CHECK POINT™
Software Technologies Ltd.

Check Point FireWall-1™ White Paper

Version 3.0 — January 1997
P/N 440–3000

http://www.checkpoint.com

Executive Summary

When you connect your local network to the Internet, the single most important measure you can take to prevent break-ins is to define a network Security Policy and establish a firewall to implement that policy. This document describes how to accurately and simply express such a Security Policy in FireWall-1 by presenting a number of typical example configurations. These examples and their design rationale will serve as a guide for your own implementation.

In addition, this document describes the architecture and unique characteristics of Check Point Software Technologies Ltd.'s FireWall-1 Internet Gateway, and outlines the major characteristics that enable Check Point FireWall-1 to establish full, transparent, and true connectivity using the entire range of Internet protocols, while ensuring network security. Encryption, User, Client and Session Authentication and other FireWall-1 features and techniques are described. Finally, performance data are presented.

The powerful combination of Check Point FireWall-1's sophisticated Stateful Inspection technology and its intuitive GUI deliver unmatched security, connectivity and performance.

In This Document:

Copyrights and Trademarks

Check Point Software Technologies Ltd.

International Headquarters:
3A Jabotinsky
Ramat Gan 52520, Israel
Tel: 972-3-613 1833
Fax: 972-3-575 9256

U.S. Headquarters:
400 Seaport Court, Suite 105
Redwood City, CA 94063
Tel: 800-429-4391
415-562-0400
Fax: 415- 562-0410

e-mail: info@checkpoint.com

HTTP://www.checkpoint.com

What’s New in FireWall-1 Version 3.0?

Content Security

This feature provides the Security Manager with finely tuned control tools for Web, Mail and FTP connections, including Anti-Virus checking for transferred files, access control for specific network resources (e.g. URLs, files etc.) and SMTP commands. The new level of security is integrated with the familiar FireWall-1 features, defined through the Security Policy and monitored via the Log Viewer to enable full auditing capabilities.

Active Network Management

Completely integrated with network security, this feature provides complete real-time control of the network configuration, including accounting, live connections monitoring, load balancing and exporting log records to an Informix database.

Synchronization and High Availability of Firewall Modules

Starting with Version 3.0, different FireWall Modules running on different machines can share state information and mutually update each other. In addition, synchronized FireWalls can take over from each other if one of them goes down.

Load Balancing

The Load Balancing feature enables distributing a processing load across any number of servers.

Live Connections

System administrators can now keep track of live connections through the FireWalled gateways and hosts using the familiar Log Viewer interface.

Accounting

Log Viewer entries can now show accounting information, for example, session length and number of bytes transferred.

Authentication

Transparent User Authentication

Users need no longer initiate separate authentication sessions on the gateway in order to be authenticated. Instead, FireWall-1 intercepts FTP, HTTP, TELNET and RLOGIN sessions passing through the gateway and folds them into the appropriate FireWall-1 Security Server.

Session Authentication

Session Authentication can be used to authenticate any service on a per-session basis. Authentication is performed by a user-written application that communicates with the FireWall Module using the FireWall-1 Session Authentication Agent Protocol. The Session Authentication Agent can run on any computer.

Across-the-Board Client/Server Architecture

The addition of X/Motif Client GUI extends the FireWall-1 Client/Server architecture to all supported Unix platforms.

Enhanced Encryption and Authentication Transparent authentication and support for state-of-the-art encryption and authentication standards (IPSec, SKIP, and RADIUS) maintain FireWall-1’s preeminent position at the forefront of network management and security technology.

Address Translation

Graphical User Interface Windows and X/Motif GUI users can now define Address Translation rules using the familiar Rule Base paradigm.

Automatic Rule Generation Address Translation rules can now be generated automatically on all platforms.

Miscellaneous New Features

Time Objects Time objects — defining time periods — can now be used to enforce rules at specific times-of-day.

Support for New Services Support for for a variety of new services, including SQL*Net, CoolTalk and Microsoft NetMeeting further extends FireWall-1’s impressive list of over 120 out-of-box supported services.

Single Interface Load for Bay Networks Routers FireWall-1 Version 3.0 enables downloading Access Control Lists to individual interfaces (or sets of interfaces) of a Bay Networks router.

Internet Firewall Technologies

Overview

When you connect your network to the Internet, securing your network against intrusion is of critical importance. The most effective way to secure the Internet link is to put a firewall system between the local network and the Internet. The firewall ensures that all communication between an organization’s network and the Internet conforms to the organization’s Security Policy.

In order to effectively provide real security, a firewall must track and control the flow of communication passing through it. To reach control decisions for TCP/IP based services (e.g., whether to pass, reject, encrypt or log communication attempts), a firewall must obtain, store, retrieve and manipulate information derived from all communication layers and from other applications.

It is not sufficient to examine packets in isolation. State information — derived from past communications and other applications—is an essential factor in making the control decision for new communication attempts. Both the communication state (derived from past communications) and the application state (derived from other applications) may be considered when making control decisions.

To summarize, control decisions require that the firewall be capable of accessing, analyzing and utilizing the following:

1. **communication information** — information from all seven layers in the packet
2. **communication–derived state** — the state derived from previous communications

   For example, the outgoing PORT command of an FTP session could be saved so that an incoming FTP data connection can be verified against it.
3. **application–derived state** — the state information derived from other applications

   For example, a previously authenticated user would be allowed access through the firewall for authorized services only.
4. **information manipulation** — the evaluation of flexible expressions based on all the above factors

Comparing Alternatives

The extent to which the available firewall technologies provide these four essential firewall capabilities is described in this section.

Routers

Routers operate at the network level and their most obvious shortcoming is their inability to provide security even for the most basic services and protocols. Routers are not secure, since they do not provide essential firewall capabilities:

1. **communication information** — routers can access only a limited part of the packet’s header information

2. **communication and application derived state** — routers are stateless, in that they cannot keep communication-derived state or application-derived state information
3. **information manipulation** — routers have a very limited capability to manipulate information.

In addition, routers are difficult to configure, monitor and manage, and do not provide adequate logging and alerting mechanisms.

Proxies

Proxies are an attempt to implement firewalls at the application level. Their primary advantage is their statefulness. Proxies provide partial communication-derived state, full application-derived state information and partial communication information. Proxies are also capable of processing and manipulating information.

However, there are obvious disadvantages in using application level proxies as firewalls, including:

1. **limited connectivity** — each service needs its own proxy, so the number of available services and their scalability are limited
2. **limited technology** — application gateways cannot provide proxies for UDP, RPC and other services from common protocol families
3. **performance** — application level implementation entails a discernible performance penalty

In addition, proxies are vulnerable to OS and application level bugs, overlook information contained in lower layers, and in the case of traditional proxies, are rarely transparent.

Historically, application level gateways suited the Internet’s common uses and needs. However, as the Internet evolved into a dynamic environment that constantly offers new protocols, services and applications, proxies are no longer able to handle the diverse types of communication on the Internet, or to fulfill the new business needs, high bandwidth and security requirements of networks.

FireWall-1 Stateful Inspection Technology

FireWall-1's Innovative Technology

In contrast, Check Point FireWall-1 introduces an innovative architecture called Stateful Inspection technology, which implements all the necessary firewall capabilities at the network level.

Employing the Stateful Inspection technology, FireWall-1's Inspection Module accesses and analyzes data derived from all communication layers. This "state" and "context" data is stored and updated dynamically, providing virtual session information for tracking connectionless protocols (e.g. RPC and UDP-based applications). Cumulative data from the communication and application states, network configuration and security rules, are used to generate an appropriate action, either accepting, rejecting or encrypting the communication. Any traffic not explicitly allowed by the security rules is dropped by default and real-time security alerts are generated, providing the system manager with complete network status.

Table 1 Technology Comparison

firewall capability	routers	proxies	Stateful Inspection
communication information	Partial	Partial	Yes
communication-derived state	No	Partial	Yes
application-derived state	No	Yes	Yes
information manipulation	Partial	Yes	Yes

To summarize, FireWall-1 combines network–level transparency, comprehensiveness, robustness and high performance with application-level flexibility to provide a superior security solution that goes far beyond the capabilities of previous solutions.

Security Policy

A FireWall-1 Security Policy is expressed in terms of a Rule Base and Properties. The Rule Base is an ordered set of rules against which each communication is tested. If a communication's source, destination and service match a rule, the action indicated in the rule (Accept, Encrypt, Reject, Drop) is performed. If a communication does not match any of the rules, then it is dropped, in accordance with the principle of "That Which Is Not Expressly Permitted is Prohibited."

FireWall-1 Inspection Module

The FireWall-1 Inspection Module is dynamically loaded into the operating system kernel, between the Data Link and Network layers (layers 2 and 3). When the first packet of a new connection arrives, the FireWall-1 Inspection

Module examines the Rule Base to determine whether or not the connection is to be allowed. Once a connection is established, FireWall-1 adds the connection to an internal connections table. For efficiency considerations, subsequent packets of the connection are verified against the connections table rather than against the Rule Base. The packet is allowed to pass only if the connection is listed in the connections table.

In a connection such as this, the entire connection is handled by the FireWall-1 Inspection Module (Figure 1).

