     |.profile..-rw---|
IP:  00000090     2d2d2d2d 20202031 20637269 73202020     |----   1 cris   |
IP:  000000a0     20207374 61666620 20202020 20203136     |  staff       16|
IP:  000000b0     33342053 65702031 34203131 3a323420     |34 Sep 14 11:24 |
IP:  000000c0     2e73685f 68697374 6f72790d 0a2d7277     |.sh_history..-rw|
IP:  000000d0     2d722d2d 722d2d20 20203120 63726973     |-r--r--   1 cris|
IP:  000000e0     20202020 20737461 66662020 20202020     |     staff      |
IP:  000000f0     20313032 34205365 70203038 2031383a     | 1024 Sep 08 18:|
IP:  00000100     31302061 0d0a2d72 772d722d 2d722d2d     |10 a..-rw-r--r--|
IP:  00000110     20202031 20637269 73202020 20207374     |   1 cris     st|
IP:  00000120     61666620 20202020 20203130 31392053     |aff       1019 S|
IP:  00000130     65702030 38203138 3a313820 620d0a64     |ep 08 18:18 b..d|
IP:  00000140     7277782d 2d2d2d2d 2d202020 32206372     |rwx------   2 cr|
IP:  00000150     69732020 20202073 74616666 20202020     |is     staff    |
IP:  00000160     20202020 35313220 53657020 30382031     |    512 Sep 08 1|
IP:  00000170     313a3337 206d6169 6c0d0a                |1:37 mail..     |
```

Figure 171. FTP data packet after authentication

When we look at the packet after authentication (see Figure 171) and compare it with the packet before authentication (see Figure 170) we see that the source and destination addresses have changed. The source address, shown in the first line (`SRC = 10.1.1.2`), is now the address of the non-secure interface of the sending firewall FW-A (`SRC = 9.24.104.98`) and the destination address, shown in the second line (`DST = 10.2.1.2`), is now the address of the non-secure interface (`DST = 9.24.104.60`) of the receiving firewall (FW-B).

The size of the data packet is no longer 355 but 399 bytes (see `ip_len=399` on the third line). Between the new IP header and the original IP header a new AH header is added. This new added header is the authentication header. The format of the authentication header can be found in RFC 2402. One of the values in the authentication header is the Security Parameter Index (SPI).

You can also see the original source and destination IP addresses included in the packet: line 00000020, data `0a010102` (10.1.1.2) and `0a020102` (10.2.1.2).

When creating a tunnel for the first time, the local SPI is automatically set to 256; when transporting packets the authentication header always contains the local SPI of the receiving tunnel endpoint. The original IP packet is encapsulated with a new IP header containing the addresses of the two endpoints of our tunnel, and an authentication header is added.

10.4.2 Encryption example

For the encryption example we will execute exactly the same FTP command as we did for the authentication, and now we are doing the same procedure using only encryption.

```
IP:     < SRC =        10.1.1.2 >
IP:     < DST =        10.2.1.2 >
IP:     ip_v=4, ip_hl=20, ip_tos=8, ip_len=355, ip_id=31014, ip_off=0
IP:     ip_ttl=58, ip_sum=f060, ip_p = 6 (TCP)
TCP:    <source port=20(ftp-data), destination port=35844 >
TCP:    th_seq=9d3802a8, th_ack=38952767
TCP:    th_off=5, flags<PUSH | ACK>
TCP:    th_win=15972, th_sum=6153, th_urp=0
TCP: 00000000     746f7461 6c203130 0d0a2d72 7778722d     |total 10..-rwxr-|
TCP: 00000010     2d2d2d2d 20202031 20637269 73202020     |----   1 cris   |
TCP: 00000020     20207374 61666620 20202020 20202032     |  staff        2|
TCP: 00000030     36372053 65702030 37203135 3a353720     |67 Sep 07 15:57 |
TCP: 00000040     2e70726f 66696c65 0d0a2d72 772d2d2d     |.profile..-rw---|
TCP: 00000050     2d2d2d2d 20202031 20637269 73202020     |----   1 cris   |
TCP: 00000060     20207374 61666620 20202020 20203136     |  staff       16|
TCP: 00000070     33342053 65702031 34203131 3a323420     |34 Sep 14 11:24 |
TCP: 00000080     2e73685f 68697374 6f72790d 0a2d7277     |.sh_history..-rw|
TCP: 00000090     2d722d2d 722d2d20 20203120 63726973     |-r--r--   1 cris|
TCP: 000000a0     20202020 20737461 66662020 20202020     |     staff      |
TCP: 000000b0     20313032 34205365 70203038 2031383a     | 1024 Sep 08 18:|
TCP: 000000c0     31302061 0d0a2d72 772d722d 2d722d2d     |10 a..-rw-r--r--|
TCP: 000000d0     20202031 20637269 73202020 20207374     |   1 cris     st|
TCP: 000000e0     61666620 20202020 20203130 31392053     |aff       1019 S|
TCP: 000000f0     65702030 38203138 3a313820 620d0a64     |ep 08 18:18 b..d|
TCP: 00000100     7277782d 2d2d2d2d 2d202020 32206372     |rwx------   2 cr|
TCP: 00000110     69732020 20202073 74616666 20202020     |is     staff    |
TCP: 00000120     20202020 35313220 53657020 30382031     |    512 Sep 08 1|
TCP: 00000130     313a3337 206d6169 6c0d0a                |1:37 mail..     |
```

Figure 172. FTP data packet before encryption

When we compare the packet in Figure 172, the packet before encryption, with the packet in Figure 169, using no tunnel, and with the packet in Figure 170 before authentication, we see no difference. The source address is still the same (`SRC = 10.1.1.2`) and it continues with the destination address `DST = 10.2.1.2`.

Figure 173 shows the packet after encryption. The size of the data packet is now 396 (see `ip_len=396` in the third line). The packet is encapsulated with a new IP header and contains the two endpoints of the tunnel (`SRC = 9.24.104.98` and `DST = 9.24.104.60`). All the headers and the data that comes after the new IP header is encrypted as we can see in the right column. The original source and destination are no longer readable nor is the data that we are transporting.

```
IP:     < SRC =     9.24.104.98 >
IP:     < DST =     9.24.104.60 >
IP:     ip_v=4, ip_hl=20, ip_tos=8, ip_len=396, ip_id=31308, ip_off=0
IP:     ip_ttl=59, ip_sum=211e, ip_p = 50 (unknown internet protocol)
IP:  00000000     00000135 00000008 f2cc2f6d ecc350cb     |...5....../m..P.|
IP:  00000010     ecf7bb25 1c3cb9b6 9509a7cc 5d796459     |...%.<......]ydY|
IP:  00000020     dd0510ac edf2ed76 408fdbc7 531de601     |.......v@...S...|
IP:  00000030     db9af973 82ee6b12 8bf13acb 78d32623     |...s..k...:.x.&#|
IP:  00000040     ec73ebb9 263deafc 1c699b6a 608875be     |.s..&=...i.j`.u.|
IP:  00000050     018b413b f527e85b a54d470b bfce5278     |..A;.'.[.MG...Rx|
IP:  00000060     7514877c 3cb1f4e8 aecc72df f8123fbd     |u..|<.....r...?.|
IP:  00000070     c3cb9e30 d19cf364 7cc18cbc 7ca162d8     |...0...d|...|.b.|
IP:  00000080     8f368a35 dfe15723 c47d4c70 4c064e4e     |.6.5..W#.}LpL.NN|
IP:  00000090     541bf084 3e2fdf77 c08b7471 79d6d31e     |T...>/.w..tqy...|
IP:  000000a0     f36af8cf 24dc6330 91101acc 440cd209     |.j..$.c0....D...|
IP:  000000b0     a27acdd4 6a0561cc 6132fcc6 2e2dd799     |.z..j.a.a2...-..|
IP:  000000c0     38984969 f1bf9462 d4ba8764 b457d9d5     |8.Ii...b...d.W..|
IP:  000000d0     184582ac 1af14e0f acb8aa78 e3b9182e     |.E....N....x....|
IP:  000000e0     4225a9a2 24ddd18e 50f38cc7 3932e54e     |B%..$...P...92.N|
IP:  000000f0     987219a1 854bc360 31094218 dac4081a     |.r...K.`1.B.....|
IP:  00000100     e367761f f93b65f1 6784e1d8 ff3b3bed     |.gv..;e.g....;;.|
IP:  00000110     0c3373a1 4d2a3ecc d3263b5c fd7950b2     |.3s.M*>..&;\.yP.|
IP:  00000120     ff2c8c1f 7cbea53a fa8534c5 12e8ef5f     |.,..|..:..4...._|
IP:  00000130     330c5462 22a997c0 f49e4853 0c479ae9     |3.Tb".....HS.G..|
IP:  00000140     bc7a51ce da342253 b7174442 594a473b     |.zQ..4"S..DBYJG;|
IP:  00000150     ba83412a 27f81eee ab442bdc f3a5d415     |..A*'....D+.....|
IP:  00000160     063181fb 54afd1ba b44927e1 9112e16a     |.1..T....I'....j|
IP:  00000170     3a9a14dd c4842744                       |:.....'D        |
```

Figure 173. Packet after encryption

10.5 Virtual Private Network scenarios

We have configured the following tunnels:

1. Between two IBM SecureWay Firewall V4.1 for AIX using static filter rules.
2. Between IBM SecureWay Firewall V4.1 for AIX and IBM eNetwork Firewall for Windows NT using dynamic filter rules.
3. Between IBM SecureWay Firewall V4.1 for AIX and AIX 4.3.2 using dynamic filter rules.

In the first scenario, we created the traffic rules on the partner tunnel manually. In the other scenarios we used dynamic filters.

When configuring a tunnel with another system different from the IBM SecureWay Firewall V4.1 for AIX, you must be sure that the authentication or encryption method you want to use is supported.

Since the VPN support in IBM SecureWay Firewall V4.1 for AIX is very similar to the VPN support in the IBM eNetwork Firewall for Windows 3.3, you can read how to establish a VPN tunnel with OS/390 and OS/400 in the redbook *Guarding the Gates Using the IBM eNetwork Firewall V3.3 for Windows NT,* SG24-5209.

10.5.1 Tunnel between two IBM SecureWay Firewall V4.1 for AIX

In this scenario we will configure one tunnel using static filter rules between our two secure networks over a non-secure network. Instead of using the predefined rules, we will need to create our own rules, since we are configuring only one specific service through this tunnel: Telnet.

This network has two firewalls, FW-A and FW-B, protecting the secure networks 10.1.1.0 and 10.2.1.0. Their non-secure IP addresses are respectively 9.24.104.98 and 9.24.104.60. The Telnet session will be established between the client 10.2.1.2 and the server 10.1.1.2, as shown in Figure 174.

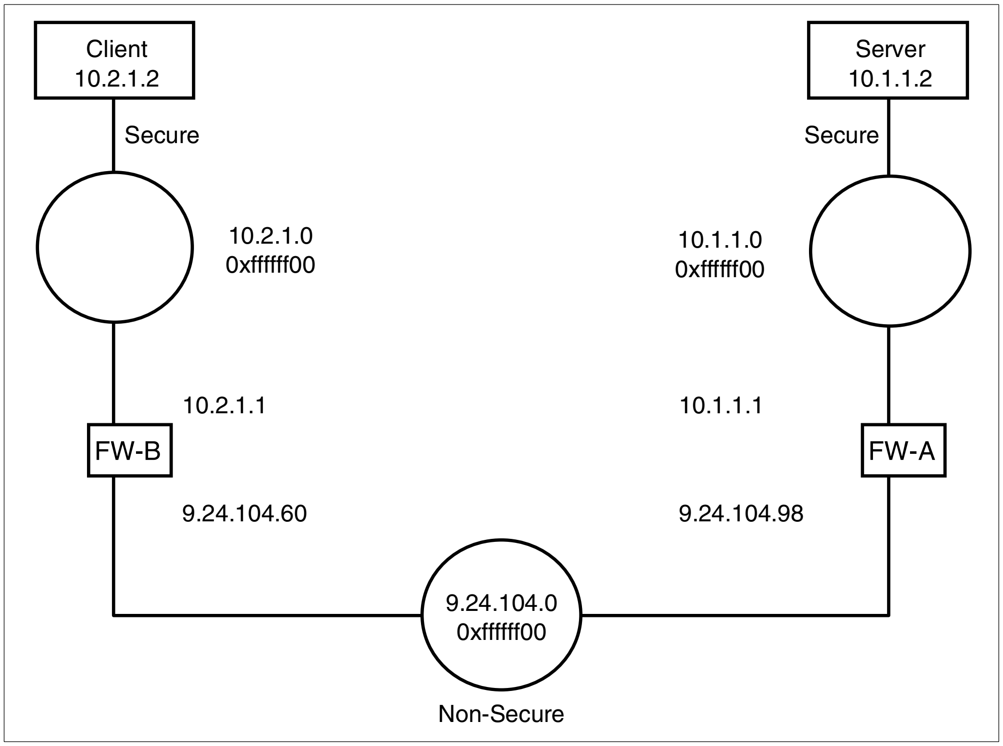

Figure 174. Network tunnel configuration between two IBM SecureWay Firewall V4.1 for AIX

We created the tunnel as described in 10.3, "Implementing the IPSec tunnel" on page 243, using authentication and encryption. See Figure 175 for the tunnel details.

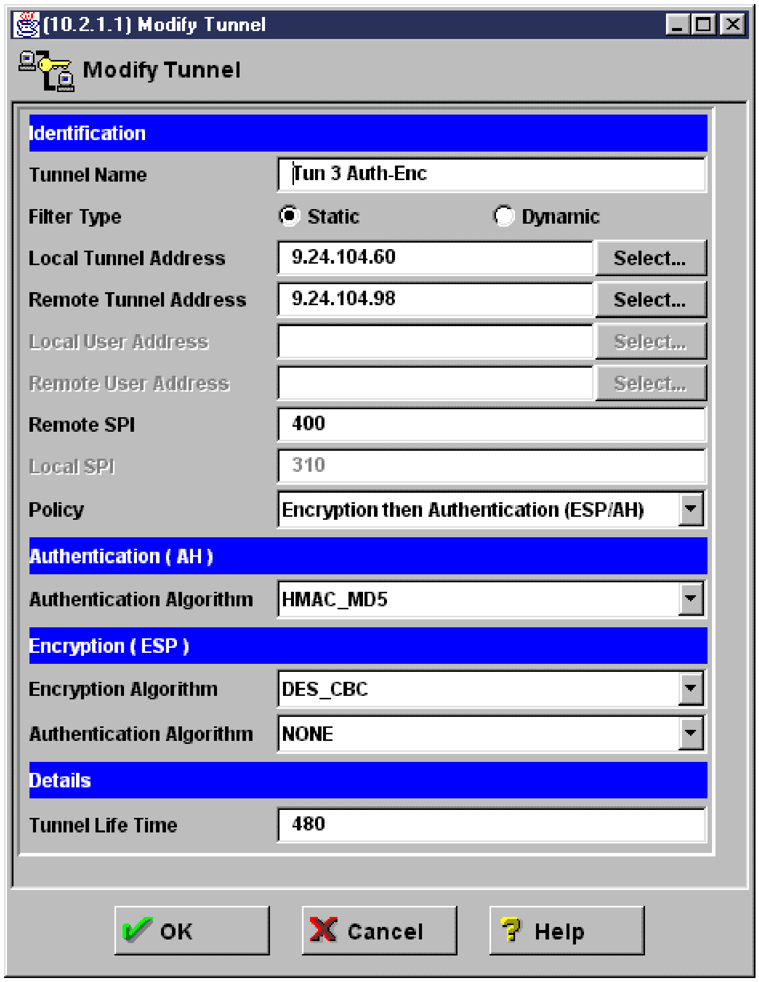

Figure 175. Tunnel characteristics

We created this tunnel in FW-B and exported it to FW-A, and we activated it on both ends.

The next step is adding the connections to allow the traffic in this tunnel. We first added the connection allowing the encapsulation traffic between both firewalls. For this connection, you can use a predefined service: "VPN encapsulation". See Figure 176.

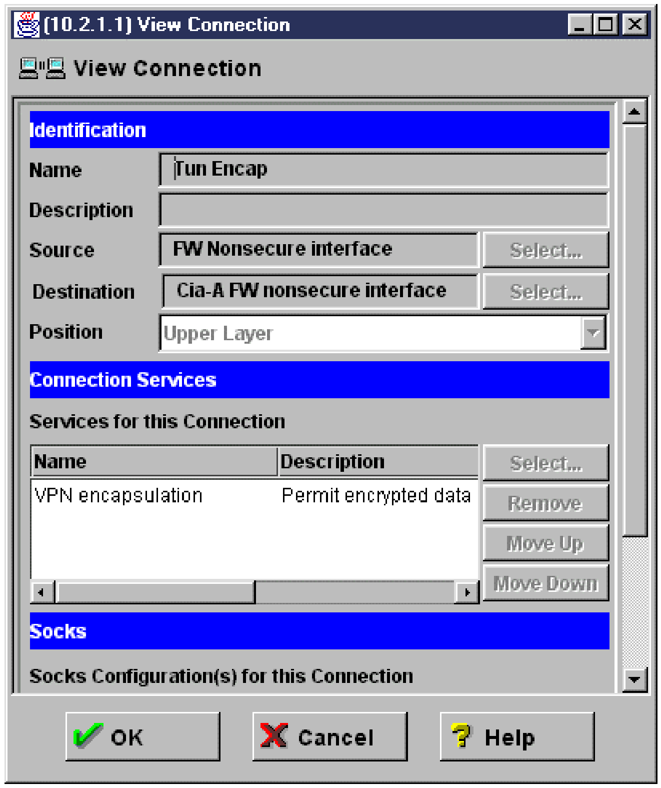

Figure 176. Connection to allow encapsulation traffic

We need to create the specific rules in both firewalls to allow only Telnet through this tunnel. Remember that the predefined services available in the firewall are generic services, so they contain rules that allow any traffic in the tunnel between the client and server machines.

The following diagram shows how you must create your rules to allow a certain service through the firewall using a VPN.

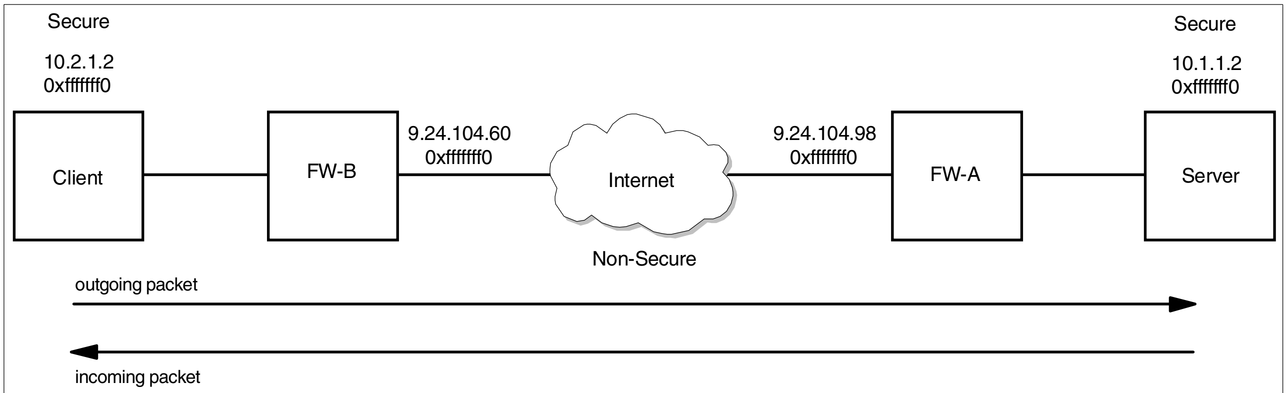

Outgoing packet:
permit <client IP address> <server IP address> <protocol> <src port> <dst port> secure route inbound T=0
permit <client IP address> <server IP address> <protocol> <src port> <dst port> nonsecure route outbound T=<number>

Incoming packet:
permit <server IP address> <client IP address> <protocol-response> <src port> <dst port> secure route outbound T=0
permit <server IP address> <client IP address> <protocol-response> <src port> <dst port> nonsecure route inbound T=<number>

Figure 177. Diagram for creating new rules for tunnel traffic

The arrows show the generic flow of packets without using tunnels, and the rules show how to split those arrows into specific rules. If the service you are configuring has more than one server port being used, you need to add two more rules for each flow, one rule using the secure interface and no tunnel, and the other using the non-secure interface and the tunnel you created.[2]

In the rules skeleton, we grouped them by the direction of the packet, so you can understand that you are going to create two rules for each packet that goes through the firewall using the tunnel. The `<protocol-response>` means that if you are using TCP, you use tcp/ack for this rule. The other protocols do not need any change.

The incoming packets are the response to the outgoing packets, so we can write the rules as shown in Figure 178, grouping them by pairs of outgoing and incoming packets.

First pair, using no tunnel:
permit <client IP address> <server IP address> <protocol> <src port> <dst port> secure route inbound T=0
permit <server IP address> <client IP address> <protocol-response> <src port> <dst port> secure route outbound T=0

Second pair, using the specific tunnel:
permit <client IP address> <server IP address> <protocol> <src port> <dst port> nonsecure route outbound T=<number>
permit <server IP address> <client IP address> <protocol-response> <src port> <dst port> nonsecure route inbound T=<number>

Figure 178. Example of rules for tunnels

In our example, we are going to allow only Telnet through that tunnel. So our rules in FW-B are:

[2] Optionally you may want to add the ICMP filters as shown in Figure 162., “Generated dynamic filter rules” on page 255