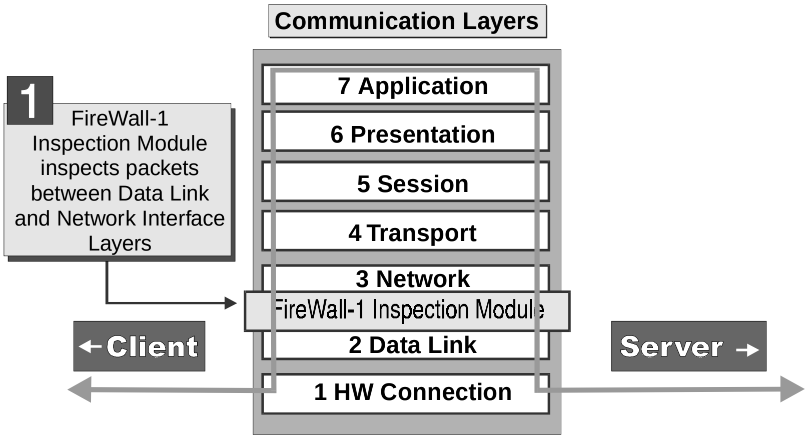
Communication Layers
1
FireWall-1 Inspection Module inspects packets between Data Link and Network Interface Layers
7 Application
6 Presentation
5 Session
4 Transport
3 Network
FireWall-1 Inspection Module
2 Data Link
1 HW Connection
Client
Server

Figure 1 A connection handled by the FireWall-1 Inspection Module

Examples

UDP

From the information in the TCP/UDP header, FireWall-1 utilizes its unique capabilities to build a communications–derived state, on which basis UDP communications are tracked and controlled (See “UDP (User Datagram Protocol)” on page 43.).

FTP

In order to track the reverse FTP-data connection, FireWall-1 extracts information from packet’s application portion. Its unique ability to utilize communication information at all layers enables FireWall-1 to build a communication–derived state, on which basis the reverse connection is allowed to pass (See “Outbound FTP Connections” on page 42.).

RPC

FireWall-1 uses all the above capabilities, including application–derived state, to track this complex protocol’s dynamic mapping between program numbers and port numbers (See “RPC (Remote Procedure Call)” on page 44.).

Security Server Authentication and Content Security

FireWall-1 enables the system administrator to define a Security Policy on a per-user basis, where not only a communication’s source, destination and service are verified, but the user is authenticated. In addition, communications can be allowed or disallowed based on their content. For example, mail to or from certain addresses can be rejected or redirected, access can be denied to specific URLs, and anti-virus checking of transferred files can be performed. Authentication and Content Security are provided by a suite of FireWall-1 Security Servers, running at the application layer (Figure 2).

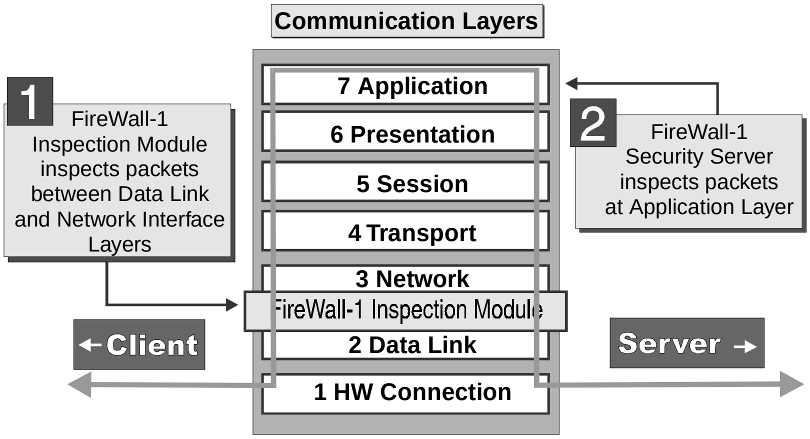
Communication Layers
7 Application
6 Presentation
5 Session
4 Transport
3 Network
FireWall-1 Inspection Module
2 Data Link
1 HW Connection
1
FireWall-1 Inspection Module inspects packets between Data Link and Network Interface Layers
2
FireWall-1 Security Server inspects packets at Application Layer
Client
Server

Figure 2 A connection mediated by a FireWall-1 Security Server

When a FireWall-1 Security Server is invoked, the Inspection Module diverts all the packets in the connection to the Security Server, which performs the required authentication and/or Content Security inspection.

There are five FireWall-1 Security Servers, as described in Table 2.

Table 2 FireWall-1 Security Servers — features

Server	Authentication	Content Security	Comments
TELNET	yes	no	
RLOGIN	yes	no	
FTP	yes	yes	
HTTP	yes	yes	
SMTP	no	yes	secure sendmail

Load Balancing

FireWall-1 enables enables distributing a processing load across any number of servers. Figure 3 depicts a configuration in which FTP and HTTP services are provided by a number of servers.

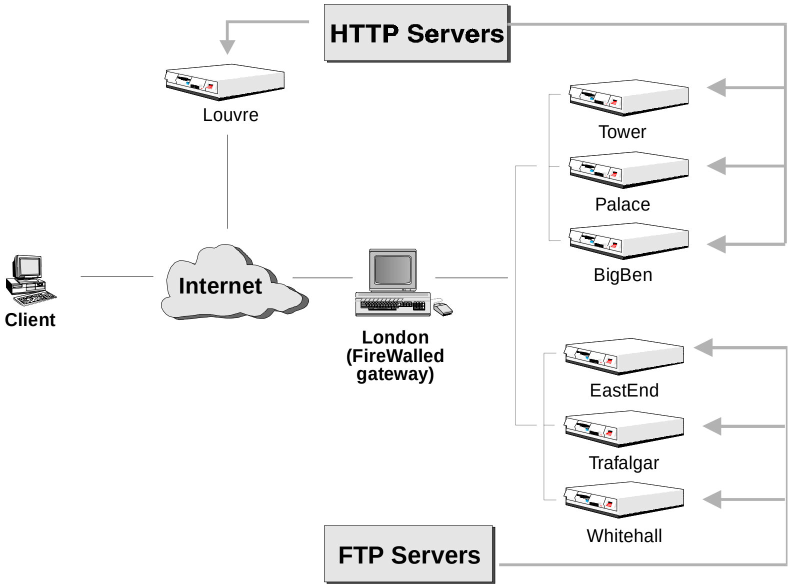

Figure 3 Load Balancing among several servers

All the HTTP servers can provide the HTTP client with the same services (though not all of the HTTP servers are behind the FireWalled gateway). In the same way, all the FTP servers can provide the FTP client with the same services.

The system administrator wishes to ensure that the service load is balanced among the servers. The client will be unaware of the different servers. From the client's point of view, there is only one HTTP server and only one FTP server. When the service request reaches the FireWall, FireWall-1 determines which of the servers will fulfill the request, based on the load balancing algorithm specified by the system administrator.

Load Balancing Algorithms

The available load balancing algorithms are:

1. **server load**

   FireWall-1 queries the servers to determine which is best able to handle the new connection. There must be a load measuring agent on the server.

2. **round trip**

   FireWall-1 uses PING to determine the round-trip times between the FireWall and each of the servers, and chooses the server with the shortest round trip time.

3. **round robin**

   FireWall-1 simply assigns the next server in the list.
4. **random**

   FireWall-1 assigns a server at random.
5. **domain**

   FireWall-1 assigns the "closest" server, based on domain names.

Authentication

FireWall-1 provides three kinds of authentication:

1. User Authentication, which enables an administrator to grant specific users special access privileges

   The User Authentication feature includes a FireWall-1 HTTP Security Server, which provides a mechanism for authenticating users of HTTP services, both incoming and outgoing.
2. Client Authentication, which provides a mechanism for authenticating users of any application, standard or custom.
3. Session Authentication, which provides a transparent per-session authentication mechanism that can be integrated with any application

Table 3 compares the features of the three FireWall-1 Authentication types.

Table 3 Comparison of Authentication Types

	User	**Client**	**Session**
services	TELNET, FTP, RLOGIN, HTTP	all services	all services
authentication performed once per	session	IP address (multiple sessions) in a separate non-transparent authentication session	session
transparent (user initiates session directly to server)?	yes	yes, but only after non-transparent authentication session	yes

Note – Client and Session Authentication are not provided by the Security Servers, but by other mechanisms described in "Client Authentication" on page 35 and "Session Authentication" on page 35.

FireWall-1 supports the following authentication schemes for each user:

1. **S/Key** — The user is challenged to enter the value of requested S/Key iteration.
2. **SecurID** — The user is challenged to enter the number displayed on the Security Dynamics SecurID card.

3. **OS Password** — The user is challenged to enter his or her OS password.
4. **Internal** — The user is challenged to enter his or her internal FireWall-1 password on the gateway.
5. **RADIUS** — The user is challenged for the response, as defined by the RADIUS server.
6. **AssureNet Pathways** — The user is challenged for the response, as defined by the AssureNet Pathways server.

Content Security

Content Security extends the scope of data inspection to the highest level of a service's protocol, achieving highly tuned access control to network resources.

FireWall-1 provides content security for HTTP, SMTP and FTP, using the Firewall-1 Security Servers. For each connection established through the FireWall-1 Security Servers, the Security Administrator is able to control specific access according to fields that belong to the specific service: URLs, file names, FTP PUT/GET commands, type of requests and more.

A major security enhancement enabled by the Content Security feature is Anti-Virus checking for files transferred. The Anti-Virus option is fully integrated with all FireWall-1 features.

Content Security is enabled by a FireWall-1 object of type Resource. A FireWall-1 Resource specification defines a set of entities which can be accessed by a specific protocol. You can define a FireWall-1 Resource based on HTTP, FTP and SMTP.

For example, you might define a URI resource whose attributes are a list of URLs and the HTTP and FTP schemes. The resource can be used in the Rule Base in exactly the same way a service can be used, and the standard logging and alerting methods are available to provide monitoring of resource usage.

HTTP

URI resources can define schemes (HTTP, FTP, GOPHER etc.), methods (GET, POST, etc.), hosts (for example, "*.com"), paths and queries. Alternatively, a file containing a list of IP addresses and servers can be specified.

SMTP

The SMTP protocol, designed to provide maximum connectivity between people all over the Internet, and enhanced to support not only E-mails but file attachments, poses a challenge to the Security Administrator who wants to maintain connectivity but keep intruders out of the internal networks.

FireWall-1 offers an SMTP server that provides highly granular control over SMTP connections. By defining SMTP resources, the Security Administrator can:

- hide outgoing mail’s **From** address behind a standard generic address that conceals internal network structure and real internal users
- redirect mail sent to given **To** addresses (for example, to root) to another mail address
- drop mail from given addresses
- strip attachments of given types from mail
- strip the **Received** information from outgoing mail, in order to conceal internal network structure
- drop mail messages above a given size

FTP

The FTP Security Server provides authentication services, and Content Security based on FTP commands (PUT/GET), file name restrictions, and Anti-Virus checking for files.

Anti-Virus

Anti-Virus inspection is an integral component of FireWall-1’s Content Security feature, and considerably reduces the vulnerability of protected hosts.

Using an integrated Anti-Virus module, the Anti-Virus option examines all files transferred. This option’s configuration (which files to check, how to handle infected files) is integrated into the Security Policy. All FireWall-1 auditing tools are available for logging and alerting when these files are encountered.

Configuring FireWall-1

This section demonstrates how to build a FireWall-1 Rule Base to implement a Security Policy for a simple network configuration using a “diode” policy, and then for a larger network using a more detailed configuration.

A Simple Configuration

The following example illustrates how to deploy FireWall-1 in the network configuration shown in the diagram below. Note that this is not a recommended configuration, but simply an example for the purposes of this document.

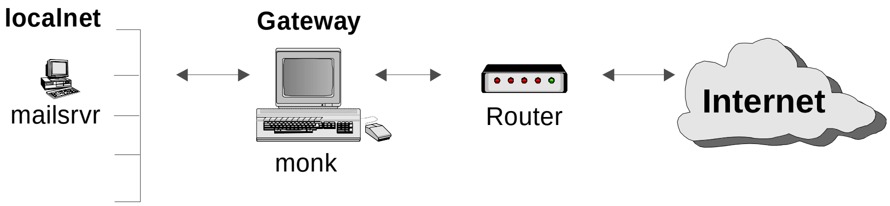

Figure 4 Example Network Configuration

In this example, FireWall-1 will be installed on the gateway computer (named “monk” in the rules that follow).

A Typical Security Policy

For the configuration shown above, a typical Security Policy might be this:

- external networks may access the local network only to send mail to local computers
- local computers may access the entire network: localnet and Internet

This policy protects the private network from non-local networks, but puts no restrictions on local computers.

You will begin by considering how this Security Policy can be implemented in FireWall-1. Next, you will see how this Security Policy can be “tightened up” so that the potential loopholes are blocked.

Implementing a Security Policy

In order to implement a Security Policy, you must perform the following actions:

1. Define the network objects used in the Rule Base.

   You do not have to define the entire network to FireWall-1 — only those objects that are used in the Rule Base. For the configuration described here, you must define the gateway (monk), the mail server (mailsrvr) and the local network (localnet).
2. Define services used in your Security Policy (optional).

You do not have to define the commonly used services. These are already defined for you in FireWall-1. You should define only those uncommon services that are part of your Security Policy.

Defining network objects and services is very straightforward. In most cases, you need only specify a name, because FireWall-1 can obtain the object’s properties from the appropriate system databases (DNS, `hosts`, etc).

3. Define the Rule Base — the rules for accepting, rejecting and logging packets.
4. Install the Rule Base — install the Inspection Code on the gateways.

The next step in implementing a Security Policy is to define the Rule Base, using the Rule Base editor. In the Rule Base editor, rules are expressed in terms of the following elements:

The first four elements describe the communication attempt.

Source — where the communication is coming from

Destination — where the communication is going

Services — the type of application

Time — the time of day of the start of the communication (new in Version 3.0)

Note – If you specify “Any” under services, all TCP, UDP and RPC based applications, even those not defined in the Service Manager, are included. Even if you use an undefined database application, FireWall-1 will secure the outbound connection and ensure that replies are passed through but nothing else.

The next two elements indicate what is to be done.

Action — what is to be done with the given communication attempt

Track — whether to log the packet or to generate an alert

The last element indicates which FireWall Module will enforce the rule defined by the first five elements.

Install On — the FireWall Module that will enforce this rule

In this simple example, there are only two rules, corresponding to the policy given above.

The first rule (non-local networks may only send mail to the mail server) can be expressed in the Rule Base editor as follows:

Source	Destination	Services	Action	Track	Install On
Any	mailsrvr	smtp	Accept	Short Log	Gateways

The second rule (local computers may access the entire network: localnet and Internet) can be expressed in the Rule Base editor as follows:

Source	Destination	Services	Action	Track	Install On
localnet	Any	Any	Accept	Short Log	Gateways

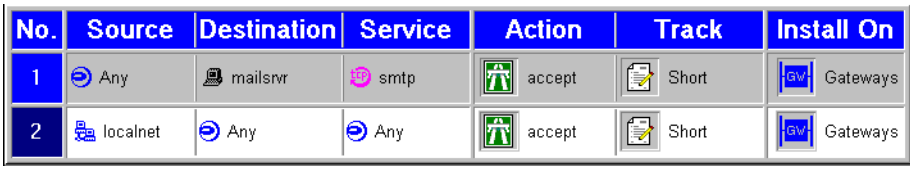

No.	Source	Destination	Service	Action	Track	Install On
1	Any	mailsrvr	smtp	accept	Short	Gateways
2	localnet	Any	Any	accept	Short	Gateways

Figure 5 Rule Base editor window showing first two rules

Implicit Drop

FireWall-1 follows the principle "That Which Is Not Expressly Permitted is Prohibited." To enforce this principle, FireWall-1 implicitly adds a rule at the end of the Rule Base that drops all communication attempts not described by the other rules.

Since the rules are examined sequentially for each packet, only packets not described by the first two rules are examined by the implicit rule. However, if you rely on the implicit rule to drop these packets, they will not be logged, because only packets which are described by a rule can be logged. To log these packets, you must explicitly define a "none of the above" rule, as follows:

Source	Destination	Services	Action	Track	Install On
Any	Any	Any	Reject	Long Log	Gateways

If you do not explicitly define such a rule, FireWall-1 will implicitly define one for you, and the packets will be dropped. In no case will FireWall-1 allow these packets to pass. The advantages of defining such a rule explicitly are:

- You can then specify logging for the rejected packets.
- You can tune your Security Policy.
- You can better understand your network behavior.

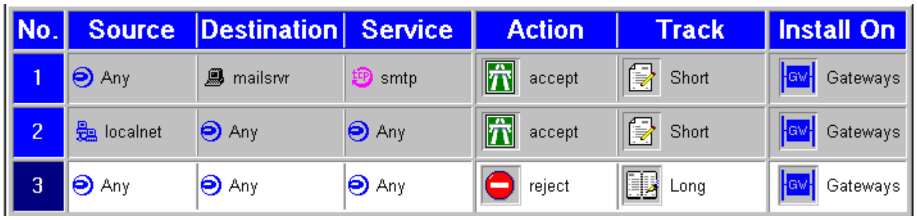

No.	Source	Destination	Service	Action	Track	Install On
1	Any	mailsrvr	smtp	accept	Short	Gateways
2	localnet	Any	Any	accept	Short	Gateways
3	Any	Any	Any	reject	Long	Gateways

Figure 6 A Complete Rule Base

"Stealthing" the Gateway

The Rule Base described above has a shortcoming — it enables localnet computers to get on the gateway (assuming that they have Unix accounts and passwords). In the given configuration, this is usually not desirable. To prevent any computers (not just local computers) from getting on the gateway, you must add a rule (before the other rules) as follows:

Source	Destination	Services	Action	Track	Install On
Any	monk	Any	Reject	Long Log	Gateways

Protecting the gateway in this manner makes it inaccessible to other users and applications (except for FireWall-1 management purposes). The gateway becomes an invisible network object that, from the point of view of the network, does not even exist.