```
Telnet tunnel 1/4:permit tcp gt 1023 eq 23 secure route inbound T=0
Telnet tunnel 2/4:permit tcp/ack eq 23 gt 1023 secure route outbound T=0
Telnet tunnel 3/4:permit tcp gt 1023 eq 23 nonsecure route inbound T=3
Telnet tunnel 4/4:permit tcp/ack eq 23 gt 1023 nonsecure route outbound T=3
```

After adding the rules, we added a service containing all rules, as shown in Figure 179:

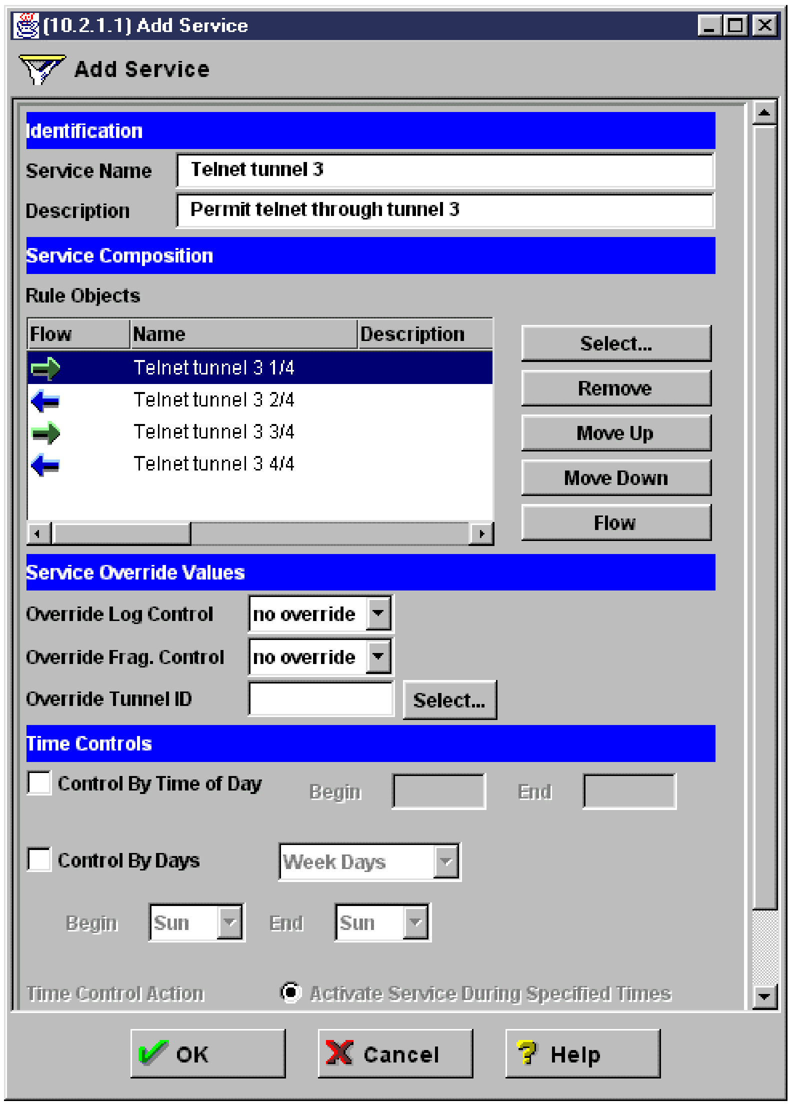

Figure 179. Service "Telnet Tunnel 3"

If you do not want to add the specific tunnel ID in the rules (so you can use the same rules for other tunnels), you can create two separate services; the first one containing the first two rules, and the second one containing the other two rules (this one should not have the tunnel ID in this case). Then you select to override the tunnel ID number in the second service.

Finally, we added a connection (as shown in Figure 180) and activated it.

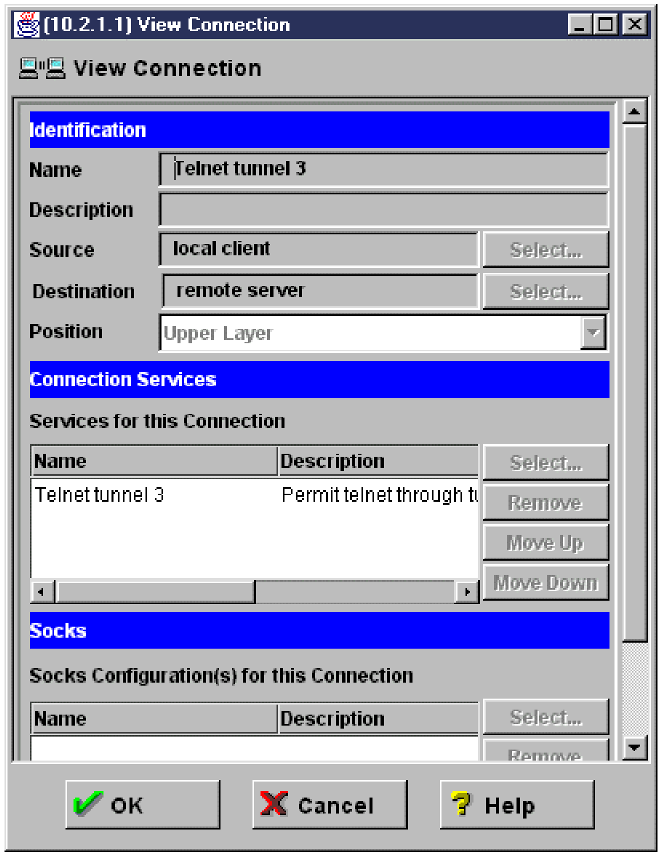

Figure 180. Connection allowing Telnet through tunnel 3

Figure 181 shows the rules listing generated by the two connections we added to allow traffic in this tunnel. Note that we turned on the logging for the rules in "Telnet tunnel 3" to have more information for debugging.

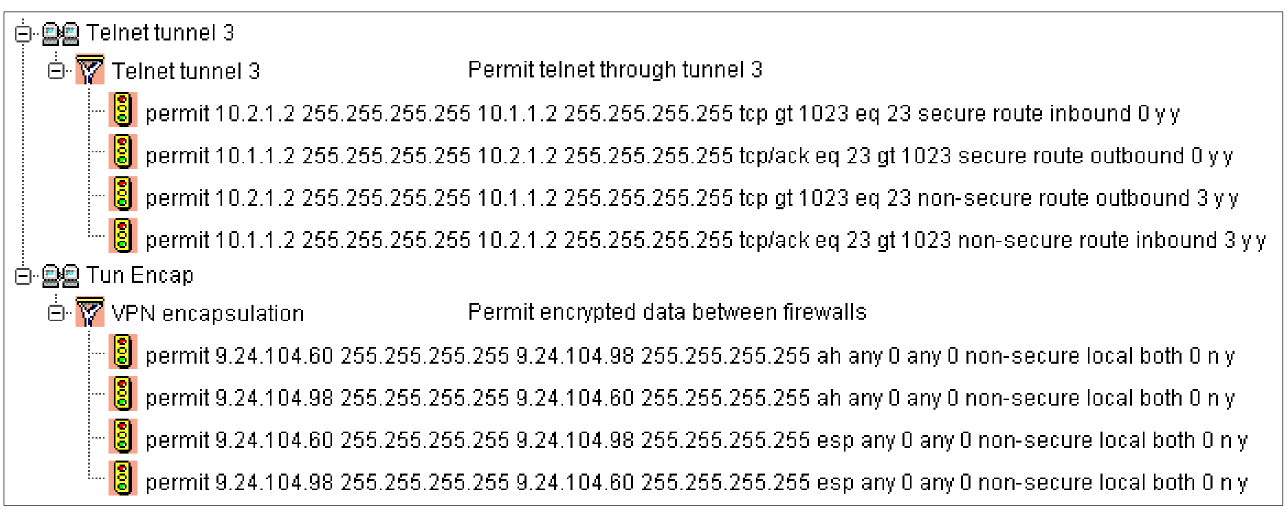

Figure 181. Rules for FW-B

After finishing the configuration in firewall FW-B, we needed to repeat the same steps in firewall FW-A. You just need to add the exact same rules, service and connections. Figure 182 shows the rules listing in firewall FW-A after doing all steps.

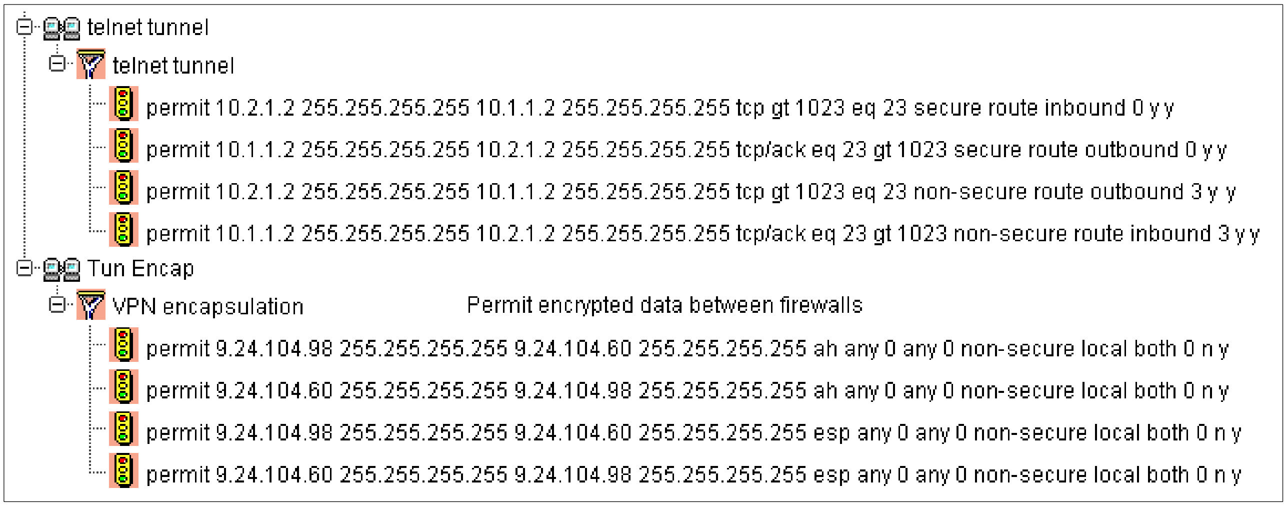

Figure 182. Rules for FW-A

Once we did all these steps, and activated the connections, we can test the tunnel by telneting from the client 10.2.1.2 to the server 10.1.1.2.

The following log shows the entries for the test we did:

```
Sep 14 19:08:32 1999 b-gateway: ICA1075i:ft:1 tid:503 sid:0 #:1 R:p i 10.2.1.1 s
:10.2.1.2 d:10.1.1.2 p:tcp sp:35865 dp:23 r:r a:s f:n T:0 l:44
Sep 14 19:08:32 1999 b-gateway: ICA1075i:ft:1 tid:503 sid:0 #:3 R:p o 9.24.104.6
0 s:10.2.1.2 d:10.1.1.2 p:tcp sp:35865 dp:23 r:r a:n f:n T:3 l:44
Sep 14 19:08:32 1999 b-gateway: ICA1075i:ft:5 tid:514 sid:0 #:1 R:p i 9.24.104.6
0 s:9.24.104.98 d:9.24.104.60 p:ah -:0 -:0 r:l a:n f:n T:0 l:108
Sep 14 19:08:32 1999 b-gateway: ICA1075i:ft:1 tid:503 sid:0 #:4 R:p i 9.24.104.6
0 s:10.1.1.2 d:10.2.1.2 p:tcp sp:23 dp:35865 r:r a:n f:n T:3 l:44
Sep 14 19:08:32 1999 b-gateway: ICA1075i:ft:1 tid:503 sid:0 #:2 R:p o 10.2.1.1 s
:10.1.1.2 d:10.2.1.2 p:tcp sp:23 dp:35865 r:r a:s f:n T:0 l:44
Sep 14 19:08:32 1999 b-gateway: ICA1075i:ft:1 tid:503 sid:0 #:1 R:p i 10.2.1.1 s
:10.2.1.2 d:10.1.1.2 p:tcp sp:35865 dp:23 r:r a:s f:n T:0 l:40
Sep 14 19:08:32 1999 b-gateway: ICA1075i:ft:1 tid:503 sid:0 #:3 R:p o 9.24.104.6
0 s:10.2.1.2 d:10.1.1.2 p:tcp sp:35865 dp:23 r:r a:n f:n T:3 l:40
Sep 14 19:08:32 1999 b-gateway: ICA1075i:ft:1 tid:503 sid:0 #:1 R:p i 10.2.1.1 s
:10.2.1.2 d:10.1.1.2 p:tcp sp:35865 dp:23 r:r a:s f:n T:0 l:55
Sep 14 19:08:32 1999 b-gateway: ICA1075i:ft:1 tid:503 sid:0 #:3 R:p o 9.24.104.6
0 s:10.2.1.2 d:10.1.1.2 p:tcp sp:35865 dp:23 r:r a:n f:n T:3 l:55
Sep 14 19:08:32 1999 b-gateway: ICA1075i:ft:5 tid:514 sid:0 #:1 R:p i 9.24.104.6
0 s:9.24.104.98 d:9.24.104.60 p:ah -:0 -:0 r:l a:n f:n T:0 l:108
Sep 14 19:08:32 1999 b-gateway: ICA1075i:ft:1 tid:503 sid:0 #:4 R:p i 9.24.104.6
0 s:10.1.1.2 d:10.2.1.2 p:tcp sp:23 dp:35865 r:r a:n f:n T:3 l:40
Sep 14 19:08:32 1999 b-gateway: ICA1075i:ft:1 tid:503 sid:0 #:2 R:p o 10.2.1.1 s
:10.1.1.2 d:10.2.1.2 p:tcp sp:23 dp:35865 r:r a:s f:n T:0 l:40
```

10.5.1.1 Using NAT with IPSec

NAT's purpose is to shield the IP addresses on the secure side of the firewall from the non-secure side (see Chapter 9, “Network Address Translation” on page 221). This solves two problems:

- It allows you to use unregistered addresses in your secure network and still access the non-secure network without conflict.
- It also keeps the non-secure network hosts from knowing about any of your secure-side host IP addresses.

To do this, NAT has to alter the source IP address of outgoing packets. When response packets come inbound on the connection, NAT reverses the translation it performed when the packet was outbound and resets the proper secure host's IP address in the packet destination fields. NAT does this even for packets that go through tunnels as we see in Figure 183.

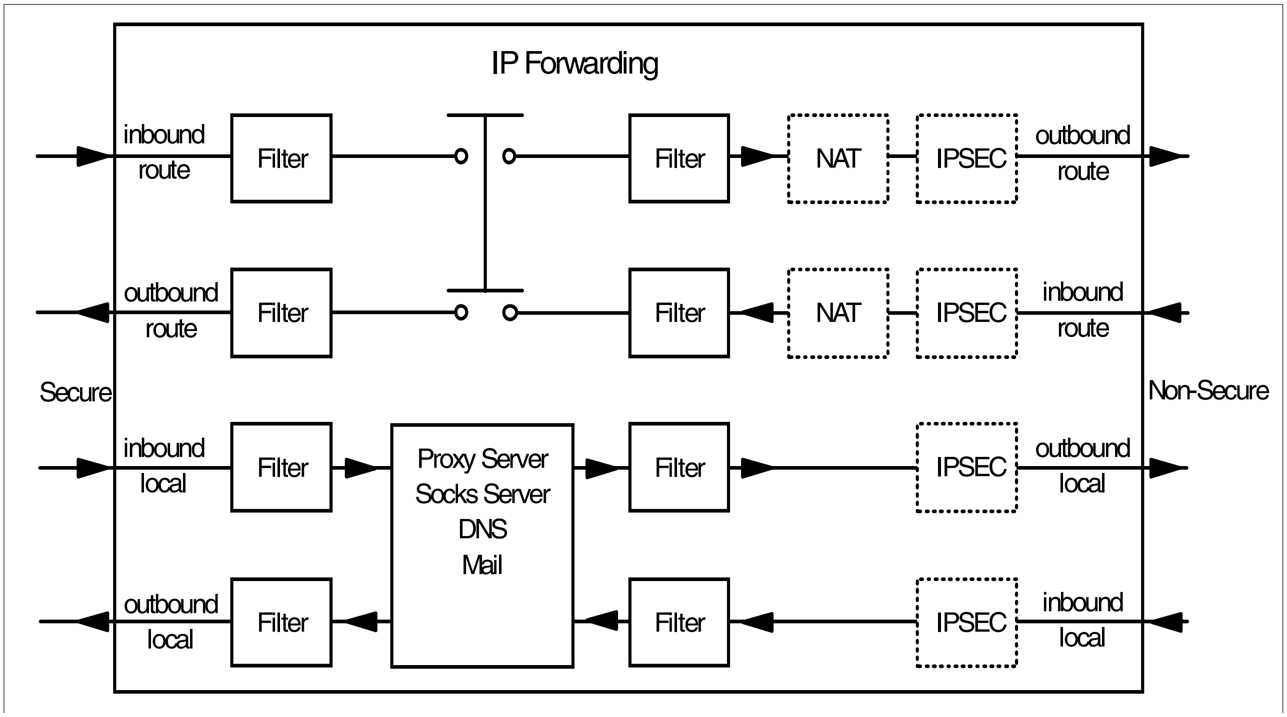

Figure 183. Filters, NAT and IPSec in the IBM SecureWay Firewall V4.1 for AIX

Note:

- When we use the proxy server together with VPN, the first connection from the secure network to the secure interface of the firewall will be the FTP proxy 1/2. The tunnel data traffic connection will be from the non-secure interface of the firewall to the remote secure network or host.
- We can create tunnels in two different ways: with dynamic or with static filter rules. When using NAT on the owning firewall we *cannot* simply use the dynamic filter rules. The reason is that dynamic filter rules are created automatically using the secure IP addresses for the traffic rules (see Figure 162 on page 255). When exporting and importing the tunnel definitions on the partner tunnel, the source and destination addresses will be reversed, but for the partner end the remote secure address is unknown because of NAT. On the partner end we *must* use the NAT address in the filter rule instead of the automatically created filter rule with the remote secure address. We could edit the export file and change the filter rules, but we think it is safer to use static filter rules.

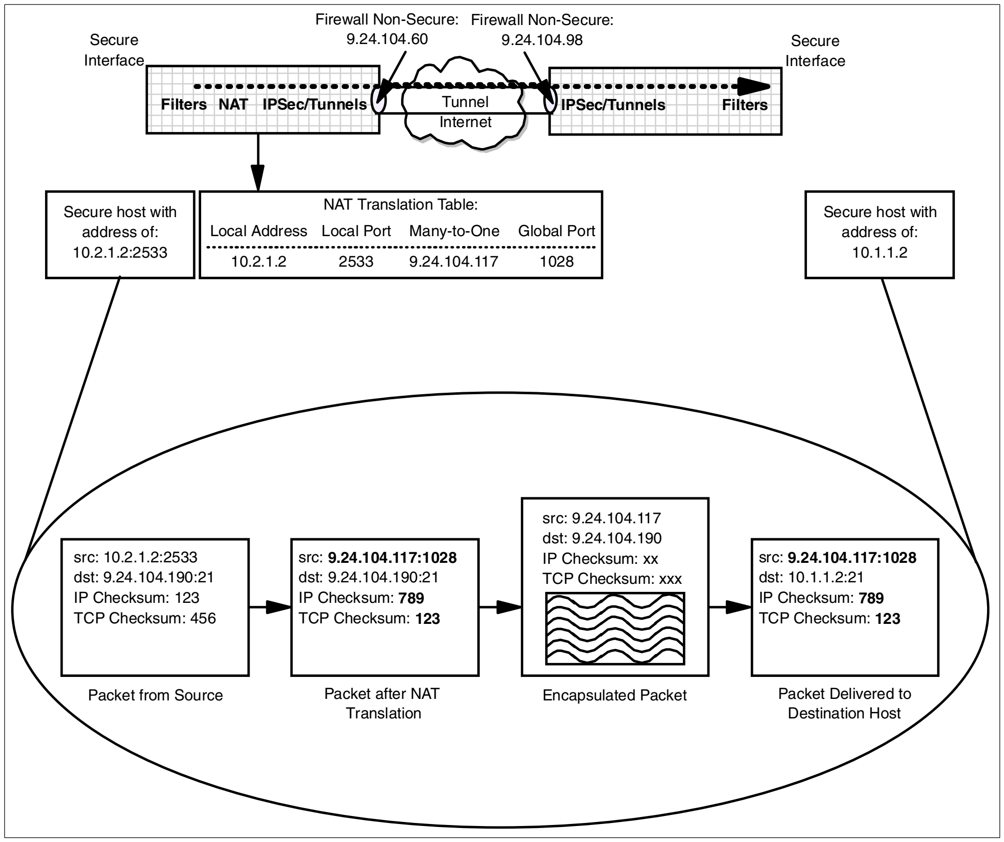

Figure 184. Interaction between NAT, filters and tunnels

Figure 184 shows us the 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. 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 partner's end of the tunnel (assuming that 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.

We are still using the configuration as shown in Figure 174 on page 267 and we have activated NAT only on FW-B with NAT address 9.24.104.190. As a result the secure network 10.2.1.0 is no longer known to the outside world. Therefore we need to modify the filter rules on FW-A, and use the NAT address 9.24.104.117 instead of the remote server’s secure IP address 10.1.1.2.

Figure 185 shows the new rules for the firewall FW-B, using the NAT address 9.24.104.190 for the destination address.

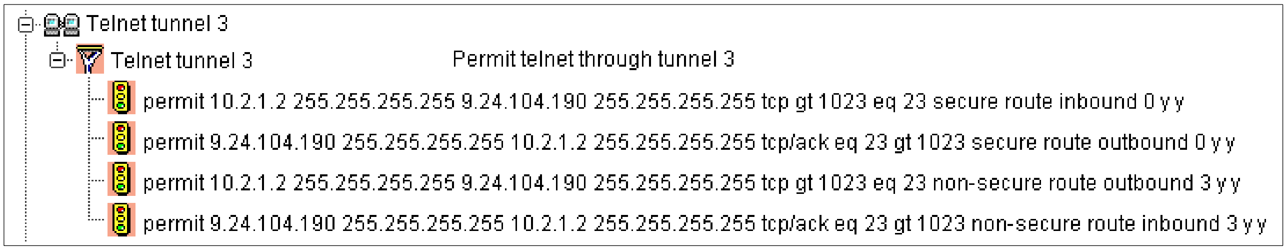

Figure 185. Filter rules on FW-B using NAT on FW-A

Figure 186 shows the rules for FW-A. This time, the source address is the NAT address (9.24.104.117), while the destination address is the actual address (10.1.1.2).

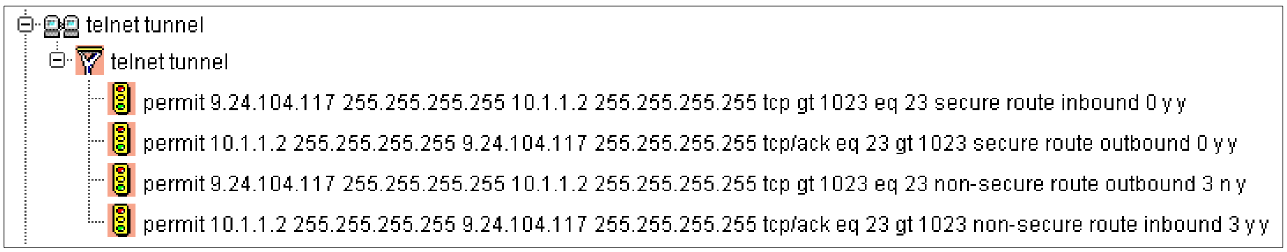

Figure 186. Filter rules on FW-A using NAT on FW-B

10.5.2 VPN between IBM SecureWay Firewall V4.1 for AIX and IBM eNetwork Firewall for Windows NT

In this test we used the IBM eNetwork Firewall for Windows NT V3.3, but we could have also used the latest version of this firewall, the IBM SecureWay Firewall for Windows NT V4.1.

In this scenario we created a tunnel with dynamic rules in an IBM SecureWay Firewall V4.1 for AIX machine and exported it to the IBM eNetwork Firewall for Windows NT. The AIX firewall has a non-secure adapter with IP address 9.24.104.60 and the NT firewall has a non-secure adapter with IP address 9.24.104.117.

We created a dynamic tunnel in the local firewall using the options as shown in Figure 187:

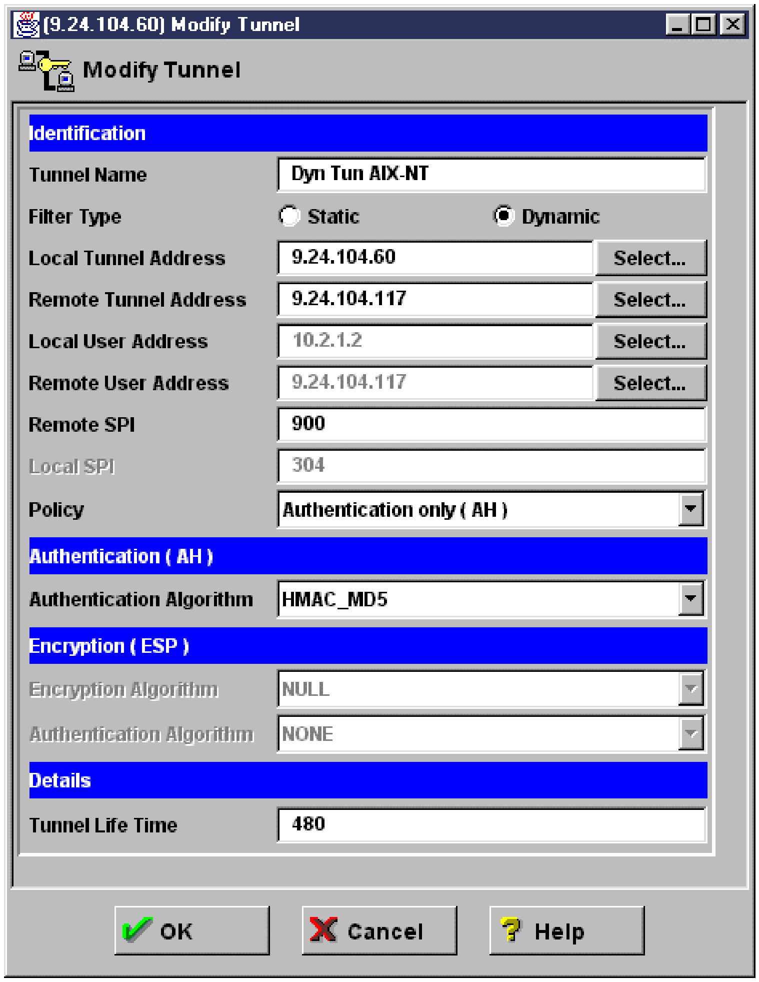

Figure 187. Tunnel definition in firewall for AIX

We exported the file, and imported it in the IBM Firewall for Windows NT machine. We did not need to convert or edit the tunnel file. Figure 188 shows the imported tunnel in IBM Firewall for Windows NT:

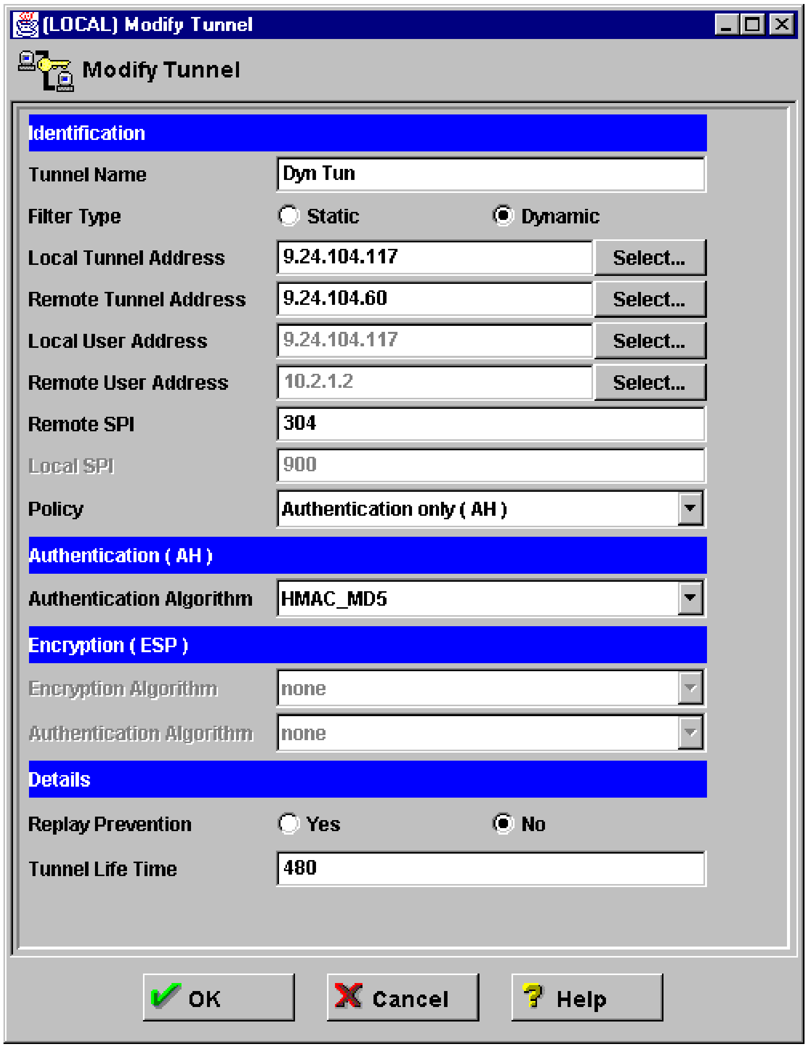

Figure 188. Tunnel definition in IBM Firewall for Windows NT after import

We activated the tunnel on both ends, and the following dynamic rules were automatically started at the firewall 9.24.104.60:

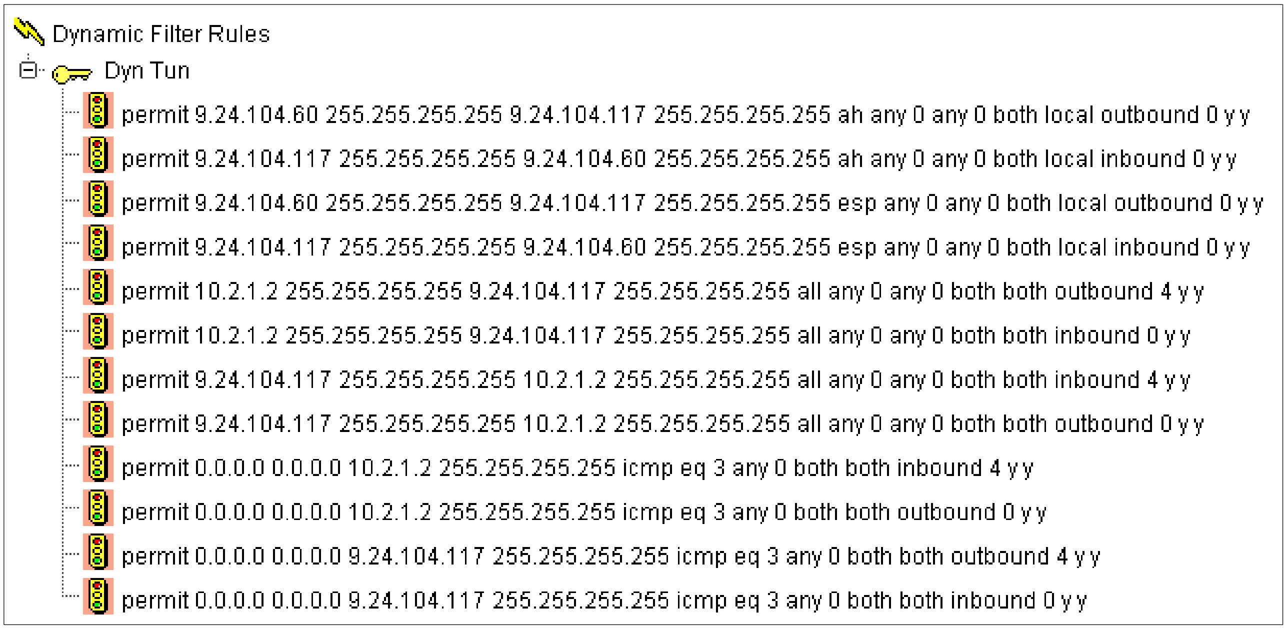

Figure 189. Dynamic rules

We did a test using ping from the machine 9.24.104.117 to the internal server 10.2.1.2. See the log entries that show the ping flowing through the tunnel:

```
Sep 13 18:55:09 1999 b-gateway: ICA1075i:ft:2 tid:4 sid:0 #:2 R:p i 9.24.104.60
s:9.24.104.117 d:9.24.104.60 p:ah -:0 -:0 r:l a:n f:n T:0 l:104
Sep 13 18:55:09 1999 b-gateway: ICA1075i:ft:2 tid:4 sid:0 #:7 R:p i 9.24.104.60
s:9.24.104.117 d:10.2.1.2 p:icmp t:8 c:0 r:r a:n f:n T:4 l:60
Sep 13 18:55:09 1999 b-gateway: ICA1075i:ft:2 tid:4 sid:0 #:8 R:p o 10.2.1.1 s:9
.24.104.117 d:10.2.1.2 p:icmp t:8 c:0 r:r a:s f:n T:0 l:60
Sep 13 18:55:09 1999 b-gateway: ICA1075i:ft:2 tid:4 sid:0 #:6 R:p i 10.2.1.1 s:1
0.2.1.2 d:9.24.104.117 p:icmp t:0 c:0 r:r a:s f:n T:0 l:60
Sep 13 18:55:09 1999 b-gateway: ICA1075i:ft:2 tid:4 sid:0 #:5 R:p o 9.24.104.60
s:10.2.1.2 d:9.24.104.117 p:icmp t:0 c:0 r:r a:n f:n T:4 l:60
```

10.5.3 VPN between IBM SecureWay Firewall V4.1 for AIX and the AIX V4.3 operating system

The next scenario is a tunnel between the IBM SecureWay Firewall V4.1 for AIX and the AIX 4.3.2 operating system. The tunnel is configured between the firewall with the non-secure interface (9.24.104.60) and the AIX interface (9.24.106.33). The tunnel is going to be established between these two machines (we are not using separate clients and servers). Figure 190 shows the configuration for this scenario.

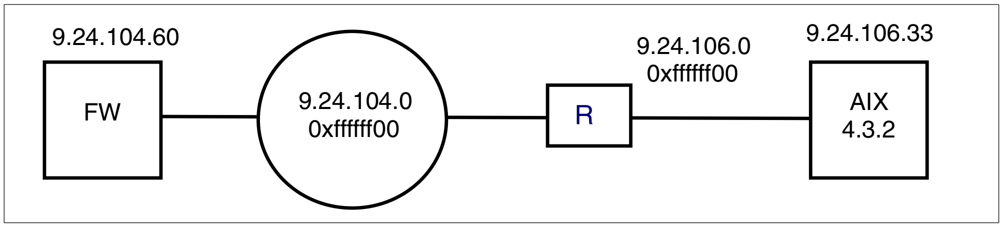

Figure 190. Network tunnel configuration between IBM SecureWay Firewall V4.1 for AIX and AIX 4.3

We executed the following steps:

- On the firewall we have created a tunnel with dynamic filter rules. This means the filter rules are put in memory and activated when we activate the tunnel (see also 10.5.1, “Tunnel between two IBM SecureWay Firewall V4.1 for AIX” on page 266).

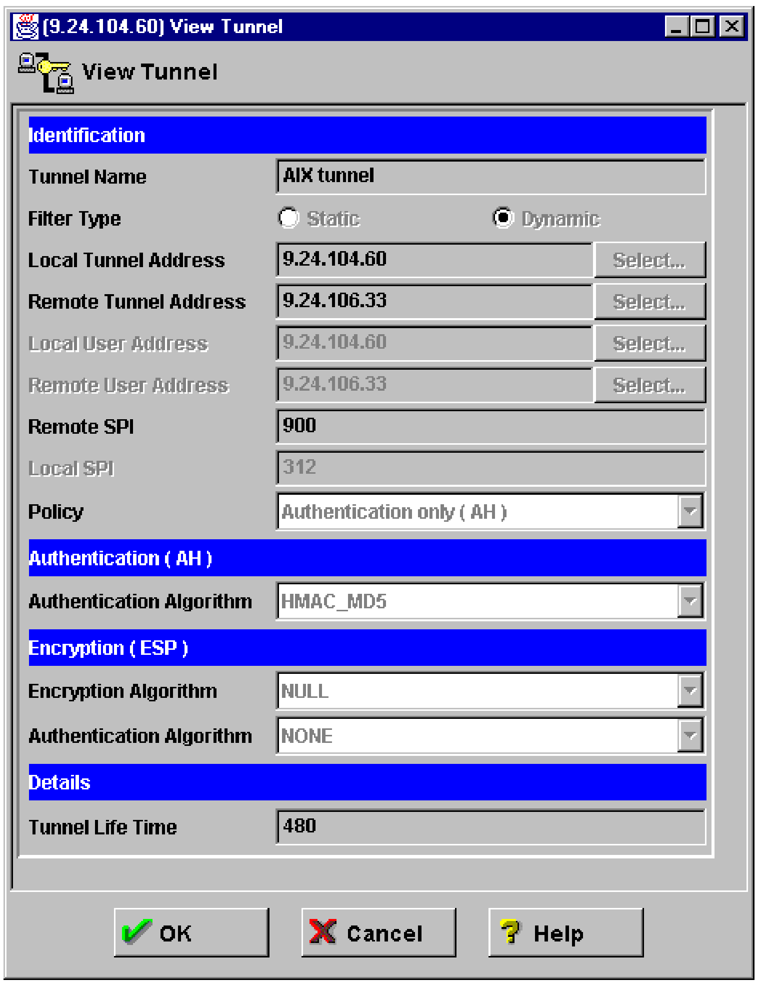

Figure 191. Tunnel definition between IBM SecureWay Firewall V4.1 for AIX and AIX 4.3

- We export the tunnel so the ipsec_tun_man.exp file is created.