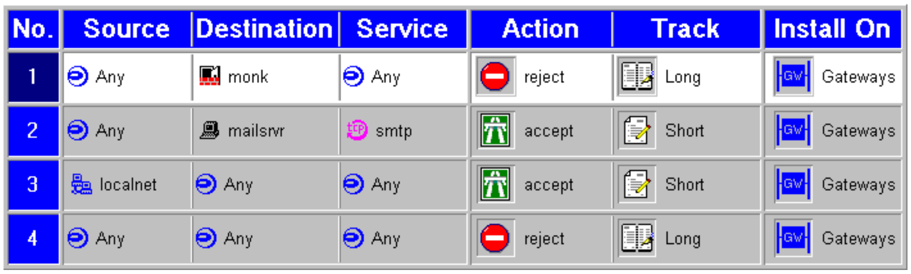

No.	Source	Destination	Service	Action	Track	Install On
1	Any	monk	Any	reject	Long	Gateways
2	Any	mailsrvr	smtp	accept	Short	Gateways
3	localnet	Any	Any	accept	Short	Gateways
4	Any	Any	Any	reject	Long	Gateways

Figure 7 Rule Base editor window with rule protecting the gateway

You may wish to confirm that the gateway is indeed inaccessible by performing the following experiment (after you have installed the Rule Base). TELNET out through the gateway and then try to TELNET back to the gateway. The second TELNET will fail.[1]

Note – The icons representing network objects convey information about the objects. Gateways look like gates, and a brick wall inside a host or gateway icon indicates that the object is FireWalled.

To further protect the gateway, you may wish to add a rule that rejects any packet which originates on the gateway, as follows:

Source	Destination	Services	Action	Track	Install On
monk	Any	Any	Reject	Long Log	Source

[1] You should be aware that if you ping the gateway (rather than TELNET back) from the outside, you will succeed, because the Accept ICMP property is specified by default as the first rule in the Properties Setup window.

The reason that Install On is specified as Source in this rule is that by default, gateways enforce rules on inbound traffic only. The rule is enforced on monk, but because monk is specified as Source, the rule is enforced only in the outbound direction (that is, it applies only to packets leaving monk which also originate on monk). If Install On were specified as Gateways, then the rule would apply only to packets entering monk which also originate on monk, in other words, to no packets at all.

A More Detailed Configuration

Consider the following configuration:

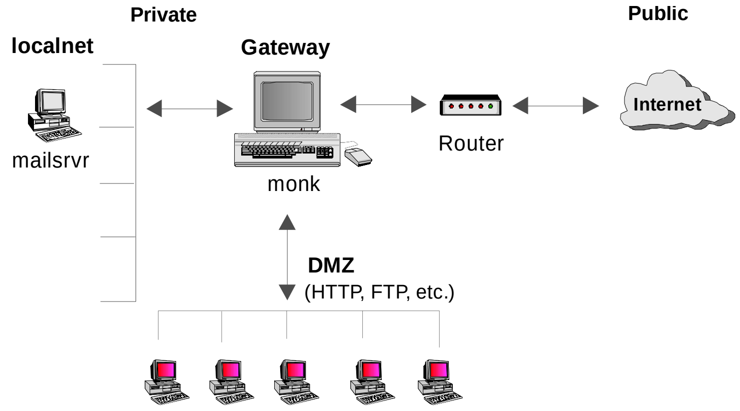

Figure 8 Example Network Configuration

This configuration is similar to the first one, except that a public server (DMZ — Demilitarized Zone) has been added. DMZ provides HTTP, FTP and other services to non-local networks, but does not initiate any traffic. DMZ is actually a third interface attached to the gateway, and might be a network, a sub-network or a host.

The Security Policy for this configuration is the same as the Security Policy for the previous configuration, with the following additional rule:

Source	Destination	Services	Action	Track	Install On
Any	DMZ	HTTP, FTP	Accept	Short Log	Gateways

Figure 9 Rule Base with DMZ added

Note that under Install On, you can specify a network object either by name or by function.

Table 4 Install On

Install On	Meaning
GW	Enforce on all network objects defined as gateways in the direction specified in the "Apply Gateway Rules to Interface Direction" property in the Control Properties/Security Policy window.
D	Enforce in the inbound direction on the FireWalled network objects defined as Destination (typically servers) in this rule.
S	Enforce in the outbound direction on the FireWalled network objects defined as Source (typically clients — initiators of traffic) in this rule.
R	Enforce on the appropriate interfaces on all routers, using FireWall-1's Auto-Scoping feature.
specific target	Enforce on the specified target object(s) only, in the inbound and outbound (eitherbound) directions.

If you specify Gateways, the rule is enforced on all the hosts that are defined as gateways (in the Workstation Properties window). The rule is enforced in the direction specified in "Apply Gateway Rules to Interface Direction" property in the Security Policy tab of the Properties Setup window.

If you specify routers, the rule is enforced on the appropriate interfaces on all routers, using FireWall-1's auto-scoping feature.[2] FireWall-1 generates an Access List for the router. It should be noted that with Access Lists only a subset of FireWall-1's FireWall Module functionality can be implemented. For example, it is not possible to secure FTP connections as described in "Outbound FTP Connections" on page 42.

If you specify an object by name, then the rule is enforced for both incoming and outgoing packets (eitherbound).

[2] If you are enforcing a Security Policy on a gateway, there is no need for Access Lists on routers. Routers are in general considered to be less secure and also have limited logging and alert capabilities, so it is recommended that they be left open and that the Security Policy be implemented on other network objects.

Authentication

After successful authentication, access privileges are determined by the user's individual properties, as defined in the User Properties window.

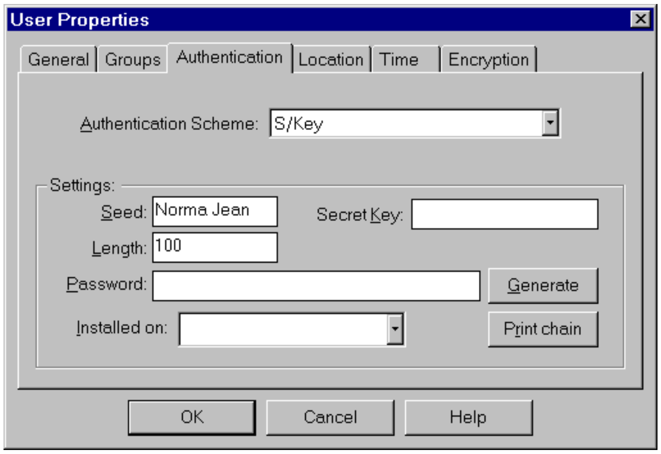

Authentication Scheme: S/Key

Settings:

Seed: Norma Jean

Secret Key:

Length: 100

Password:

Installed on:

Figure 10 Authentication Method for User

These properties include allowed sources, destinations and time-of-day.

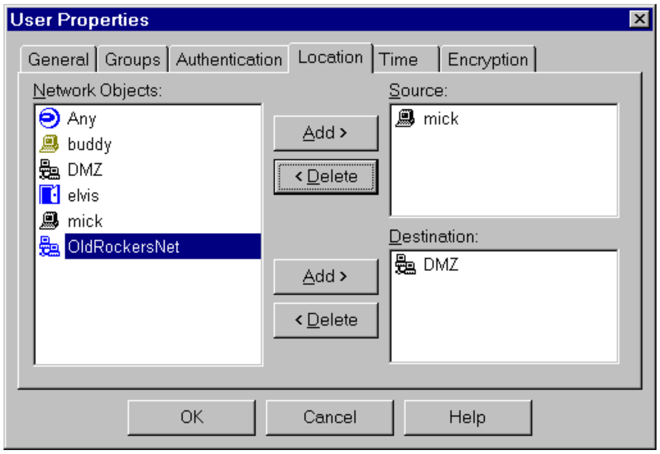

Network Objects:
- Any
- buddy
- DMZ
- elvis
- mick
- OldRockersNet

Source:
- mick

Destination:
- DMZ

Figure 11 Allowed Sources and Destinations for User

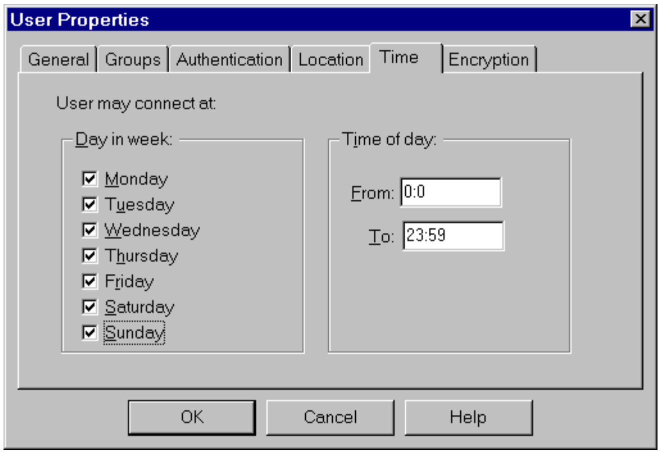

User may connect at:

Day in week:
- ☑ Monday
- ☑ Tuesday
- ☑ Wednesday
- ☑ Thursday
- ☑ Friday
- ☑ Saturday
- ☑ Sunday

Time of day:
From: 0:0
To: 23:59

Figure 12 Allowed Time-of-Day for User

Anti-Spoofing

Spoofing is a technique where an intruder attempts to gain unauthorized access by altering a packet's IP address to make it appear as though the packet originated in a part of the network with higher access privileges. For example, a packet originating on the Internet may be disguised as a localnet packet. If undetected, this packet might then have unrestricted access to the localnet.

FireWall-1 has a sophisticated anti-spoofing feature which detects such packets by requiring that the interface on which a packet enters the gateway correspond to its IP address. For example, FireWall-1 would identify a packet entering the gateway from the Internet which carried a localnet IP address, and would reject that packet and issue an alert. This type of effective anti-spoofing can only be achieved by an integrated system such as FireWall-1 that has access to both the interface through which the packet arrived and the IP address it claims to have, and that is capable of logging and alerting these events.

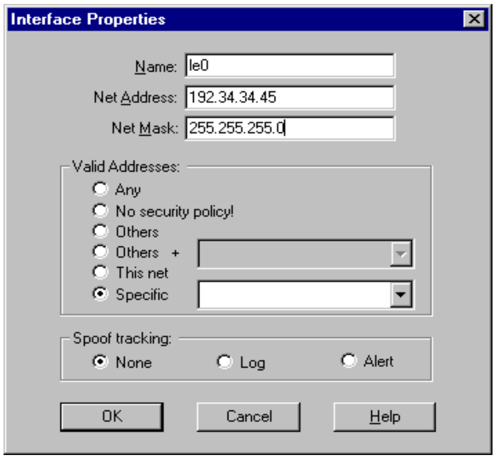

Name: le0
Net Address: 192.34.34.45
Net Mask: 255.255.255.0

Valid Addresses:
- ○ Any
- ○ No security policy!
- ○ Others
- ○ Others +
- ○ This net
- ● Specific

Spoof tracking:
● None ○ Log ○ Alert

Figure 13 Spoof Tracking defined

Anti-spoofing is defined in the property window of the network object which enforces it, for example, in the gateway’s or router’s Workstation Properties window (OpenLook GUI - Figure 14) or in the Interface Properties window (Windows GUI - Figure 13 on page 21).

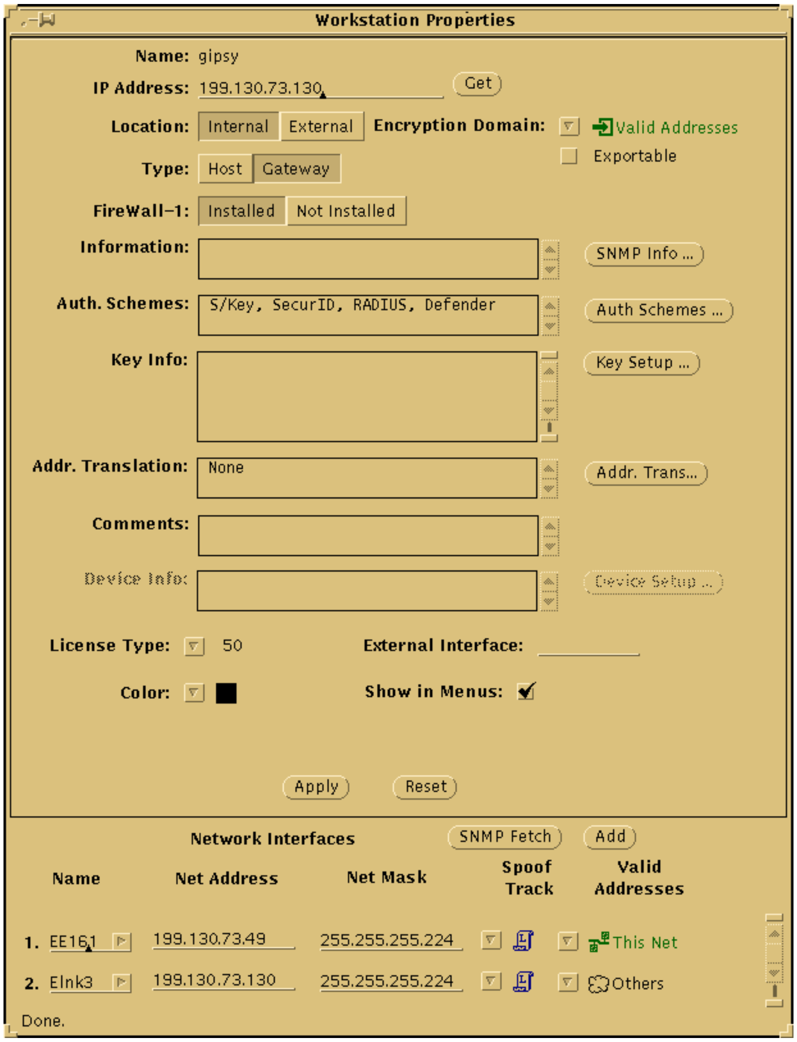

Name: gipsy

IP Address: 199.130.73.130

Location: Internal External

Encryption Domain: Valid Addresses

□ Exportable

Type: Host Gateway

FireWall-1: Installed Not Installed

Information:

Auth. Schemes: S/Key, SecurID, RADIUS, Defender

Key Info:

Addr. Translation: None

Comments:

Device Info:

License Type: 50

External Interface:

Color: ■

Show in Menus: ✓

Network Interfaces

	Name	Net Address	Net Mask	Spoof Track	Valid Addresses
1.	EE161	199.130.73.49	255.255.255.224		This Net
2.	Elnk3	199.130.73.130	255.255.255.224		Others

Done.

Figure 14 Workstation Properties window (OpenLook GUI)

Logging and Alerting

You can record events, including acceptance and rejection of packets, in a log. The FireWall-1 Log Viewer (Figure 24 on page 31) enables you to examine the log, filtering and searching the log in a variety of different ways, so that you can quickly and efficiently extract the information you need.You can issue an alert when FireWall-1 rejects or accepts a packet.

The alert can take several forms. You can display a message on the master console, or you can send a mail message to some pre-determined address, or you can issue an SNMP trap. In fact, you can specify any OS command to be executed. All alerts are also logged.

In addition, you can choose to view accounting log entries or active (live) connections by making the appropriate selection from the drop down list in the Log Viewer toolbar.

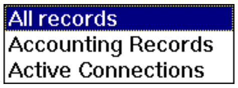

Figure 15 Selecting the type of log entries to display

Installing a Rule Base

When you install a Rule Base on the FireWall Modules which are to enforce it, FireWall-1 verifies that the Rule Base is logical and consistent, and then generates and downloads the Inspection Code to the specified FireWall Modules. If the specified network object is a router, FireWall-1 generates and downloads the appropriate access list rather than Inspection Code.

Properties

A Security Policy is defined not only by the Rule Base, but also by parameters specified in the Properties Setup window.

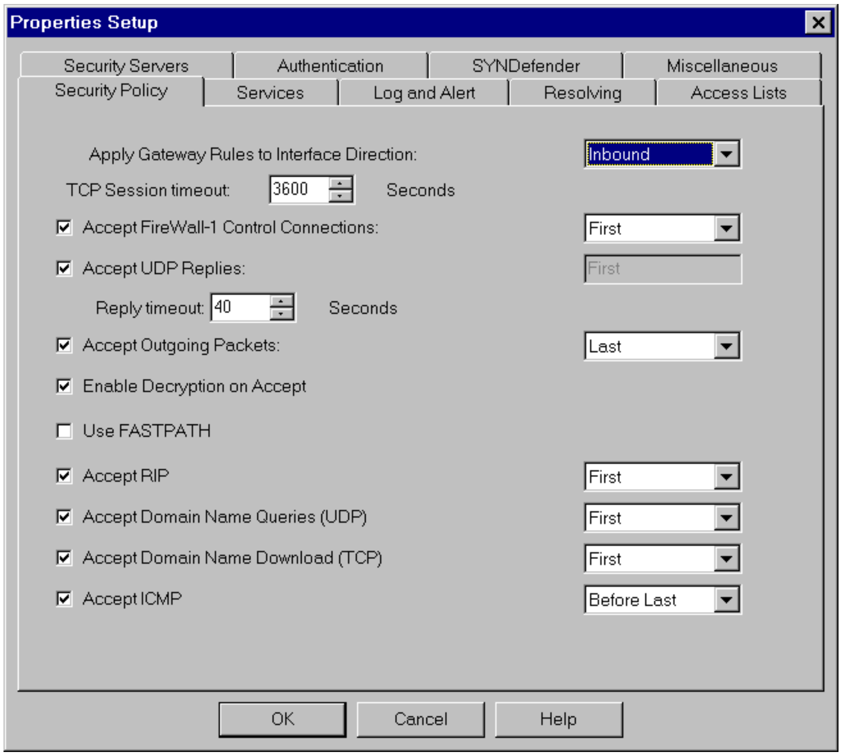

Apply Gateway Rules to Interface Direction: Inbound

TCP Session timeout: 3600 Seconds

☑ Accept FireWall-1 Control Connections: First

☑ Accept UDP Replies: First

Reply timeout: 40 Seconds

☑ Accept Outgoing Packets: Last

☑ Enable Decryption on Accept

□ Use FASTPATH

☑ Accept RIP First

☑ Accept Domain Name Queries (UDP) First

☑ Accept Domain Name Download (TCP) First

☑ Accept ICMP Before Last

Figure 16 Properties Setup window - Security Policy tab

These parameters enable the user to control all aspects of a packet's inspection, while at the same time freeing the user of the need to specify repetitive detail in the Rule Base Editor.