```
#-------------------------------------
4
9.24.104.60
9.24.106.33
5
900
900
312
312
NULL
0
0x
NULL
0
0x
HMAC_MD5
16
0xA0CDA1F218E92567D4B474A1C2014F24
HMAC_MD5
16
0x928C8D1189C01E5330636E406589C41A
0
28800
tunnel
tunnel
axax
0
1
NONE
0

NONE
0

0
-
-
AIX tunnel
1
9.24.104.60
255.255.255.255
9.24.106.33
255.255.255.255
```

Figure 192. Export file

- Before importing the file on the AIX we need to convert the file format of the export file to an AIX file format. The `conv_export_file` utility will execute the conversion for us. The command syntax is:

  `conv_export_file dir=ddddd`

  where `dir=ddddd` specifies the directory of the location of the export file to be converted, as shown in the following:

```
# conv_export_file dir=/export
Command completed successfully.
```

Conversion

The conversion utility comes with the IBM SecureWay Firewall V4.1 for AIX and runs on AIX. The `conv_export_file` utility can only be used for a tunnel between the IBM SecureWay Firewall V4.1 for AIX and AIX. The converter utility will modify the existing file. It does not create a new file.

```
#-------------------------------------
4
9.24.104.60
9.24.106.33
5
900
900
312
312
NULL
0
0x
NULL
0
0x
HMAC_MD5
16
0xA0CDA1F218E92567D4B474A1C2014F24
HMAC_MD5
16
0x928C8D1189C01E5330636E406589C41A
0
28800
tunnel
tunnel
axax
0
1
NONE
0

NONE
0

0
-
-
```

Figure 193. Converted export file

The only difference we noticed between the original file and the converted file is the deletion of the filter lines at the bottom of the file.

When the export file is successfully converted we can import the tunnel on the AIX server. To do it, follow the next steps:

1. Take the converted tunnel file to the AIX machine.
2. Run `smit`, and select the following menus to start IP Security on this machine: **Communications Applications and Services -> TCP/IP -> Configure IP Security (IPv4) -> Start/Stop IP Security**.
3. Go to the Basic IP Security configuration menu and import the tunnel file.

4. Check the filter rules (see Figure 194).
5. Activate the rules.
6. Activate the tunnel.

```
1 *** Dynamic filter placement rule for IKE tunnels *** no
2 permit 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 yes ah any 0 any 0 both both no all packets 0 all
3 permit 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 yes esp any 0 any 0 both both no all packets 0 all
4 permit 9.24.104.60 255.255.255.255 9.24.106.33 255.255.255.255 yes all any 0 any 0 both inbound no all
packets 1 all
5 permit 9.24.106.33 255.255.255.255 9.24.104.60 255.255.255.255 yes all any 0 any 0 both outbound no all
packets 1 all
6 permit 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 yes all any 0 any 0 both both yes all packets 0 all
```

Figure 194. AIX tunnel rules

The AIX filter rules syntax is a little different from those in the firewall; there are no secure or non-secure interfaces. Note that:

1. Title 1 and rule 6: Two default rules automatically created by the system. Rule 6 opens the system to any traffic; this is not a firewall, so there is no need to restrict the traffic, but you may change it if you want.
2. Rules 2 and 3: Automatically created when the IP Security is started on the system.
3. Rules 4 and 5: Automatically created when importing the export file. The number 1 refers to the tunnel ID.

In our example, we connected from the firewall machine to AIX. If you are going to connect from a client behind the firewall to AIX, you must add filters in AIX allowing this connection though the respective tunnel.

Chapter 11. Logging, monitoring, and reporting

Logging is essential to the day-to-day operation of the IBM SecureWay Firewall V4.1 for AIX. Unless you log the activity on your firewall and generate alerts for suspicious activity, you could be under attack without even realizing it. Worse, in the event of an attack, you would be seriously hampered in your attempts to determine the origin and target of the attack.

This chapter describes how to configure the logging facilities, how to monitor the logging of alerts in real time, and how to build useful reports on top of the logged data. The firewall monitors the messages sent to its 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.

11.1 Configure logging

It is very important to configure the firewall to log the information you need. If you get too much information on the logs you may overlook important data.

We are going to show some examples, but it is important that you configure the logging according to your needs and your environment.

11.1.1 Logging priority levels

Let's start with a quick look at *syslog* to understand the basics of logging. Syslog is a daemon common to the UNIX environment. It is used to centralize the logging activities from different applications. AIX provide eight *priority levels* to specify the amount of logging activity for a specific service. They are:

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

At the *debug* level all activity for that facility is logged. At the *emergency* level, very little activity (only severe messages) would be logged at all.

On AIX, the syslog activity is handled by the daemon syslogd. The configuration file for this daemon is /etc/security/syslog.conf.

You should, however, use the configuration client to configure the log facilities, because it takes archive settings into account and refreshes the process for you.

The firewall only uses only five of the available priority levels, which are shown in Table 29.

Table 29. Firewall log priority levels

Level	Description
debug	All messages are logged.

Level	Description
information	Only messages with the priority levels information, warning, error, and critical are logged.
warning	Only messages with the priority levels warning, error, and critical are logged.
error	Only messages with the priority levels error and critical are logged.
critical	Only messages with the priority level critical are logged.

11.1.2 Log facilities

The log facility determines the type and source of information that is logged. The firewall uses the log facilities shown in Table 30.

Table 30. Firewall log facilities

Facility	Internal Name	Description
Firewall Log	local4	General firewall log, including IP filter logging, SOCKS and proxy usage, and mail events.
Alert Log	local1	Log monitor threshold violation warnings. The messages displayed in the Alert Display (see Figure 210 on page 304) are taken from this facility. You get messages in the Alert Display only if you create a alert log file.
Audit Log	local0	All firewall administrator functions are logged here.

11.1.3 Manage log facilities

You can manage the three log facilities from Table 30 with the configuration client. The Audit Log facility can be configured and browsed only from the command line (see 11.1.5, “Manage the audit log facility” on page 292).

To create new log facilities or modify existing ones select **System Administration -> System Logs -> Log Facilities** from the configuration client navigation tree.

The window in Figure 195 is displayed.

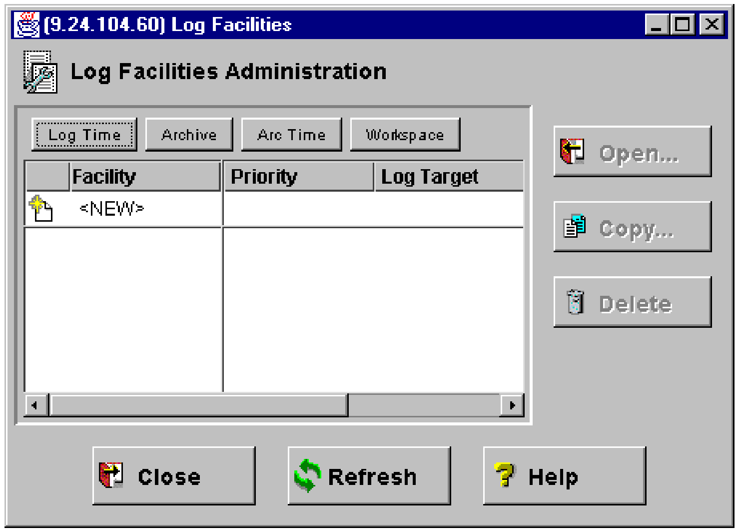

Figure 195. Log facilities administration

To create a new log facility to monitor activity on the firewall, double-click **<NEW>** in Figure 195. The window shown in Figure 196 is displayed.

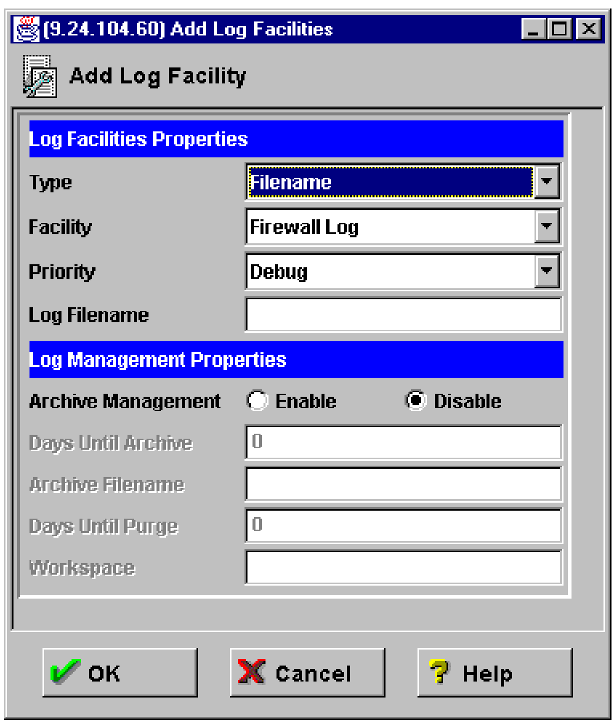

Figure 196. Add alert log facility

This window provides two main panels. The first one lets you set up the basic options for logging. The second one is related to archiving log files. These options are explained in 11.1.4, “Archive log files” on page 289.

You have the following options to specify the log facility properties:

Type It determines how syslog will handle the log entries. You can choose **Filename** (it saves the entries to a file), **Hostname** (it sends the entries to other machine) or **User ID** (it sends the entries to a user).

Facility The log facility determines the type and source of information that is logged (see 11.1.2, “Log facilities” on page 286). You cannot select the Audit Log facility. This facility can be created only via the command line (see 11.1.5, “Manage the audit log facility” on page 292).

Priority Specifies the log priority level for the selected facility. The log priority levels are listed in order of increasing severity (see 11.1.1, “Logging priority levels” on page 285). The priority you select will be the minimum level that gets logged.

Log Filename Fill in the log filename. The log filename must have an absolute path (beginning with the drive). The path to the file must exist.

We decided to log at the debug priority level for the following reasons:

- We have plenty of disk space.
- We will archive regularly to efficiently use the space available.
- We need to capture everything while we test the firewall.

How much should you log?

From the time you set up the firewall to the time you finish the setup we advise you to log at the debug level or you risk missing vital information. After successful completing the testing you could decrease the log priority level.

We repeated the above steps for the Alert Log log facility.

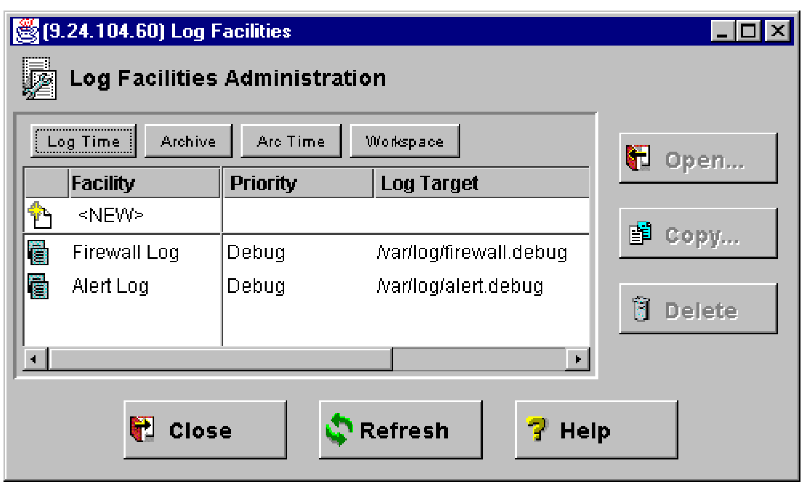

Figure 197. Completed log facilities

As you can see, we also selected the debug level of priority for the Firewall Log facility.

The configuration file for log facility properties is /etc/security/syslog.conf. The following example shows the syslog.conf file for the log facilities.

```
local4.debug /var/log/firewall.debug
local1.debug /var/log/alert.debug
```

The terms `local1` and `local4` in this file are coming from the syslog facilities. They are user-definable log facilities. Each line has the following format:

```
log_facility.log_priority log_file
```

Where `log_facility` is the Log Facility type, `log_priority` is a valid log priority, and `log_file` is the file name of the log file.

11.1.4 Archive log files

The firewall can also manage the size of your log files, which can increase heavily over time. The size of your log files depends on:

- What priority level you are logging.
- Which filter rules are being logged.
- The amount of traffic going through your firewall.

Therefore, we highly recommend that you use the log archive capabilities of the firewall. The archival process removes qualifying records from an active log file, places them in a separate file, compacts the resulting file, and places the new file into an archive directory.

Configuring the archive management involves two steps:

1. Enable archive management.
2. Add commands to the AIX crontab for periodic execution.

11.1.4.1 Enable archive management

To enable the archive management, double-click on one of your defined logging facilities in Figure 197.

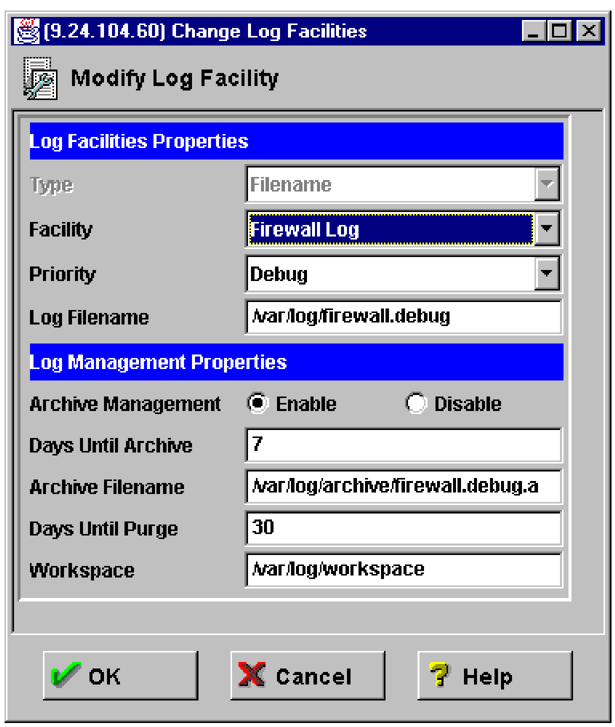

Figure 198. Modify archive properties

Now focus on the Log Management Properties. They give you the following options:

Archive Management	When enabled, the log file will be processed according to the settings described below.
Days Until Archive	Specifies the number of days until the log files are archived. The number of days until archive must be zero or greater.
Archive Filename	Specifies the name of the file where the archived data will be written. An absolute path name must be specified.
Days Until Purge	Specifies the number of days until the log files are purged. The number of days to keep the log files must be zero or greater. Log management does not count the current day when calculating the number of days to keep the file.

We enabled archive management and set our log to archive every seven days, clearing after 30 days. If you find disk space a problem, daily archiving can help.

The configuration file for the log management is /etc/security/logmgmt.cfg.

This file is filled out by the configuration client. You can check here to make sure your archiving is set right. The following is the logmgmt.cfg file for the log facilities:

```
var/log/firewall.debug 7 /var/log/archive/firewall.debug.a 30
/var/log/workspace
/var/log/alert.debug 7 /var/log/archive/alert.debug.a 30 /var/log/workspace
```

Each line has the following format:

```
log_file days_in_log_file archive_file days_in_archive
```

Where `log_file_name` is the name of the log file, `archive_file` is the name of the compressed archive file, `days_in_log_file` is the number of days until log entries are archived, and `days_in_archive` is the number of days until entries are purged from the archive file.

You can also review your settings with the `fwlog` command. Figure 199 shows an example output.

```
# fwlog cmd=list
1       facility    =   local4
        priority    =   debug
        logfile     =   /var/log/firewall.debug
        logtime     =   7
        arcfile     =   /var/log/archive/firewall.debug.a
        arctime     =   30
        workspace   =   /var/log/workspace

2       facility    =   local1
        priority    =   debug
        logfile     =   /var/log/alert.debug
        logtime     =   7
        arcfile     =   /var/log/archive/alert.debug.a
        arctime     =   30
        workspace   =   /var/log/workspace
```

Figure 199. Log file properties using the fwlog cmd=list command

11.1.4.2 Schedule the archiving

While updating the three log facilities you will probably see a panel similar to Figure 200.

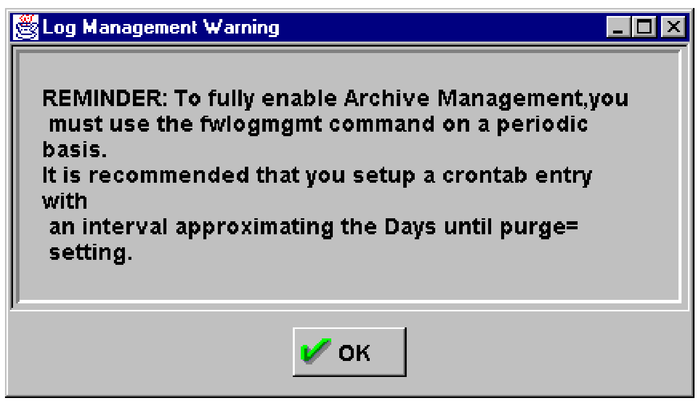

Figure 200. Archiving warning message

Archive management also requires that the `fwlogmgmt` command be submitted on a periodic basis.

Unfortunately it will not stop the message in Figure 200 from being displayed, but you will be able to safely ignore it.

11.1.4.3 Using AIX crontab

To complete configuration of archiving you need to add two commands to the AIX crontab for periodic execution.