Principles of Operation

Introduction

This section describes the architecture and unique characteristics of the Check Point FireWall-1 Internet gateway, and outlines the major characteristics that enable Check Point FireWall-1 to establish full, transparent, and true Internet connectivity using the entire range of Internet protocols, while ensuring network security. In addition, Authentication, Encryption and other FireWall-1 features and techniques are described. Finally, performance data are presented.

Note – In FireWall-1 terminology, the term "gateway" is used to describe a computer used primarily to route traffic coming into and leaving a network. In some literature, the term "router" is used to describe a gateway. In FireWall-1 terminology, "router" means a Cisco or Bay Networks router.

FireWall-1 Architecture

FireWall-1 is comprised of two primary modules:

- Control Module

  The Control Module includes the GUI and the Management Module.

  The GUI is the front end to the Management Module, which manages the FireWall-1 database: the Rule Base, network objects, services, users etc.
- FireWall Module

  The FireWall Module includes the Inspection Module, FireWall-1 daemons and Security Servers.

  The FireWall Module implements the Security Policy, logs events, and communicates with the Control Module using the daemons.

A FireWall-1 Security Policy is defined in terms of network objects, services, users, and the rules that govern the interactions between them. Once these have been specified with the Control Module, Inspection Code is generated and then installed on the FireWalls that will enforce the Security Policy.

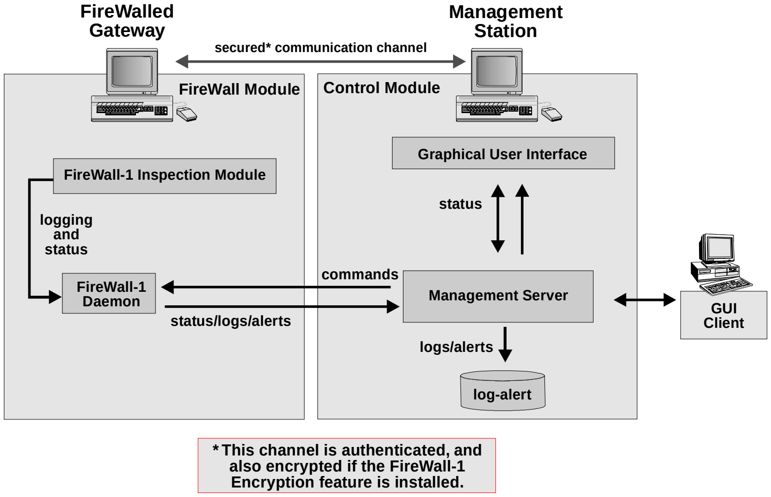

Figure 17 FireWall-1 Components

Client/Server Model

The Control Module can also be deployed in a Client/Server configuration. The client can run either a Windows 95, Windows NT or X/Motif Graphic User Interface, and controls a server running on any of the supported platforms (Windows NT, SunOS 4.1.3, Solaris 2.x, HP-UX 9 and 10).

The client interacts with the user via the GUI, but all the data (the database and configuration files) is maintained on the server. The client’s connection to the Management Module is mediated by the server. For example, if the client requests that the Security Policy be installed, then the server passes the request to the Management Module, which carries it out and notifies the server of the result. The server in turn notifies the client.

As in previous versions of FireWall-1, an OpenLook GUI is also available. The OpenLook GUI includes both the server and the client and can be installed on any of the supported OpenLook platforms (SunOS 4.1.3 and higher, Solaris 2.x, HP-UX 9 and 10).

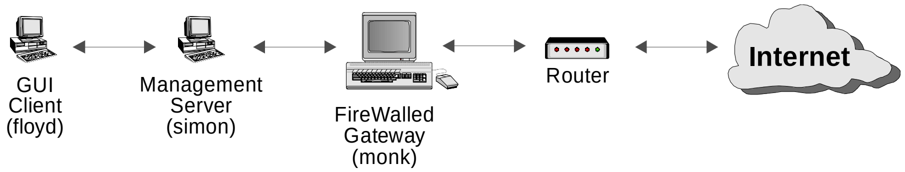

Figure 18 FireWall-1 Client/Server configuration

In the configuration depicted in Figure 18, the functionality of the Control Module is divided between two workstations (simon and floyd). The

Management Module, including the FireWall-1 database is on simon, the server. The GUI is on floyd.

The user, working on floyd, maintains the FireWall-1 Security Policy and database, which reside on simon. The FireWall Module is installed on monk, the FireWalled gateway, which enforces the Security Policy and protects the network.

The connection between the client and server is secured, enabling true remote management. For example, floyd can be situated outside the local network. In this way, several sites can be managed through one Control Module.

The Control Module GUI client, the Control Module server and the FireWall Module can be installed on different computers, or on the same computer if its platform supports all three components. The system administrator uses the Control Module on the Management Center to define the Security Policy, but it is the FireWalled Gateway (where the FireWall Module is installed) that enforces the Security Policy.

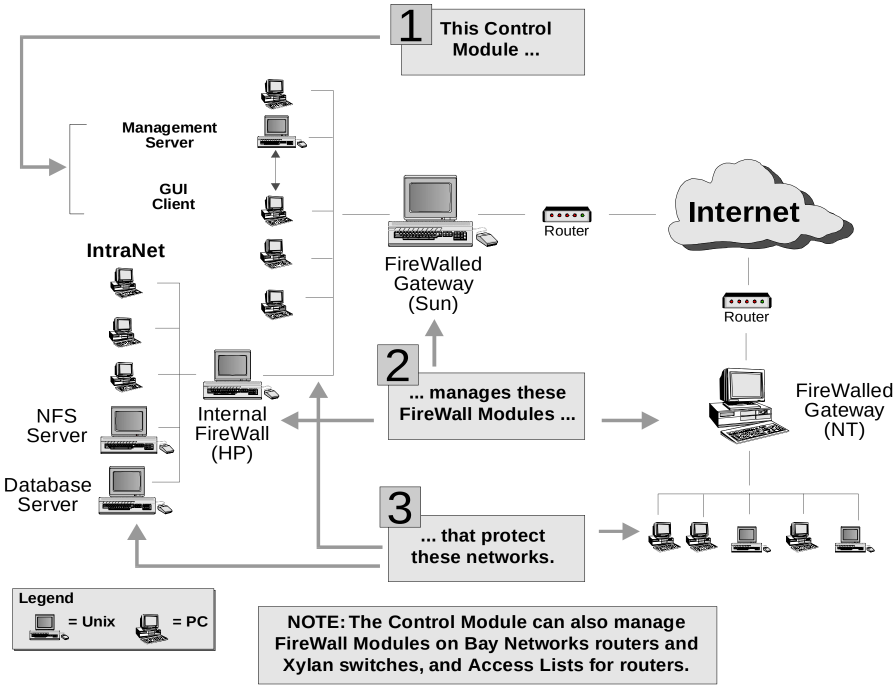
1 This Control Module ...
Management Server
GUI Client
IntraNet
NFS Server
Database Server
Internal FireWall (HP)
FireWalled Gateway (Sun)
Router
Internet
Router
FireWalled Gateway (NT)
2 ... manages these FireWall Modules ...
3 ... that protect these networks.
Legend
= Unix
= PC
NOTE: The Control Module can also manage FireWall Modules on Bay Networks routers and Xylan switches, and Access Lists for routers.

Figure 19 Distributed FireWall-1 Configuration

Figure 19 depicts a distributed configuration, on which a Control Module (in the Client/Server implementation) controls three FireWall Modules, each of which is on a different platform, which in turn protect three heterogeneous networks.

Control Module

Rule Base

Once the network administrator has defined a Security Policy — a Rule Base and the properties of the objects (networks, services, hosts, and users) used in the Rule Base — it is converted into an Inspection Script. Inspection Code, compiled from the Inspection Script, is then transmitted on a secured control channel from the FireWall-1 Management Station — the computer on which the FireWall-1 database is maintained — to the FireWall-1 daemons on the FireWalls that will enforce the policy. The FireWall-1 daemon loads the Inspection Code into the FireWall-1 Inspection Module. A network object on which the FireWall-1 Inspection Module is installed is known as a “FireWalled system.”

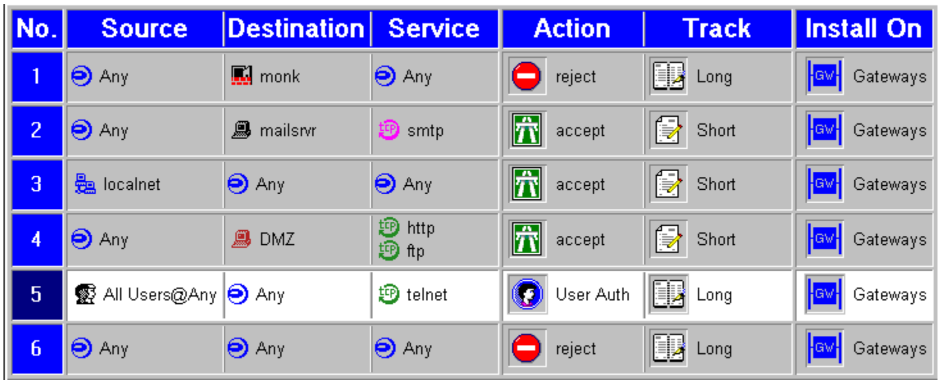

No.	Source	Destination	Service	Action	Track	Install On
1	Any	monk	Any	reject	Long	Gateways
2	Any	mailsrvr	smtp	accept	Short	Gateways
3	localnet	Any	Any	accept	Short	Gateways
4	Any	DMZ	http ftp	accept	Short	Gateways
5	All Users@Any	Any	telnet	User Auth	Long	Gateways
6	Any	Any	Any	reject	Long	Gateways

Figure 20 Example of a Rule Base

A FireWalled system enforces those parts of the Inspection Code that are relevant to itself, but all logging and alerts are sent to the network object designated as the Master. The Master also maintains the most recent Inspection Code for each of the FireWalled systems it controls. If a FireWalled system loses its Inspection Code for any reason, it can retrieve an up-to-date copy from the Master. In practice, the Master and Management Station are always on the same system. Failover Masters can be defined, which will take over if the primary Master goes down.

Communication between the Inspection Module hosts and the Management Station is secured. The Inspection Module host reports its status to the Management Station using an SNMP agent.

The FireWall-1 deployment is completely integrated. In other words, though the Security Policy may be enforced by more than one network object and, as explained below, implemented at more than one layer, there is still only one Security Policy, one Rule Base, and one centralized log.[3]

[3] In addition to a single integrated Security Policy, the system administrator can, if required, maintain different Rule Bases to be implemented, for example, at different times of day.

Router Extension Module

In the case of routers, the Access Lists derived from the Security Policy are installed by FireWall-1 on the routers. For Cisco routers, FireWall-1 downloads the access list using an Expect session that emulates a TELNET session into the router. For Bay Networks routers, FireWall-1 uses SNMP.

Network Object Manager

The Network Object Manager defines the entities that are part of the Security Policy. Only those objects that are used in the Security Policy must be defined. These may include:

- networks and sub-networks
- hosts , gateways and servers (FireWalled or not)
- routers
- Internet domains
- logical servers (among which a processing load can be distributed automatically)
- groups

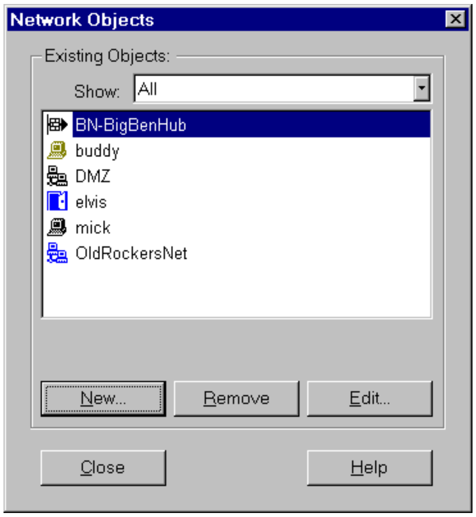

Network Objects

Existing Objects:

Show: All

BN-BigBenHub
buddy
DMZ
elvis
mick
OldRockersNet

Figure 21 Network Object Manager window

Every object has a set of attributes, such as network address, subnet-mask, etc. Some of these attributes are specified by the user, while others are extracted by FireWall-1 from the network databases, like the `hosts` and `networks` files, Network Information Services (NIS/Yellow Pages), network databases and the Internet domain service. SNMP agents are used for extracting additional information, including the interface and network configuration of hosts, routers and gateways. Objects can be combined in groups and hierarchies.

User Manager

FireWall-1 enables access privileges to be defined for users on an individual or group basis. User groups can be created, and access privileges, including allowed sources and destinations as well as user authentication schemes, can be defined (see “Authentication” on page 20).

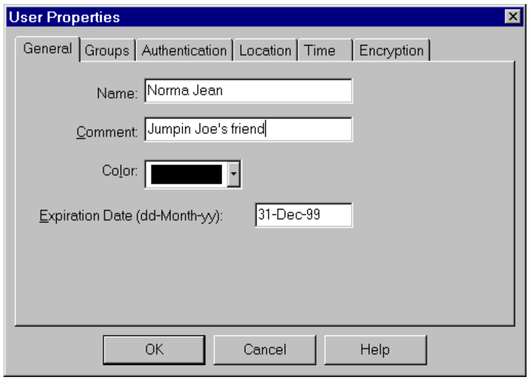

Name: Norma Jean
Comment: Jumpin Joe's friend
Color: ■
Expiration Date (dd-Month-yy): 31-Dec-99

Figure 22 User Properties window

Service Manager

The Service Manager defines the services known to the system and used in the Security Policy. All network services are screened and controlled, even those that are not defined. A comprehensive set of TCP/IP and Internet services is pre-defined, including the following:

- Standard arpa-services: TELNET, FTP, SMTP, etc.
- Berkeley r-services: rlogin, rsh, etc.
- SunRPC services: NIS/yellow pages, NFS, etc.
- Advanced Internet protocols such as HTTP, Gopher, Archie and many others
- IP services: Internet Control Message Protocol (ICMP), Routing Internet Protocol (RIP), SNMP, etc.

New services can be defined by selecting the service type and setting the service’s attributes. Service types include:

- Transmission Control Protocol (TCP)
- User Datagram Protocol (UDP)
- Remote Procedure Call (RPC)
- Internet Control Message Protocol (ICMP)
- Others — enables definition of services and protocols that do not conform to the standard set of attributes. Services are defined using simple expressions and macros.

Services can be grouped in families and hierarchies. Examples: NFS (the `mount` program, NFS-server, lock manager), NIS/Yellow Pages (`ypserv/ypbind`), and WWW (HTTP, FTP, Archie, Gopher, etc.).

System Status Monitor

The FireWall-1 Inspection Module includes robust status, auditing, and alerting capabilities. The System Status window displays a snapshot of all the FireWall Modules and Bay Networks routers at any given time. Status includes FireWall Module status as well as packet statistics (accepted, blocked, logged, etc.).

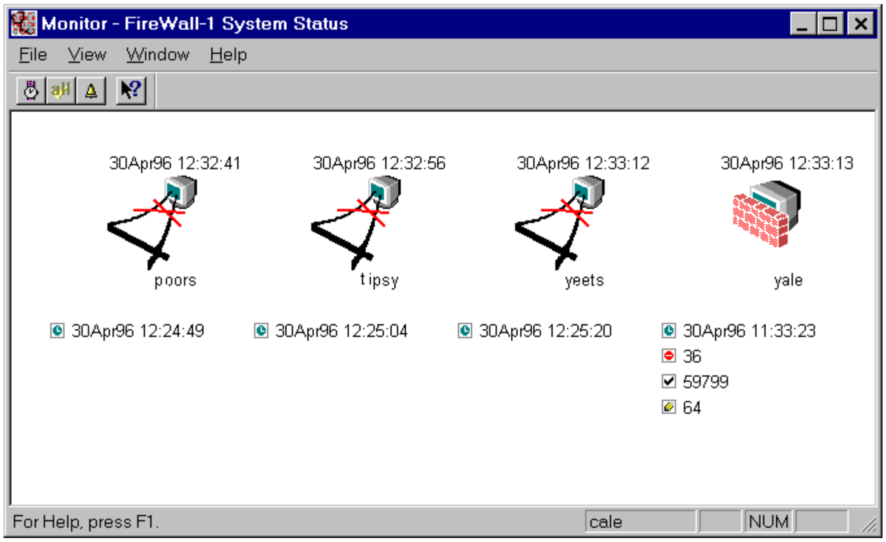

poors	tipsy	yeets	yale
30Apr96 12:32:41	30Apr96 12:32:56	30Apr96 12:33:12	30Apr96 12:33:13
30Apr96 12:24:49	30Apr96 12:25:04	30Apr96 12:25:20	30Apr96 11:33:23 36 59799 64

Figure 23 System Status Window

On Unix, FireWall-1 installs its own SNMP daemon. The FireWall-1 SNMP daemon uses either the standard SNMP MIB, or the extended FireWall-1 MIB. On Windows NT, FireWall-1 extends the built-in SNMP daemon by adding an extension agent to support the FireWall-1 MIB. The SNMP agent used by the Firewall Modules exports information to and integrates with other network management platforms. In addition, FireWall-1 can issue an SNMP trap as part of an alert.

Log Viewer

In the Rule Base, the administrator can specify logging and alerting for any communication attempt, whether the attempt is allowed or not. Log and alert formats and actions are open and can be configured by the user. The standard formats contain the source and destination of the communication, the service attempted, protocol used, time and date, source port, action (communication accepted, rejected, encryption key exchange, address translation), log and alert type, rule number, user, and the Firewall Module that originated the log entry. Any information about any communication attempt can be logged or used to trigger an alert (for example, pop up a window, send a mail message, activate a user defined action — a program or a trap).

The Log Viewer (Figure 24 on page 31) displays every logged event, including communication attempts, Security Policy installations, system shutdowns, etc. For every event, the relevant information is displayed. Fields can be displayed or hidden. Colors and icons attached to events and fields create an easily understood visual representation.