```
0 3 * * * /usr/bin/fwlogmgmt -l
0 4 * * * /usr/bin/fwlogmgmt -a
```

The first command will archive your logs every day of the week at three o'clock in the morning and the second will purge entries over the age you specified earlier, every day at four o'clock in the morning. Log file archiving is a processor-and disk-consuming task, so be sure you choose a time frame where your firewall will probably not be very busy. Also choose different start times for both commands.

You can also use both log management commands from the command line, but we recommend you configure them as scheduled services. This decreases the chances of lost logging entries due to a full hard disk.

11.1.5 Manage the audit log facility

Audit log files must be managed from the command line on a local machine. This is because all log files that are configured with the configuration client can also be browsed with it. The audit log file contains sensitive administrative data and should therefore be browsed only by the primary firewall administrator.

To enable this log, first edit the file /etc/syslog.conf and add the following line:

```
local0.debug /var/log/audit.debug
```

Now you need to create the file `/var/log/audit.debug` and refresh the syslogd daemon. Run the following commands:

```
# touch /var/log/audit.debug
# refresh -s syslogd
0513-095 The request for subsystem refresh was completed successfully.
```

11.1.6 Examine the firewall log files

You can use the firewall Log Viewer to display your firewall log files. To open the Log Viewer click the **Log Viewer** button at the bottom right of the configuration client. Now select one of your defined log files with the File Name pull-down menu.

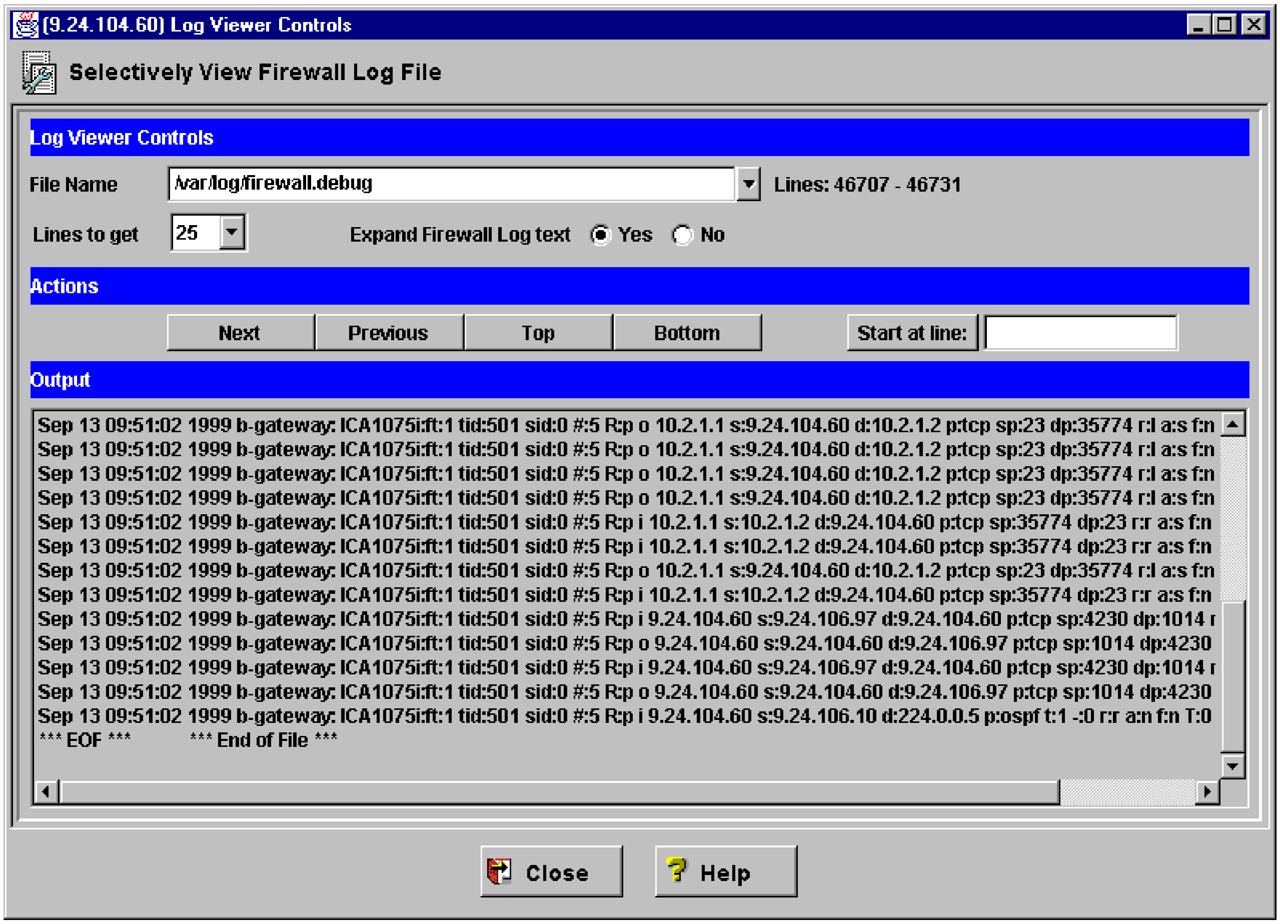

Figure 201. Log viewer

You can navigate in the log file by clicking the **Next**, **Previous**, **Top**, and **Bottom** buttons. You can change the number of lines displayed at one time by changing the value in the Lines to get scroll box. The Start at line: box allows you to jump to a specific line given in the entry field behind the button. The Expand Firewall Log Text buttons allow you to change between a log output with and without textual messages.

Examining large log files

Depending on the speed and memory of your machine, examining large log files with the Log Viewer can be a boring task. So do not give up waiting if you do not receive a response immediately.

11.1.7 Configure logging sources

Some of the firewall services do not log by default. Therefore you have to manually enable the logging for these services if you want to receive more information.

11.1.7.1 Connection rules

Before any denied or permitted IP packets are logged you have to enable connection rules logging with the Connection Activation panel. To open the panel double-click **Traffic Control -> Connection Activation** in the navigation tree.

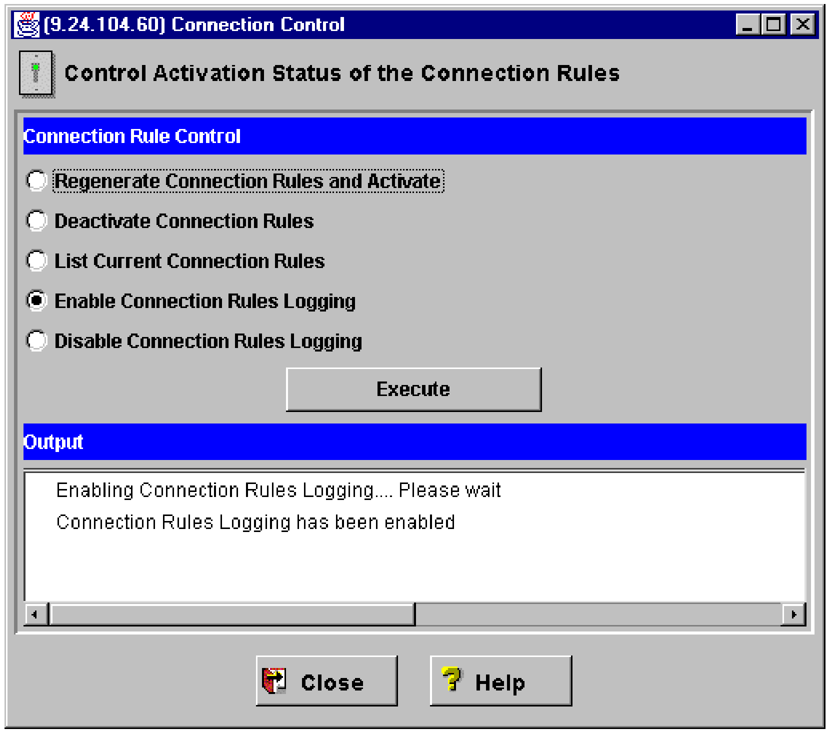

Figure 202. Enable connection rules logging

Check the **Enable Connection Rules Logging** box and click the **Execute** button. If the logging was successfully enabled, you get a message in the Output part of the panel.

Connections rules logging

The connection rules logging is disabled after installing the firewall. Unless you manually enable the logging, you do not see any IP packets in the Firewall Log file.

NAT also has a separate option to activate the logging. In the configuration client main window, click **NAT** and then **Activation**. The window in Figure 203 is displayed.

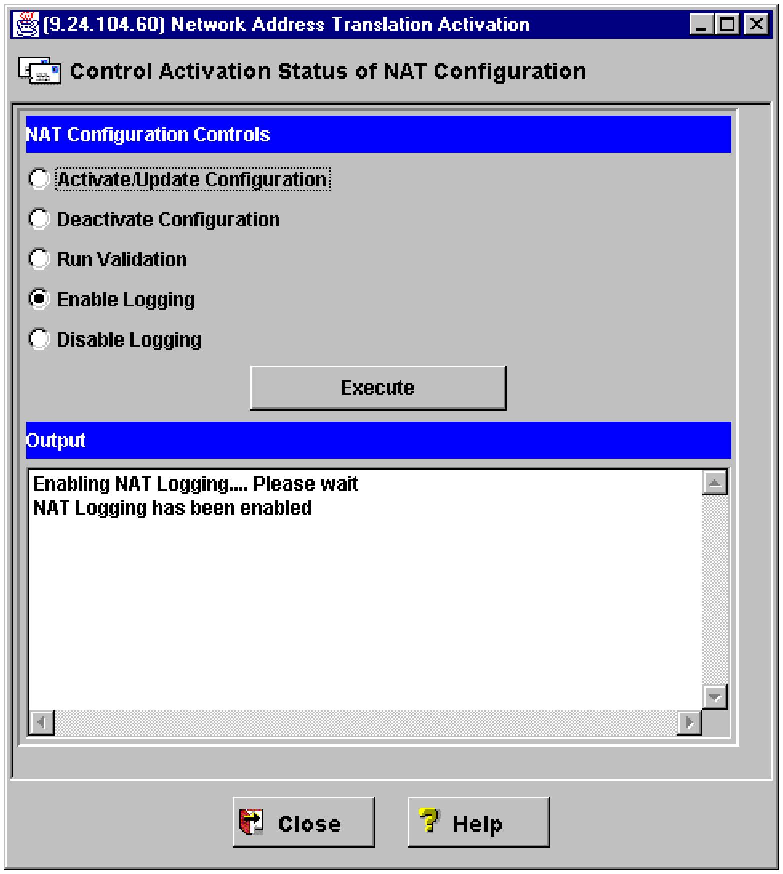

Figure 203. NAT control window

Select **Enable Logging** and click **Execute**. The log entries will be sent to the syslog local4 facility.

11.2 Monitoring and alerts

There are two ways to extract important information from the log files:

- Generating real-time alerts.
- Building analysis reports.

This section covers the generation of real-time alerts. This is the best way to begin extracting information from your logs. Once you become proficient with this facility you will be able to customize a variety of alerts to notify you in case of firewall misuse. Building analysis reports is covered in 11.4, “Building reports” on page 309.

11.2.1 Log monitor thresholds

The generation of real-time alerts is based on the violation of configurable thresholds. A threshold consists of a *count* and a *time* parameter.

A threshold is violated if a number of specific events (count) is exceeded in the specified amount of time (time). You can configure the following types of thresholds:

Table 31. Types of log monitor thresholds

Type	Description
Total authentication failures threshold	Counts all authentication failures, regardless of the originating user ID or host.
Per user threshold	Authentication failures against any particular user ID.
Per host threshold	Authentication failures originating from any particular host.
Message threshold	Occurrences of a ICA message tag in the firewall log file.

You will have to experiment with the threshold settings in your environment so you do not get swamped with too many alerts, but do not lose any important information.

While this list may seem limited, the ability to set a message threshold gives you a lot of options. An overview of messages is in 11.2.3, “Alert message delivery methods” on page 297. For a complete breakdown see *IBM SecureWay Firewall Reference Version 4 Release 1, SC31-8418*.

11.2.2 Alert messages

Alert messages will be delivered by the Log Monitor to keep you informed of firewall use and misuse.

The format of all ICA messages is as follows:

```
CAxxxxa
```

Where `ICA` is a fixed three-byte identifier, `xxxx` is the message number (see Table 32), and `a` is the message severity indicator (see Table 33).

The message numbers are classified into the following categories:

Table 32. ICA message number categories

Numbers	Category
0000 - 0999	Intrusion alarm
1000 - 1999	Filters
2000 - 2999	Proxy-related messages.
3000 - 3999	SOCKS-related messages.
4000 - 4999	Pager-related messages
5000 - 5999	Secure Socket Layer
6000 - 6999	Virtual Private Network
7000 - 8999	Available for Future Use
9000 - 9999	General/others

There are four levels of severity, which correspond to the log priority levels as shown in Table 33.

Table 33. ICA message severity levels

Message severity indicator	Description	Corresponding log priority level
		debug
i	info	information
w	warning	warning
e	error	error
s	severe	critical

11.2.3 Alert message delivery methods

Every time a threshold is violated the firewall generates a real-time alert message (ICA message). The delivery of these messages can take place in four ways. You can combine the delivery methods as you want:

Table 34. Types of alert message delivery

Type	Description
Mail notification	Mails the ICA message to a user or a list of users. Execute Command executes a user-defined command with the ICA message as the first parameter.
Pager notification	Pages a message to a defined pager.
Log entry	Logs the ICA message in an Alert Log file. This file is displayed in the Alerts Display from the configuration client.

11.2.4 Log monitor administration

Open the **System Administration -> System Logs -> Log Monitor Thresholds** document from the configuration client navigation tree. In Figure 204 you see a list of four predefined log monitors.

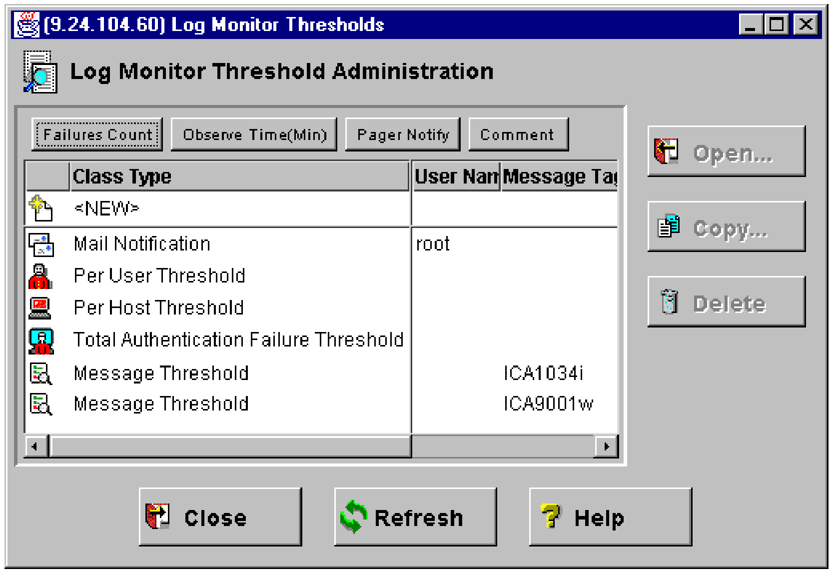

Figure 204. Log monitor threshold administration

You can open the Add Log Monitor panel by double-clicking the **<NEW>** entry in the list shown in Figure 204.

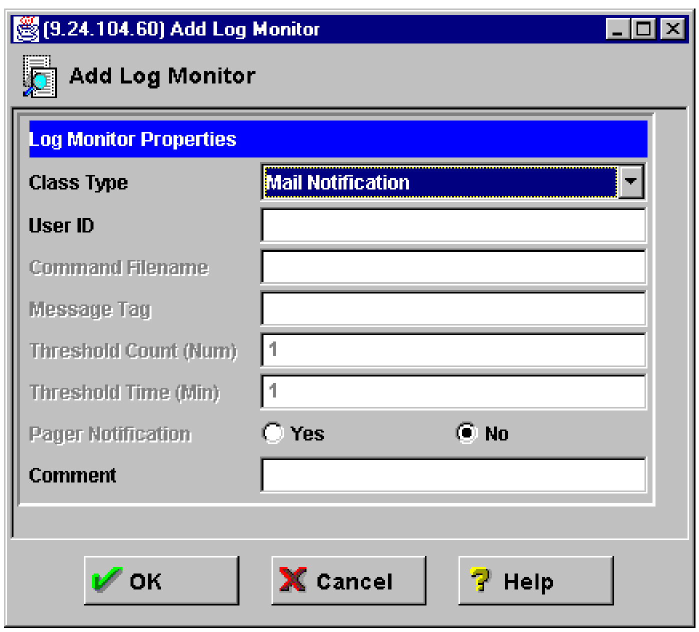

Figure 205. Add log monitor

This window allows you to define new threshold monitors (see 11.2.5, “Configure threshold monitors” on page 299) and new alert message delivery monitors (see 11.2.6, “Configure delivery monitors” on page 301). With the Class Type field you indicate the type of monitor. Click the pull-down menu to choose from the list of available options. This list contains entries to define threshold monitors (see Table 31) and entries to specify delivery monitors (see Table 34). You cannot select the delivery monitor types Log Entry and Pager Notification with this menu.

The other fields are explained in the corresponding sections. Be aware that some fields are not available for all monitors.

11.2.5 Configure threshold monitors

Threshold monitors contain a type and a threshold. In case of threshold violation they trigger an alarm message. As we discussed in 11.2.1, "Log monitor thresholds" on page 295, there are four types of thresholds:

- Total Authentication Failures Threshold
- Per User Threshold
- Per Host Threshold
- Message Threshold

The first three are authentication failure threshold monitors and they have the same options. Therefore we describe only the Total Authentication Failure Threshold monitor and the Message Threshold monitor. You should check the three predefined threshold monitors in Figure 204 and change them if necessary.

There is also one predefined message tag threshold monitor, but you can define as many message tag monitors as you want. You find a complete list of message tags in the *IBM SecureWay Firewall Reference Version 4 Release 1, SC31-8418*.

11.2.5.1 Configure authentication failure monitors

Double-click the **Total Authentication Failure Threshold** monitor in Figure 204 on page 298. The Modify window is displayed.

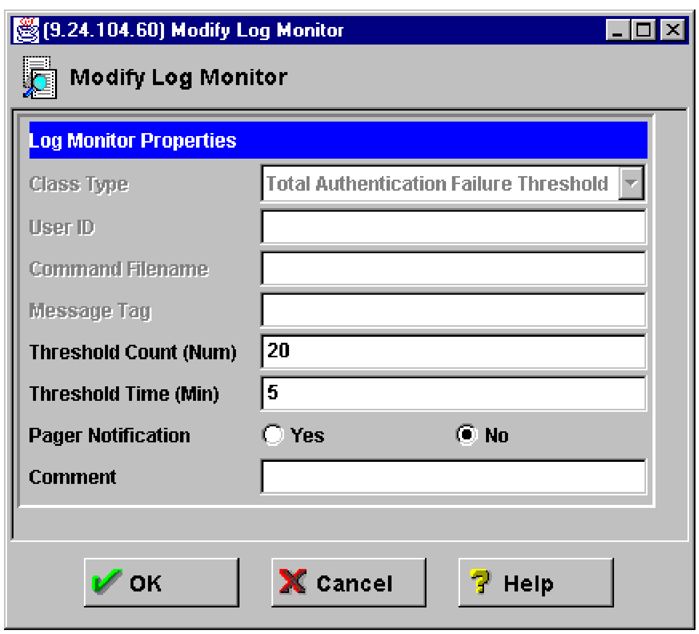

Figure 206. Total authentication failure threshold monitor

You have the following options:

Threshold Count (Num.) Specifies the maximum number of occurrences of a specified log message tag. If the number of occurrences exceeds the threshold count within the specified time (in minutes), the log monitor sets off an alarm. Threshold count cannot exceed 99999.

Threshold Time (Min.) Specifies the time period that defines when a certain number of events has exceeded the threshold. If the threshold has been exceeded, the log monitor sets off an alarm. Threshold time cannot be greater than 99999. A value of 0 for time indicates an unlimited time period.

Pager Notification Enables the pager notification for this monitor when the threshold is exceeded.

Comment Contains an optional textual description of the monitor.

To create new authentication threshold monitors double-click **<NEW>** in Figure 204 on page 298. Choose one of the three authentication threshold monitors from the Class Type pull-down menu and fill out the other fields.

11.2.5.2 Configure message tag monitors

You can define a new message tag threshold monitor by double-clicking **<NEW>** as shown in Figure 204 on page 298.

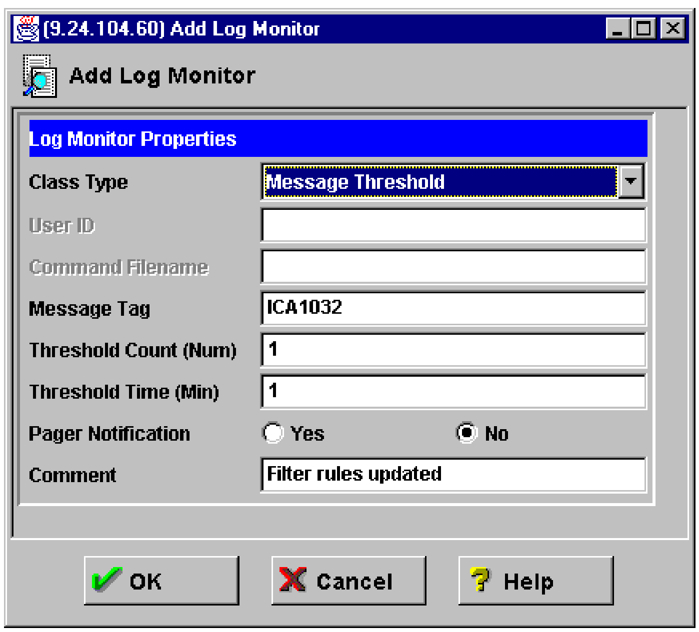

Figure 207. Message tag threshold monitor

Choose **Message Threshold** from the Class Type pull-down menu. Now you have the following options:

Message Tag You specify here the identifying message number. A list of all messages, including associated tags and descriptions, is provided in the *IBM SecureWay Firewall Reference Version 4 Release 1, SC31-8418*.

Threshold Count (Num.) Specifies the maximum number of occurrences of a specified log message tag. If the number of occurrences exceeds the threshold count within the specified time (in minutes), the log monitor sets off the alarm. Threshold count cannot be greater than 99999.

Threshold Time (Min.) Specifies the time period that defines when a certain number of events has exceeded the threshold. If the threshold has been exceeded, the log monitor sets off an alarm. Threshold time cannot be greater than 99999. A value of 0 for time indicates an unlimited time period.

Pager Notification Enables the pager notification for this monitor when the threshold is exceeded.

Comment Contains an optional textual description of the monitor.

11.2.6 Configure delivery monitors

As we discussed in 11.2.3, “Alert message delivery methods” on page 297 there are four alert message delivery monitors:

- Mail Notification
- Pager Notification
- Execute Command
- Log Entry

You can use them to keep you informed of threshold violations from the firewall.

11.2.6.1 Configure mail notification monitors

You can define a mail notification monitor by double-clicking **<NEW>** in Figure 204 on page 298.

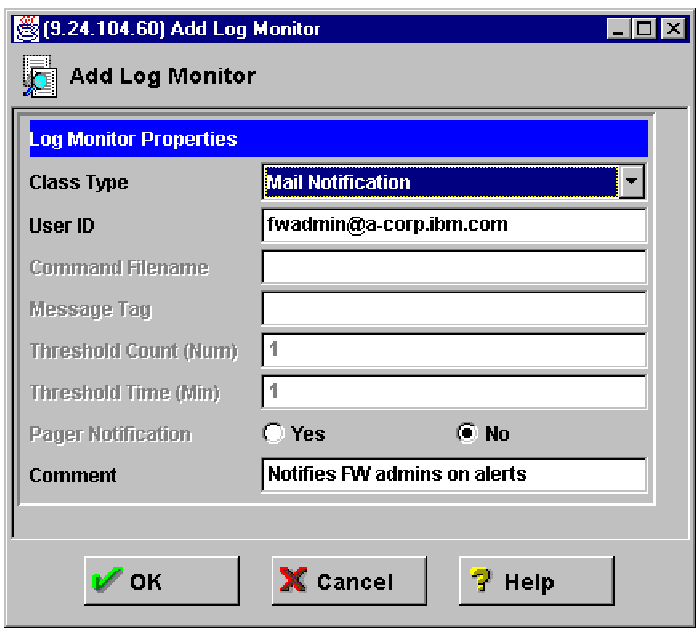

Figure 208. Mail notification delivery monitor

Choose **Mail Notification** from the Class Type pull-down menu. Now you have the following options:

User ID Specify here the e-mail address of the recipient.

Comment Contains an optional textual description of the monitor.

Mailing lists

You can define multiple mail notifications but we recommend you define only one. Create a mailing list for your firewall administrators at your internal mail server and use the address of this list as the recipient. Now you can easily configure this list in case of holidays or the exchange of an administrator.

This mail is sent using AIX sendmail, so you need to add one entry on sendmail configuration file (/etc/sendmail.cf) so it will redirect the mail to an internal server.

Find the line DS in /etc/sendmail.cf. This macro defines the smart relay host (the machine that receives all e-mail that is not local to sendmail). In our example, our internal server is b-mail.b-corp-secure.ibm.com. Our DS macro is:

```
#DRmailer:relayhostname
DRsmtp:b-mail.b-corp-secure.ibm.com
```

If you are using Secure Mail Proxy, then sendmail is not running on your firewall, so you do not need to restart it.

Sendmail as overflow server

In case you are using sendmail as the overflow server on the same machine as the firewall, you do not need to make any change in the sendmail configuration file. For more information on configuring sendmail as an overflow server, see 8.3.5.5, “Overflow server on the firewall” on page 208.

11.2.6.2 Configure pager notification

The pager notification is specified on a per threshold monitor basis. You can decide for every threshold monitor whether you want to send a page or not in case of a violation of this monitor. But you can define only one fixed message that is sent to you in case of an alert. See 11.2.6.3, “Configure command execution monitors” on page 302 for an example to overcome that limitation.

See 11.2.5, “Configure threshold monitors” on page 299 for details on how to enable paging for a threshold monitor. You will find details about pager setup in 11.2.7, “Pager setup” on page 304.

11.2.6.3 Configure command execution monitors

You can define a command execution monitor by double-clicking **<NEW>** in Figure 204 on page 298.

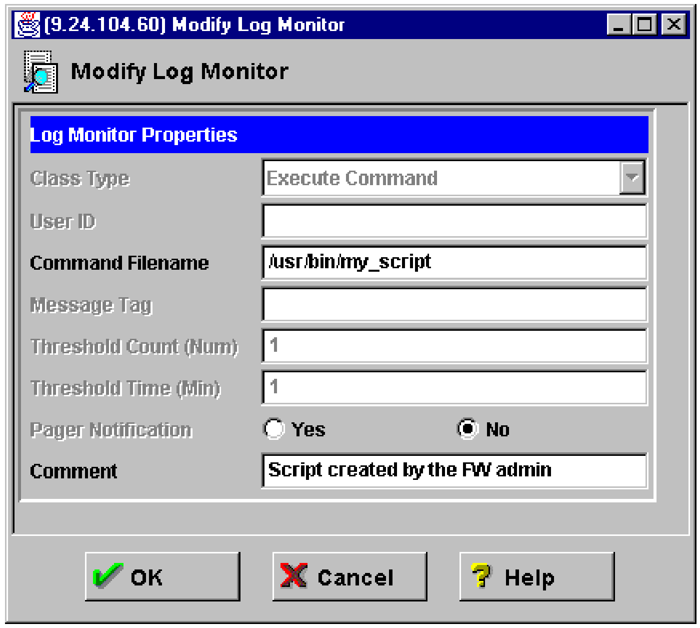

Figure 209. Command execution monitor

Choose **Execute Command** from the Class Type pull-down menu. Now you have the following options:

Command Filename Indicates the file name of the program that will be executed when a threshold is exceeded.

Comment Contains an optional textual description of the monitor.

Serializing

The program you specify is launched in a new process every time an alert is generated and it is possible that multiple instances of the program can be running at the same time. So the program is responsible for serializing access to shared resources.

A descriptive alert message is sent to the program. This enables you do advanced functions, for example:

- Execute programs depending on the ICA message tag.
- Forward messages to your systems management environment.
- Automatically update of your rule base to disable logging of specific packets.

11.2.6.4 Configure log entry notification

The logging of alerts in a special log file takes place if you create an Alert Log facility (see 11.1.3, “Manage log facilities” on page 286). Once you have created this facility, all violations of defined thresholds are displayed in the Alerts Display of the firewall configuration client.

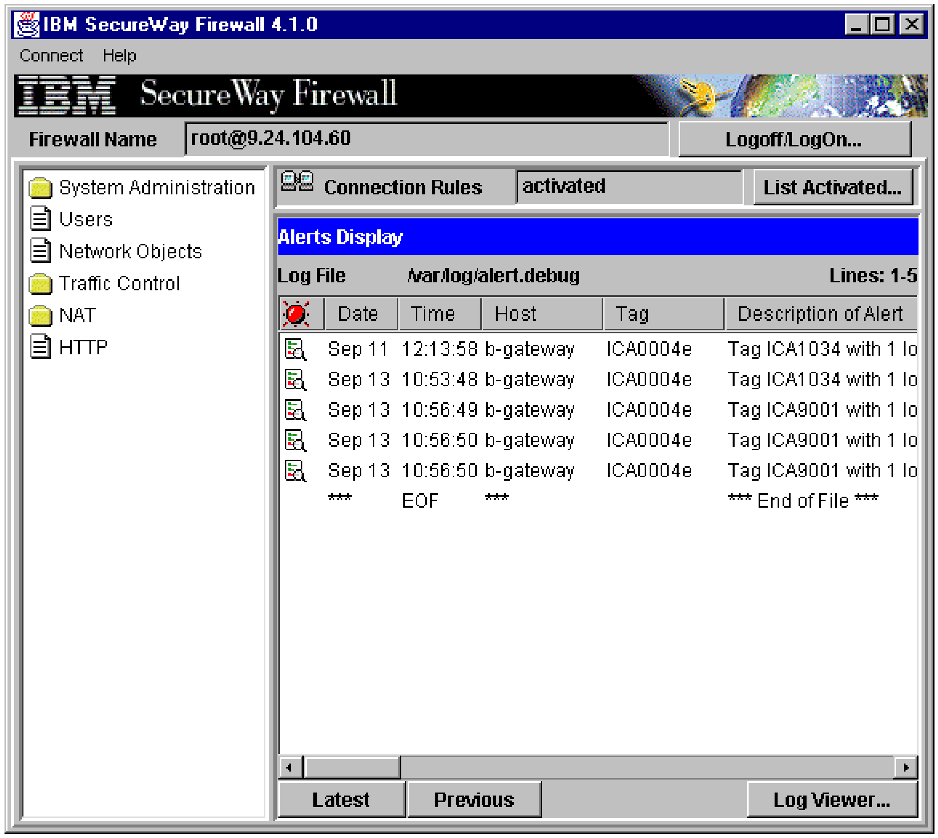

Figure 210. Alerts display in the configuration client

You can get the latest alerts by clicking the **Latest** button in the bottom area. To see old entries use the **Previous** button.

11.2.7 Pager setup

In this section we describe how to configure a pager to work with the firewall. You must configure three components:

- Command customization
- Carrier administration
- Modem administration

Before starting, you need to make sure that you have a carrier where you can dial-in and send pages. Your carrier should be able to page to the locations where your administrators are. You also need to get the correct modem phone numbers, pager ID, and modem parameters from your carrier. The carrier must support the TAP protocol. TAP stands for Telocator Alphanumeric Protocol. It is an industry-standard protocol for sending a page via a modem. See the following URL for details on this protocol:

`http://www.mot.com/MIMS/MSPG/pcia_protocols/tap_v1p8/index.html`

For more information about the pager setup, see *IBM SecureWay Firewall for AIX User's Guide Version 4 Release 1, GC31-8419*.

11.2.7.1 Command customization

To configure the command that is sent to the pager if a threshold is violated, open the **System Administration -> System Logs -> Pager Setup** document in the navigation tree.

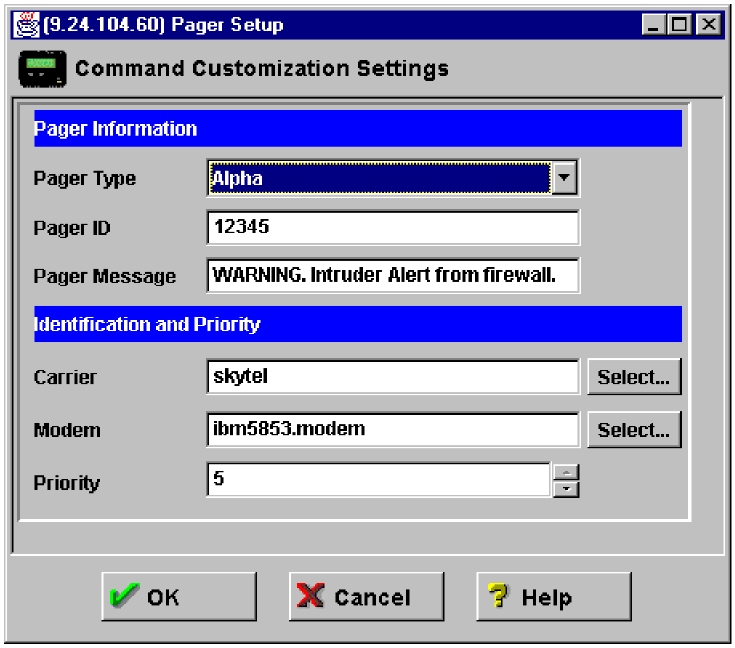

Figure 211. Pager customization

You can specify the following options:

Pager Type: Select the appropriate value from the list: Alpha (alphanumeric) or Numeric.

Pager ID This is usually a unique PIN number assigned to your pager by your carrier company.

Pager Message Specifies the pager message that appears on the pager when the log facilities issue the page. Do not use double quotes at the start and end of your message.

Carrier Name Click **Select** to select or define a carrier.

Modem Name Click **Select** to select or define a modem.

Priority Select the priority for sending the page. The highest value is 5 and the lowest is -1.

You can specify only one message for your defined threshold monitors. For numeric pagers, this must be a number only. For alphanumeric pagers, this can be a text message. Do not exceed the maximum message length for alphanumeric pagers or your message might be truncated.

11.2.7.2 Carrier administration

If you click the **Select** carrier button in Figure 211, you get the Carrier Administration window.

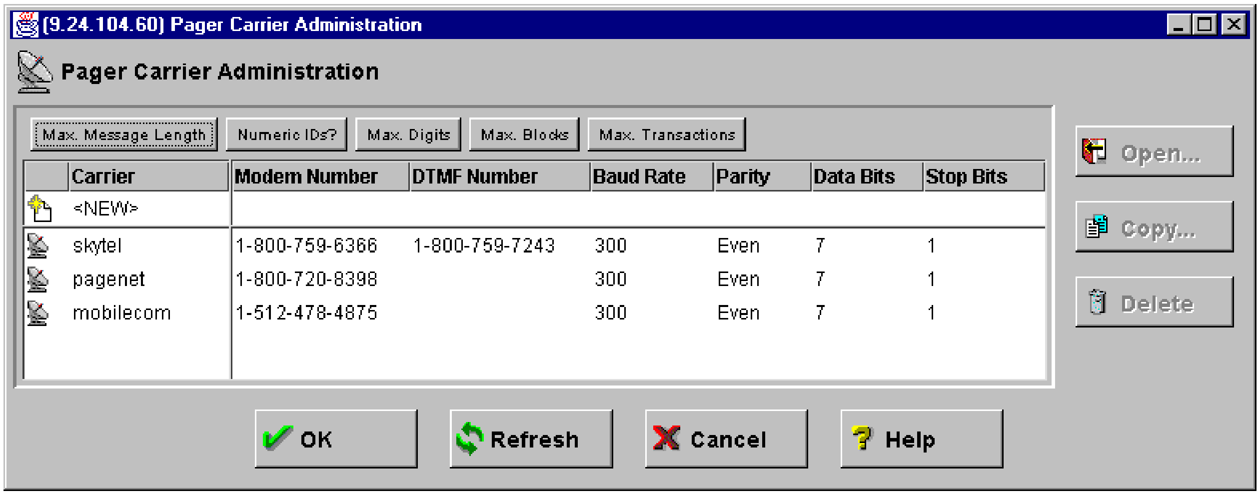

Figure 212. Carrier administration

Select one of the predefined carriers or define a new one by double-clicking **<NEW>**.

11.2.7.3 Modem administration

If you click the **Select** modem button in Figure 211 on page 305, you get the Modem Administration window.

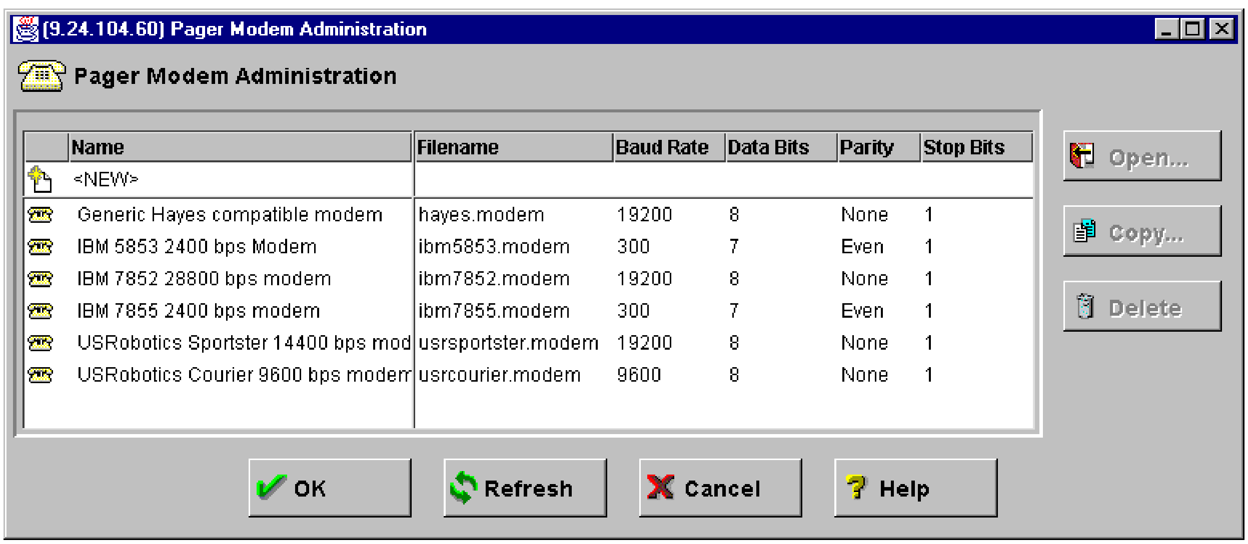

Figure 213. Modem administration

Select one of the predefined modems or define a new one by double-clicking **<NEW>**.

It is quite possible you will need to adjust the init string to achieve proper communication between your modem and your carrier's modem.

11.3 Log file formats

You need to understand the log files and their format if you want to extract further information from them. As described in 11.1.2, “Log facilities” on page 286 there are three log facilities using their own formats:

- Firewall Log
- Alert Log
- Admin Audit Log

11.3.1 Firewall log format

The entries in the firewall log files have the following format:

```
date time fw_name:year;pid: Amsg_num;msg_ID;var_1;...;var_n;
```

The fields are explained in Table 35.

Table 35. Firewall log entry fields

Field	Description
date	Month and day when entry occurs. The format is `MMM DD`.
time	Time when entry occurs. The format is `HH:MM:SS`.
fw_name	The name of the firewall that logged this entry.
year	`year` is the four-character year.
pid	`pid` is the process ID to which the entry applies.
Amsg_num	`msg_num` is a sequential integer that the Report Utilities use to access the appropriate, translated message text from the fw_log.cat file. The number is preceded by a log level indicator letter. This indicator distinguishes both the platform that originated the log entry and any differences in log format.
msg_ID	`msg_ID` is the external number of the message, such as `ICA0001e` (see 11.2.2, “Alert messages” on page 296).
var_n	`var_1 to var_n` represents the values of message variables, where n is the number of variables in the message definition.

Figure 214 shows a few example entries in the short form.

```
Sep 13 11:18:47 b-gateway : 1999;5422: 2151;ICA1075i;1;501;0;5;p;o;10.2.1.1;9.24
.104.60;10.2.1.2;tcp;sp:;23;dp:;35774;l;s;n;0;43;
Sep 13 11:18:47 b-gateway : 1999;5422: 2151;ICA1075i;1;501;0;5;p;i;10.2.1.1;10.2
.1.2;9.24.104.60;tcp;sp:;35774;dp:;23;r;s;n;0;40;
Sep 13 11:18:47 b-gateway : 1999;34022: 18143;ICA9071i;11:18:47;09-13-1999;
Sep 13 11:18:57 b-gateway : 1999;5938: 9;ICA0004e;ICA1034;1;
Sep 13 11:18:57 b-gateway : 1999;34022: 2069;ICA1034i;11:18:57;09-13-1999;
Sep 13 11:19:07 b-gateway : 1999;32494: 18065;ICA9032i;11:19:07;Sep-13-1999;
Sep 13 11:19:07 b-gateway : 1999;5422: 18091;ICA9045i;9.24.104.117;0;10.2.1.3;0;
```

Figure 214. Firewall log example entries

You can also convert this log into a more readable form using the `fwlogtxt` command:

```
Sep 13 11:18:47 1999 b-gateway: ICA1075i:ft:1 tid:501 sid:0 #:5 R:p o 10.2.1.1 s
:9.24.104.60 d:10.2.1.2 p:tcp sp:23 dp:35774 r:l a:s f:n T:0 l:43
Sep 13 11:18:47 1999 b-gateway: ICA1075i:ft:1 tid:501 sid:0 #:5 R:p i 10.2.1.1 s
:10.2.1.2 d:9.24.104.60 p:tcp sp:35774 dp:23 r:r a:s f:n T:0 l:40
Sep 13 11:18:47 1999 b-gateway: ICA9071i: Kernel level filters logging stopped a
t 11:18:47 on 09-13-1999
Sep 13 11:18:57 1999 b-gateway: ICA0004e: ALERT - Tag ICA1034 with 1 log entries
.
Sep 13 11:18:57 1999 b-gateway: ICA1034i: Filter support deactivated at 11:18:57
 on 09-13-1999
Sep 13 11:19:07 1999 b-gateway: ICA9032i: NAT configuration updated at 11:19:07
on Sep-13-1999.
Sep 13 11:19:07 1999 b-gateway: ICA9045i: NAT allocated address:port 9.24.104.11
7:0 for secured address:port 10.2.1.3:0
```

Figure 215. Converted firewall log example entries

11.3.2 Alert log format

The entries in the Alert Log files have the following format:

```
date time fw_name:year;pid: msg_num;msg_ID: descr
```

The fields are explained in Table 36.

Table 36. Alert log entry fields

Field	Description
date	Month and day when entry occurs. The format is MMM DD.
time	Time when entry occurs. The format is HH:MM:SS.
fw_name	The name of the firewall which logged this entry.
year	year is the four-character year.
pid	pid is the process ID to which the entry applies.
Amsg_num	msg_num is a sequential integer that the Report Utilities use to access the appropriate, translated message text from the fw_log.cat file. The number is preceded by a log level indicator letter. This indicator distinguishes both the platform that originated the log entry and any differences in log format.
msg_ID	Contains one of the four threshold violation message IDs (ICA0001, ICA0002, ICA0003, or ICA0004).
descr	Shows the message tag which has been violated and the number of occurrences.

Figure 216 shows a few example entries.

```
Sep 11 12:13:58 b-gateway : ICA0004e: ALERT - Tag ICA1034 with 1 log entries.
Sep 13 10:53:48 b-gateway : ICA0004e: ALERT - Tag ICA1034 with 1 log entries.
Sep 13 10:56:49 b-gateway : ICA0004e: ALERT - Tag ICA9001 with 1 log entries.
Sep 13 10:56:50 b-gateway : ICA0004e: ALERT - Tag ICA9001 with 1 log entries.
```

Figure 216. Alert log example entries

11.3.3 Audit log format

The entries in the Admin Audit Log files have the following format:

```
date time fw_name:user_ID;action; var_1=a ... var_n=z[;rc=n]
```

The fields are explained in Table 37.

Table 37. Audit log entry fields

Field	Description
date	Month and day when entry occurs. The format is `MMM DD`.
time	Time when entry occurs. The format is `HH:MM:SS`.
fw_name	The name of the firewall which logged this entry.
user_ID	The administrator user ID which performs the action.
action	The administrative action performed.
var_n=x	Shows the settings of the variables used for the action. The format is `var_name=value`.
rc=n	The return code of the action. This parameter is optional.

Figure 217 shows a few example entries.

```
Sep 13 11:56:41 b-gateway : root;fwListPager; logonmode=host admin_userid=root;r
c=1
Sep 13 11:56:45 b-gateway : root;fwGetFilterStatus; logonmode=host admin_userid=
root;rc=115
Sep 13 12:01:40 b-gateway : root;fwChangeConnection; logonmode=host admin_userid
=root name=$b+FTP in:SI>SN include=0 overwrite=1 index=519;rc=1
Sep 13 12:01:42 b-gateway : root;fwUpdateFilter; logonmode=host admin_userid=roo
t name=$b+FTP in:SI>SN include=0 overwrite=1 index=519;rc=1
```

Figure 217. Audit log example entries

11.4 Building reports

The firewall allows you to log the different events happening in your firewall. For example, you can log denied or permitted IP traffic, SOCKS and proxy usage, and mail events.

As we have said, there are two ways to extract important information out of the log files:

- Generating real-time alerts.
- Building analysis reports.

Generating a real-time alert is covered by 11.2, “Monitoring and alerts” on page 295. You should also build useful reports from your firewall logging. The reports may provide information about attack rates and types of attacks.

You may also be interested in information of resource usage from the secure network, for example, information about the amount and duration of Internet traffic, or total number of sessions per given period, regardless of what session type or total number of bytes transferred by FTP per given period, and so on.

The firewall log files contain this information but it is not easily derived from them. With the Report Utilities it is possible to convert the firewall log files into import files for database managers such as DB2/6000, DB2, or Oracle. In this way you can use all the power of the Structured Query Language (SQL), or other tools such as IBM's Visualizer or Query Management Facility to query the data and generate reports.

11.4.1 Report utilities

You can use the firewall Report Utilities to create full text log files or import files for databases out of log files from the firewall log facility. We describe here only the way to create the files necessary for creating reports with DB2. Please refer to the *IBM SecureWay Firewall for AIX User's Guide Version 4 Release 1, GC31-8419* and the *IBM SecureWay Firewall Reference Version 4 Release 1, SC31-8418* for usage of the Report Utilities for other purposes.

The Report Utilities can be used with the configuration client or via the command line. The command line programs are also installed with a configuration client on a remote machine.

Where to use?

We recommend that you transfer the log files to a remote machine and run the Report Utilities from the command line. Do not forget to create a connection from the firewall to your remote machine for the file transfer.

The Report Utilities consists of the programs and scripts listed in Table 38. In AIX, the commands are installed in the /usr/bin directory, and in Windows (NT/95/98) you can find the sample SQL scripts in C:\ProgramFiles\IBM\Firewall\sample\report\ and the commands in C:\ProgramFiles\IBM\Firewall\bin. Note that not all commands are available for AIX. Check the following table for all commands and their availability for each version of the configuration client (AIX and Windows).

Table 38. Firewall report utilities

File	Description	AIX	Windows
fwlogcvrt	Program to convert a Windows NT firewall log files to AIX firewall log files.	No	Yes
fwlogtxt	Program to generate full-text messages from a firewall log file.	Yes	Yes
fwlogtbl	Program to generate database import files, in DEL (delimited) format, from a firewall log file.	Yes	Yes
fwschema.ddl	File of SQL Data Definition Language (DDL) statements, suitable for defining the database tables.	No	Yes
fwimport.dat	File of DB2 import statements, suitable for importing the DEL files into the database tables.	No	Yes
fwqrysmp.dml	File of SQL Data Manipulation Language (DML) statements, suitable for generating sample reports.	No	Yes

The SQL scripts are specific to the DB2 family, but you can modify them for other database managers.

11.4.1.1 Using the configuration client

To create the DB2 import files with the configuration client select **System Administration -> System Logs -> Report Utilities** from the navigation tree. You will see the window shown in Figure 218.

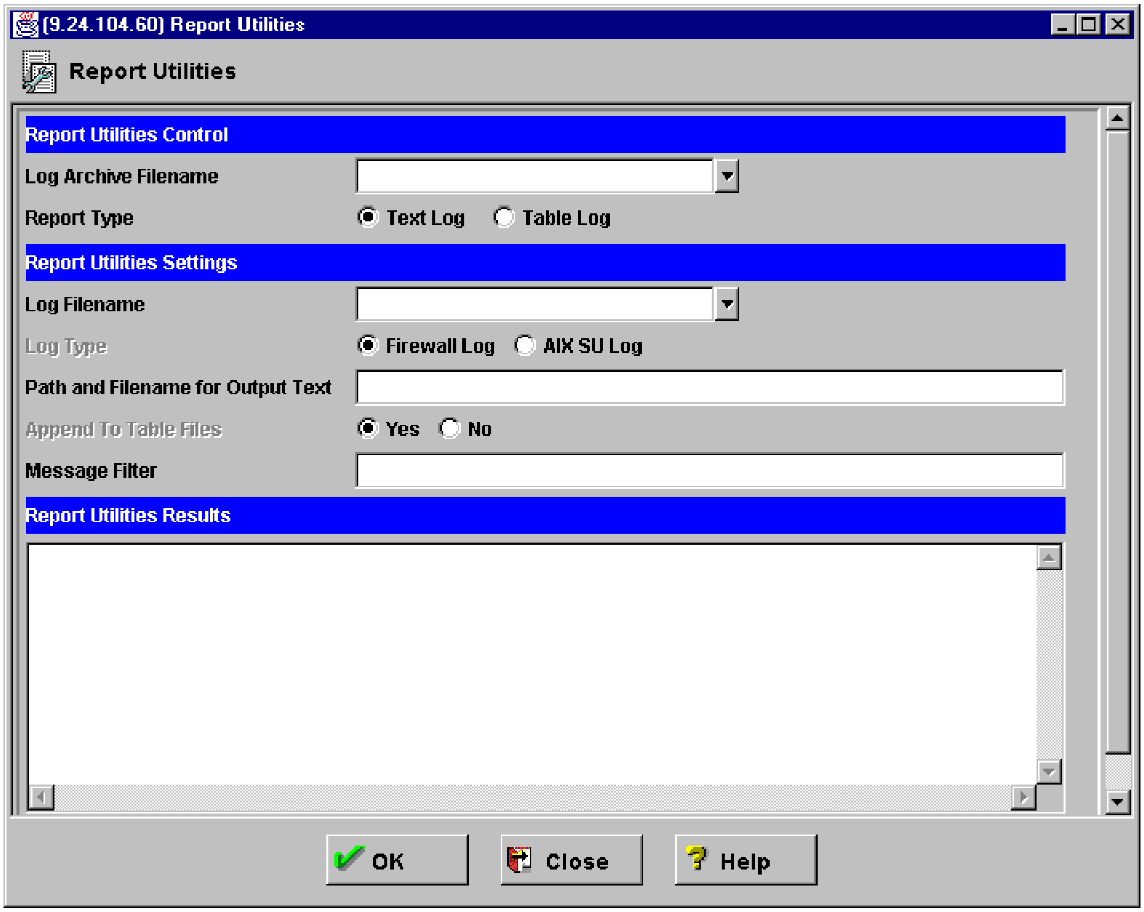

Figure 218. Report utilities

The different fields are:

Log Archive Filename The log archive is the name of the archive file that you want to extract the text log file from. If you want a log file that is not archived, leave this field blank.

Report Type To produce the expanded log message text, select **Text Log**. To create tabulated files for DB2 usage, select **Table Log**.

Log Filename The log filename is any one of the compressed archived log files. If you made an entry in the log archive filename field, you can select the button in the Log Filename field to choose which log to work with. If you do not enter a log archive, the log file name you enter here must be the name of a valid, uncompressed firewall log file. You must specify a full path.

If you change the Report Type radio button to **Table Log** you have the following options.

Log Type Select the log type.

Directory for Output Files The Report Utilities create a bunch of files that can be imported to DB2. Specify here an output directory for these files.

Append To Table Files Select **Yes** to append the results of a table log request to existing tabulated files or **No** to replace the existing files.

Message Filter This field works like the command `grep` on AIX: it looks for lines that contain the word you typed, and ignores the other lines. You can use this if you want to get only specific entries.

Finally click **OK** to run the Report Utilities. After successful completion you find the import files in the directory specified.

11.4.1.2 Using the command line

The import files can also be generated from the command line using the following command:

```
# fwlogtbl -w -d . /var/log/firewall.debug
```

The parameter -w specifies that eventually existing output files should be replaced. -d is followed by the output directory. The last parameter is the input log file. Please remember, you can use firewall log files only.

11.4.1.3 Database import files

After using the Report Utilities either with the configuration client or via the command line you find several files created in your output directory. The resulting files are in delimited ASCII (DEL) file format, with no character string delimiters, and using semicolon (;) as the column delimiters. The extension for these files is .tbl.

The following table shows the relationship between the files and the tables created by the SQL scripts.

Table 39. Report utilities database tables

Database table	Import file	Description
ADMIN_ALERT	a_alert.tbl	Messages related to intrusion alerts.
FILTER_ACTIVE_RULE	f_rule.tbl	Active IP filter rules.
FILTER_INFO	f_info.tbl	Error or general information messages related to IP filters.
FILTER_MATCH	f_match.tbl	Matched IP filter rules.
FILTER_STATUS	f_stat.tbl	Information on status changes of IP filters.
PAGER_INFO	pgr_info.tbl	Information related to the paging feature of the firewall for those pager messages that are mapped to the database.
PROXY_FTP	p_ftp.tbl	FTP action information from FTP sessions.
PROXY_HTTP	p_http.tbl	HTTP action information from proxy sessions.
PROXY_INFO	p_info.tbl	Error or general information messages related to proxy.

Database table	Import file	Description
PROXY_LOGIN	p_login.tbl	Information (primarily regarding authentication) about successful proxy logins.
PROXY_STATUS	p_stat.tbl	Proxy status information.
SERVER_INFO	srv_info.tbl	Information about Configuration Server status and activities.
SESSION	session.tbl	SOCKS and proxy session start/stop information.
SOCKS_FTP	s_ftp.tbl	SOCKS FTP action information from FTP sessions.
SOCKS_INFO	s_info.tbl	Error or general information messages related to SOCKS.
SSL_INFO	ssl_info.tbl	Information about SSL status and activities.
TUNNEL_CONTEXT	t_cntxt.tbl	Active tunnel context specifications.
TUNNEL_POLICY	t_policy.tbl	Tunnel policy statements.
TUNNEL_STATUS	t_stat.tbl	Information on status changes of tunnels.

Chapter 12. Network Security Auditor

Once an Internet firewall has been installed, it must be maintained correctly to allow it to continue to provide sufficient security. Auditing the firewall to be sure that no doors are open to be exploited by hackers is an important task that should be done on a regular basis.

Network Security Auditor, which is part of the IBM SecureWay Firewall V4.1 for AIX, is a tool that scans target machines and tries to exploit detected services of known weaknesses. For example, it scans ports to find security leaks. Some obvious passwords are tried on services that ask for a password such as Telnet and FTP, or commands are tried that are considered dangerous or risky on sendmail, such as `DEBUG`, `VRFY` and `EXPN`.

NSA has been enhanced in IBM SecureWay Firewall V4.1 for AIX.

12.1 NSA enhancements

Some of the new features are:

- New vulnerability tests:
  - NetScape Enterprise Server buffer overflow
  - AMD buffer overflow
  - HP/UX remwatch buffer overflow
  - AIX instsrv vulnerability
  - MS IIS buffer overflow
  - FTP does user name/password testing
  - SMB user name/password testing
- Backdoor server detection (Trojan Horses) for Acid Shiver, GirlFriend, Hack'a'Tack, Portal of Doom, NetSphere, Gatecrasher, EvilFTP, phAse-Zero, SubSeven, NetBus, BackOrifice, DeepThroat
- New functions:
  - Finding filters that allow you to screen out known false positives during report generation without manually editing the reports.
  - Add *split* subcommand to split a findings database along scan groups.
  - Added ability for NSA report generation to incorporate *raw* formatted records into the report. This allows you to export the scan results as a raw report, process them, then import them back into a report.
  - New underlying storage library. Provides more reliable storage while providing faster access with smaller databases; however, can still process old databases.
  - Host discovery using TCP and UDP ports. This is similar to pingfirst (which uses ICMP Echo Requests), only TCP and UDP packets are used to find "live" hosts.
  - Added ability to specify *names* to which policy violation scores maps. This allows you to have high, medium, or low in policy violation reports.

- Added control over how many ICMP Echo Requests can be outstanding. The flow of ICMP Echo Requests has also been tuned to provide a more consistent flow of packets (performance improvement).
- New server recognitions (not including SunRPC servers or backdoor servers):
  - Local Mail Transport Protocol (LMTP)
  - Lightweight Directory Access Protocol (LDAP)
  - CSO nameserver
  - HP/UX remwatch
  - AIX instsrv

12.2 Implementation

We recommend that you install NSA on a machine on the non-secure network, so it does not need to go through any router to be able to scan the firewall (it would take longer to complete the scanning and it would impact on the performance of the network).

You can use NSA by running the command `nsa`. For example, to scan the machine 9.24.104.60 using the default scanning options, run:

```
# nsa scan 9.24.104.60
```

In this example, since we did not inform the output file, it will send the report to the screen.

You can also choose the scan type that you want, that means, the ports and services you want to scan on the target machine. The available scan types are defined in the file /etc/nsa/scannerdefs.

We did a test using the scan type "firewall", which is the most complete predefined scan type. We ran it to a firewall machine with filters allowing all traffic (it had no filter protection). We used the following command:

```
nsa scan --scantype firewall --outfile report1 9.24.104.60
```

The results of the scanning will be saved in the file report1 in the current directory. See the following report generated by the preceding example:

```
Network Services Audit Report

   Report Date: Thursday, September 16, 1999 18:32

      o Name: b-gateway.gw.itso.ral.ibm.com (9.24.104.60)
        Operating System: UNIX IBM AIX 4.3
        Audit Date: `Thursday, September 16, 1999 18:24'
        Auditor: `root@jupiter.itso.ral.ibm.com'

       Security Audit Summary
```

```
          o Configuration Settings - 1
          o Potential Vulnerabilities - 1

        Security Audit Breakdown

         Configuration Settings - 1
          o Access Control Configuration - 1
             o User Account Configuration - 1

         Potential Vulnerabilities - 1
          o Problems with SMTP Service - 1

        Security Audit Findings
          o Configuration Settings
             o Access Control Configuration
                o User Account Configuration
          o Potential Vulnerabilities
             o Problems with SMTP Service
                o [SMTP at port 25] Server accepts sender address of
                  `<"|/usr/bin/id">' with response
                  250 |/usr/bin/id Okay
          o Active Network Servers
             o telnet is active on TCP port 23.
             o DNS is active on TCP port 53.
             o FTP is active on TCP port 21.
             o socks5 is active on TCP port 1080.
             o SMTP is active on TCP port 25.
             o X is active on TCP port 6000.
             o HTTP is active on TCP port 8080.
             o DNS is active on UDP port 53.

          o Available Network Services
             o User Login Services
                o telnet is active on TCP port 23.
                o FTP is active on TCP port 21.
          o Operating system is `UNIX IBM AIX 4.3'.
          o Server Version Strings
             o [53/UDP] DNS server version is `UNKNOWN'.
             o [53/TCP] DNS server version is `UNKNOWN'.
             o [8080/TCP] HTTP server version is `IBM-PROXY-FW/2.0'.
             o [21/TCP] FTP server version is `4.1 Tue Sep 8 15:35:59 CDT
               1998'.
          o Network Transport Information
             o IP Transport Information
o Host responded to ICMP Echo Request.
          o Port Scan Information
             o TCP Port Scan Data
                o The following TCP ports were scanned:
                  1-65535
                o The following TCP ports were visible:
                  1-65535
                o The following TCP ports were active:
                  21, 23, 25, 53, 1014, 1080, 6000, 8080
                o The servers on these TCP ports could not be identified:
                  1014
                o The servers on these TCP ports terminated immediately:
                  None
             o UDP Port Scan Data
```

```
            o The following UDP ports were scanned:
              1-65535
            o The following UDP ports were visible:
              1-513, 515-65535
            o The following UDP ports were active:
              53
            o The following UDP ports did not respond:
              514
            o The servers on these UDP ports could not be identified:
              None
      o Server Banners
         o [23/TCP] telnet server banner -

 telnet (b-gateway.gw.itso.ral.ibm.com)

login:

         o [25/TCP] SMTP server banner -

220-b-gateway.gw.itso.ral.ibm.com Connection Established.
220 ESMTP

         o [21/TCP] FTP server banner -

220 b-gateway.gw.itso.ral.ibm.com FTP server (Version 4.1 Tue Sep 8 15:35:59
CDT 1998) ready.
```

After that, we repeated the scanning on the firewall machine, but this time we removed the "permit all" configuration, so this scanning can show us which services are "reachable" from the outside. The only ports we should allow to machines from the non-secure network (this includes the Internet) are mail and DNS. This time we used the scan type default.

We got the following report from our firewall:

```
Network Services Audit Report

   Report Date: Friday, September 17, 1999 08:14

      o Name: b-gateway.gw.itso.ral.ibm.com (9.24.104.60)

        Operating System: Unknown
        Audit Date: `Friday, September 17, 1999 08:12'
        Auditor: `root@jupiter.itso.ral.ibm.com'

       Security Audit Findings

         o Active Network Servers
            o SMTP is active on TCP port 25.
            o DNS is active on UDP port 53.

         o Server Version Strings
            o [53/UDP] DNS server version is `UNKNOWN'.

         o Port Scan Information
            o TCP Port Scan Data
```

```
o The following TCP ports were scanned:
                21-23, 25, 109-111, 139, 143, 512-514, 6000
              o The following TCP ports were visible:
                25
              o The following TCP ports were active:
                25
              o The servers on these TCP ports could not be identified:
                None
              o The servers on these TCP ports terminated immediately:
                None

           o UDP Port Scan Data
              o The following UDP ports were scanned:
                53, 69, 111, 137, 161, 177
              o The following UDP ports were visible:
                53
              o The following UDP ports were active:
                53
              o The following UDP ports did not respond:
                69, 111, 137, 161, 177
              o The servers on these UDP ports could not be identified:
                None
o Server Banners
           o [25/TCP] SMTP server banner -

220-b-gateway.gw.itso.ral.ibm.com Connection Established.
220 ESMTP
```

While the scan is being performed, you can press Ctrl+X and it shows the estimated time for completion of the command:

```
# nsa scan 9.24.104.60
Network Services Audit (NSA)
Version 1.2.0; Jun 11 1999 07:01:40 [AIX 4.1.3.0 power]

Copyright (c) 1996-1999 International Business Machines Corp.
This product is proprietary material belonging to International
Business Machines Corp.  Disclosure to, or possession by,
unauthorized entities is prohibited.

Password:
<CTRL-X>
**********************************************************************
9.24.104.60 39671 TCP left, 48735 UDP left;   1 task  [11:48:47]
Scan completed on 0 hosts.


Estimated time of completion: Saturday, 08:00

**********************************************************************
```

You can create your own predefined scan type by editing the file /etc/nsa/scannerdefs. We added the following lines to the end of this file:

```
define my-scantype
  tcpports 21,23,25,110,6000
  udpports 53
  options no-user-login
end
```

We used this scan type by running the following command:

```
# nsa scan --scantype my-scantype --outfile my-report 9.24.104.60
```

You can find more information about NSA on the online documentation in the directory /usr/lpp/nsauditor/doc, or at:

http://dr.watson.ibm.com/nsa

Appendix A. Special notices

This publication is intended to help network security specialists who plan, install, implement and manage firewalls to protect corporate intranets. The information in this publication is not intended as the specification of any programming interfaces that are provided by IBM SecureWay Firewall V4.1 for AIX, Program Number 5697-F48. See the PUBLICATIONS section of the IBM Programming Announcement for IBM SecureWay Firewall V4.1 for AIX, Program Number 5697-F48 for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these 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 IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter 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 the IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785.

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, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM ("vendor") products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

DB2	OS/390
eNetwork	OS/400
Hummingbird	RS/6000
IBM	S/390
IMS	SecureWay
Netfinity	WebSphere
OS/2	XT

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc. in the United States and/or other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and/or other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States and/or other countries.

PC Direct is a trademark of Ziff Communications Company in the United States and/or other countries and is used by IBM Corporation under license.

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States and/or other countries.

UNIX is a registered trademark in the United States and/or other countries licensed exclusively through X/Open Company Limited.

SET and the SET logo are trademarks owned by SET Secure Electronic Transaction LLC.

Other company, product, and service names may be trademarks or service marks of others.

Appendix B. Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

B.1 International Technical Support Organization publications

For information on ordering these ITSO publications see "How to get IBM Redbooks" on page 327.

- *Guarding the Gates Using the IBM eNetwork Firewall V3.3 for Windows NT,* SG24-5209
- *The Domino Defense: Security in Lotus Notes and the Internet,* SG24-4848
- *The Technical Side of Being an Internet Service Provider,* SG24-2133
- *A Comprehensive Guide to Virtual Private Networks, Volume I,* SG24-5201
- *A Comprehensive Guide to Virtual Private Networks, Volume II,* SG24-5234
- *A Comprehensive Guide to Virtual Private Networks, Volume III,* SG24-5309

B.2 Redbooks on CD-ROMs

Redbooks are also available on the following CD-ROMs. Click the CD-ROMs button at `http://www.redbooks.ibm.com/` for information about all the CD-ROMs offered, updates and formats.

CD-ROM Title	Collection Kit Number
System/390 Redbooks Collection	SK2T-2177
Networking and Systems Management Redbooks Collection	SK2T-6022
Transaction Processing and Data Management Redbooks Collection	SK2T-8038
Lotus Redbooks Collection	SK2T-8039
Tivoli Redbooks Collection	SK2T-8044
AS/400 Redbooks Collection	SK2T-2849
Netfinity Hardware and Software Redbooks Collection	SK2T-8046
RS/6000 Redbooks Collection (BkMgr Format)	SK2T-8040
RS/6000 Redbooks Collection (PDF Format)	SK2T-8043
Application Development Redbooks Collection	SK2T-8037
IBM Enterprise Storage and Systems Management Solutions	SK3T-3694

B.3 Other publications

These publications are also relevant as further information sources:

- *AIX Version 3.2 & V4 Performance Monitoring and Tuning Guide,* SC23-2365
- *IBM SecureWay Firewall for AIX User's Guide Version 4,* GC31-8419

  This book ships with the firewall product in electronic form. It contains basic installation and configuration instructions.
- *IBM SecureWay Firewall Reference Guide for AIX Version 4,* SC31-8418

  This book ships with the firewall product in electronic form. It contains advanced configuration options on the IBM Firewall.
- *Building Internet Firewalls,* D. Brent Chapman and Elizabeth D. Zwicky, (O'Reilly Associates, Inc.), ISBN 1565921240.

This book covers a great deal about firewalls, starting from firewall architectures to actually implementing them. Of great value is its extensive treatment of firewall filtering rules, most of which you can apply to the IBM SecureWay Firewall V4.1 for AIX.

- *DNS and BIND* by Albitz and Liu (O'Reilly and Associates, Inc.), ISBN 1565925122

  Hands down, this is the bible for DNS and BIND.

- *Sendmail* by Costales and Allman, (O'Reilly and Associates, Inc.), ISBN: 1565922220

  A cryptic book for a cryptic program, which is why it was indispensable for configuring Sendmail.

- *Actually Useful Internet Security Techniques* by Hughes (New Riders Publishing), ISBN 1562055089

  This book provides some insight into Internet security.

B.4 Referenced Web sites

- `http://www.redbooks.ibm.com`
- `http://dr.watson.ibm.com/nsa`
- `http://www.software.ibm.com/security/firewall/support/fixes/fwaix.html`
- `http://www.cert.org/advisories/index.html`
- `http://www.ers.ibm.com/tech-info/advisories/sva/index.html`
- `http://service.software.ibm.com/cgi-bin/support/rs6000.support/downloads`
- `http://www.webshopper.com/graphics/buttons/showproducts.gif`
- `http://home.netscape.com/escapes/search/netsearch_6.html`
- `http://www.pc.ibm.com/msprotect/ncommerce3/ExecMacro/ccadmin.d2w/report`
- `http://static.wired.com/advertising/blipverts/WebMD/aging.gif`
- `http://www1.raleigh.ibm.com/pics/PicsRULZ.html`
- `http://www.w3.org/PICS`
- `http://www.rsac.org/ratingsv01.html`
- `http://www.ibm.com/`
- `http://www.aventail.com`
- `http://www.socks.nec.com`
- `http://www.microsoft.com/ie`
- `http://www.software.ibm.com/webservers/wte`
- `http://www.socks.nec.com/socks5.html`
- `http://www.hummingbird.com`
- `http://www.socks.nec.com/sockscap.html`
- `http://www.socks.nec.com/how2socksify.html#runsocks`
- `http://www.hummingbird.com/products/socks/install.html`
- `http://www.icq.com/firewall/icqsocks5.html`

- http://publib.boulder.ibm.com/pubs/html/as400/v4r4/ic2924/info/INFOCENT.HTM
- http://www.ietf.org/html.charters/nat-charter.html
- http://www.ietf.org/html.charters/ipsec-charter.html
- http://www.mot.com/MIMS/MSPG/pcia_protocols/tap_v1p8/index.html
- http://w3.itso.ibm.com
- http://w3.ibm.com
- http://www.elink.ibmlink.ibm.com/pbl/pbl

How to get IBM Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, redpieces, and CD-ROMs. A form for ordering books and CD-ROMs by fax or e-mail is also provided.

- **Redbooks Web Site** `http://www.redbooks.ibm.com/`

  Search for, view, download, or order hardcopy/CD-ROM redbooks from the redbooks Web site. Also read redpieces and download additional materials (code samples or diskette/CD-ROM images) from this redbooks site.

  Redpieces are redbooks in progress; not all redbooks become redpieces and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

- **E-mail Orders**

  Send orders by e-mail including information from the redbooks fax order form to:

	e-mail address
In United States	usib6fpl@ibmmail.com
Outside North America	Contact information is in the “How to Order” section at this site: `http://www.elink.ibmlink.ibm.com/pbl/pbl/`

- **Telephone Orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	Country coordinator phone number is in the “How to Order” section at this site: `http://www.elink.ibmlink.ibm.com/pbl/pbl/`

- **Fax Orders**

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
Outside North America	Fax phone number is in the “How to Order” section at this site: `http://www.elink.ibmlink.ibm.com/pbl/pbl/`

This information was current at the time of publication, but is continually subject to change. The latest information may be found at the redbooks Web site.

IBM Intranet for Employees

IBM employees may register for information on workshops, residencies, and redbooks by accessing the IBM Intranet Web site at `http://w3.itso.ibm.com/` and clicking the ITSO Mailing List button. Look in the Materials repository for workshops, presentations, papers, and Web pages developed and written by the ITSO technical professionals; click the Additional Materials button. Employees may access `MyNews` at `http://w3.ibm.com/` for redbook, residency, and workshop announcements.

IBM Redbook fax order form

Please send me the following:

Title	Order Number	Quantity

First name Last name

Company

Address

City Postal code Country

Telephone number Telefax number VAT number

☐ Invoice to customer number

☐ Credit card number

Credit card expiration date Card issued to Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

List of abbreviations

AH	Authentication Header
AIS	Advanced Imaging Software
AIX	advanced interactive executive
APAR	Authorized Program Analysis Report
API	application programming interface
ARP	address resolution protocol
ASCII	American National Standard Code for Information Interchange
CBC	cypher block chaining
CDE	common desktop environments
CDMF	Commercial Data Masking Facility algorithm
CD-ROM	compact disk read only memory
CERT	Computer Emergency Response Team
CGI	Common Gateway Interface
CPU	central processing unit
CRAM	Challenge Response Authentication Method
DARPA	Defense Advanced Research Projects Agency
DES	data encryption standard/system
DF	disk free
DMZ	demilitarized zone
DNS	domain name system
DOS	disk operating system
EFM	Enterprise Firewall Management
ERS	Emergency Response Service
ESP	Encapsulating Security Payload
FTP	file transfer protocol
GB	gigabyte
GIF	graphic interchange format
GUI	graphical user interface
HMAC	Hashed Message Authentication Code
HTML	Hypertext Markup Language
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol
IANA	Internet Assigned Number Authority
IBM	International Business Machines
ICMP	Internet control message protocol
ICRA	Internet Content Rating Organization
IETF	Internet Engineering Task Force
IMAP4	Internet Mail Access Protocol
IP	Internet protocol
IPL	initial program load
IPSec	Internet security protocol
ISP	Internet service provider
ITSO	International Technical Support Organization
KB	kilobyte
Kbps	kilobytes per second
KISS	keep it simple and secure
L2F	Layer 2 Forwarding
L2TP	Layer 2 Tunneling Protocol
LAN	local area network
LDAP	Lightweight Directory Access Protocol
LED	light emitting diode
LMTP	Local Mail Transport Protocol
MAC	medium access control
MAS	Nways Multiprotocol Access Services
MD5	Message Digest 5
MIB	management information base
MIME	Multipurpose Internet Mail Extensions
MP	multi-processor
MRS	Nways Multiprotocol Routing Services
MTA	Mail Transport Agent
MTU	maximum transmission unit
MX	mail exchanger

NAPT Network Address Port Translation

NAT Network Address Translation

NNTP NetNews transfer protocol

NSA Network Security Auditor

OS operating system

OSPF open shortest path first (TCP/IP)

PDF portable document format

PICS Platform for Internet Content Selection

PIN personal identification number

PMTU path maximum transmission unit

POP3 Post Office Protocol 3

PPTP purchase pilot test plan

PTF program temporary fix

RTSP flight computer realtime support software

SHA Secure Hash algorithm

SMP symmetric multi-processor

SMTP simple mail transfer protocol

SNMP simple network management protocol

SOA start of authority

SOCKS software common knowledge IR system

SPAM unsolicited e-mail

SPI security parameter index

SQL structured query language

SSH Secure SHell

SSL secure sockets layer

TAP telocator alphanumeric protocol

TCP/IP transmission control protocol/Internet protocol

TFTP trivial file transfer protocol

UDP user datagram protocol

URL universal resource locator

VPN virtual private network

WAIS wide area information servers

WTE Web Traffic Express

Index

Symbols

Numerics

A

B

C

D

E

F

X

IBM Redbook evaluation

A Secure Way to Protect Your Network: IBM SecureWay Firewall for AIX V4.1
SG24-5855-00

Your feedback is very important to help us maintain the quality of ITSO redbooks. **Please complete this questionnaire and return it using one of the following methods:**

- Use the online evaluation form found at `http://www.redbooks.ibm.com/`
- Fax this form to: USA International Access Code + 1 914 432 8264
- Send your comments in an Internet note to `redbook@us.ibm.com`

Which of the following best describes you?
_ Customer _ Business Partner _ Solution Developer _ IBM employee
_ None of the above

Please rate your overall satisfaction with this book using the scale:
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction __________

Please answer the following questions:

Was this redbook published in time for your needs? Yes___ No___

If no, please explain:

What other redbooks would you like to see published?

Comments/Suggestions: (THANK YOU FOR YOUR FEEDBACK!)

SG24-5855-00

Printed in the U.S.A.

A Secure Way to Protect Your Network: IBM SecureWay Firewall for AIX V4.1

SG24-5855-00