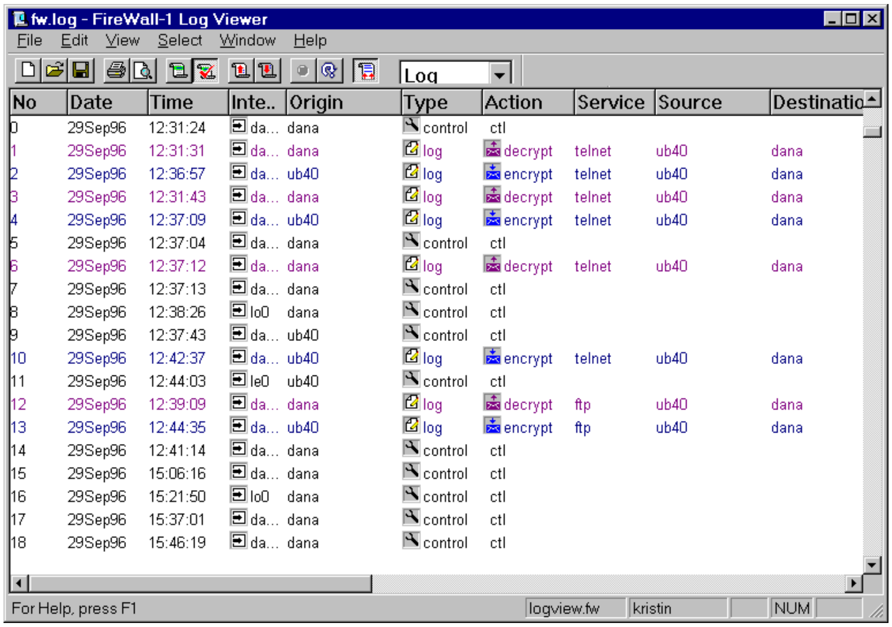
fw.log - FireWall-1 Log Viewer
File Edit View Select Window Help
Log

No	Date	Time	Inte..	Origin	Type	Action	Service	Source	Destinatio
0	29Sep96	12:31:24	da...	dana	control	ctl			
1	29Sep96	12:31:31	da...	dana	log	decrypt	telnet	ub40	dana
2	29Sep96	12:36:57	da...	ub40	log	encrypt	telnet	ub40	dana
3	29Sep96	12:31:43	da...	dana	log	decrypt	telnet	ub40	dana
4	29Sep96	12:37:09	da...	ub40	log	encrypt	telnet	ub40	dana
5	29Sep96	12:37:04	da...	dana	control	ctl			
6	29Sep96	12:37:12	da...	dana	log	decrypt	telnet	ub40	dana
7	29Sep96	12:37:13	da...	dana	control	ctl			
8	29Sep96	12:38:26	lo0	dana	control	ctl			
9	29Sep96	12:37:43	da...	ub40	control	ctl			
10	29Sep96	12:42:37	da...	ub40	log	encrypt	telnet	ub40	dana
11	29Sep96	12:44:03	le0	ub40	control	ctl			
12	29Sep96	12:39:09	da...	dana	log	decrypt	ftp	ub40	dana
13	29Sep96	12:44:35	da...	ub40	log	encrypt	ftp	ub40	dana
14	29Sep96	12:41:14	da...	dana	control	ctl			
15	29Sep96	15:06:16	da...	dana	control	ctl			
16	29Sep96	15:21:50	lo0	dana	control	ctl			
17	29Sep96	15:37:01	da...	dana	control	ctl			
18	29Sep96	15:46:19	da...	dana	control	ctl			

For Help, press F1 | logview.fw | kristin | NUM

Figure 24 Log Viewer

Logs and log records can be filtered and searched. Data can be exported to other applications and printed reports can be generated. Searching capabilities enable the user to quickly locate and track events of interest. Reports are generated by applying selection criteria to chosen fields, providing compound and comprehensive views. Reports can be viewed, exported in an ASCII format, or printed in PostScript.

On-line viewing features enable real-time monitoring of communication activities and alerts.

User Interface

FireWall-1 includes a powerful and intuitive GUI (graphic user interface), as well as a command line interface. The FireWall-1 Control Center runs under the Windows, OPEN LOOK X11R5, X/Motif and Solaris Open Windows graphical user interfaces. All FireWall-1 features and functions are also available from a command line interface, which enables operation from a standard terminal.

The command line interface can also be used in conjunction with standard system tools such as `awk` and `grep`, to invoke user-defined automatic procedures, such as log and status scanners, accounting, *etc*. Similarly, different Rule Bases can be deployed at different times of day and/or days of the week by using the command line interface together with `crontab`.

INSPECT Language

In FireWall-1, a Security Policy is described by an Inspection Script in the FireWall-1 INSPECT language. Inspection Scripts are easily understood ASCII files and can also be written using a text editor. The Graphical User Interface generates Inspection Scripts in INSPECT. In either case, the Inspection Scripts are compiled into Inspection Code to be loaded and executed by the FireWall Modules.

The ability to directly edit Inspection Scripts facilitates debugging and enables users to tailor Inspection Scripts to their specialized requirements, though in practice, this capability is rarely needed.

Firewall Module

The Firewall Module comprises two FireWall-1 daemons (snmpd and fwd), Security Servers and the Inspection Module. FireWall-1 is usually installed on a dual homed host (a gateway) but can also be installed on a server. Since FireWall-1 is loaded in the operating system kernel, it is not necessary to stop routing (disable IP Forwarding) because FireWall-1 intercepts packets before they are forwarded. In addition, processes and daemons on the gateway need not be killed, since FireWall-1 controls connections to them at the lowest layer, the network layer. FireWall-1 implements full security with connectivity.

Windows NT Support

The FireWall Module can be installed on Windows NT machines. All the functionality of the Unix implementation is available with the Windows NT implementation, including the Client/Server model.

Bay Networks Routers and Xylan Support

FireWall-1 Stateful Inspection can now be implemented on Bay Network routers and Xylan switches. The standard Firewall-1 features, with the exception of encryption, authentication and address translation, are supported.

This feature enables a greatly improved level of security to be implemented at the router and switch levels.

Inspection Module Architecture

When installed on a gateway, the FireWall-1 Inspection Module controls traffic passing between networks. The Inspection Module is dynamically loaded into the operating system kernel, between the Data Link and the Network layers (layers 2 and 3). Since the data link is the actual network interface card (NIC) and the network link is the first layer of the protocol stack (for example, IP), FireWall-1 is positioned at the lowest software layer.

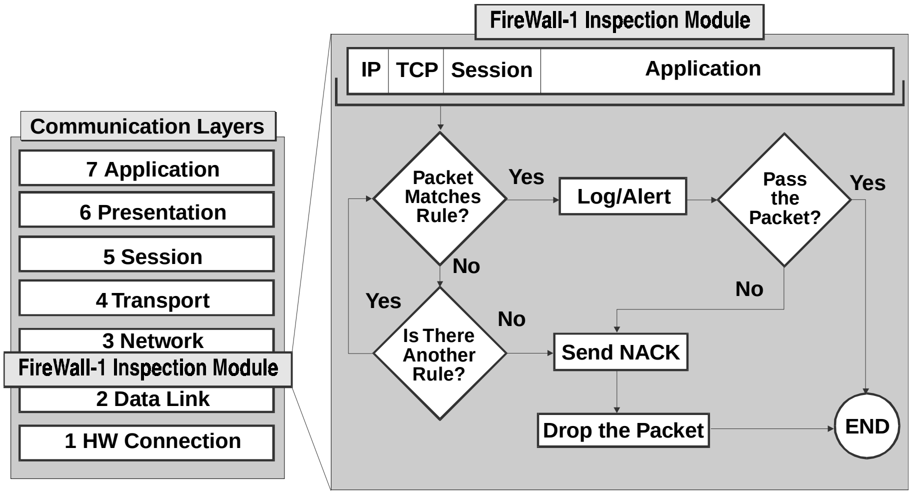

Figure 25 Inspection Module Architecture

By inspecting at this layer, FireWall-1 ensures that the Inspection Module intercepts and inspects all inbound and outbound packets on all interfaces. No packet is processed by any of the higher protocol stack layers, no matter what protocol or application the packet uses, unless the Inspection Module first verifies that the packet complies with the Security Policy.

Because the Inspection Module has access to the “raw message,” so to speak, it can inspect all the information in the message, including information relating to all the higher layers, as well as the message data itself (the application content of the packet). The Inspection Module examines IP addresses, port numbers, and any other information required in order to determine whether packets should be accepted, in accordance with the Inspection Code compiled from the Rule Base.

FireWall-1 understands the internal structures of the IP protocol family and applications built on top of them. For stateless protocols such as UDP and RPC, FireWall-1 creates and stores context data (see “UDP (User Datagram Protocol)” on page 43 and “RPC (Remote Procedure Call)” on page 44). FireWall-1 is able to extract data from the packet’s application content and store it to provide context in those cases where the application does not provide it. Moreover, FireWall-1 is able to dynamically allow and disallow connections as necessary. These dynamic capabilities are designed to provide the highest level of security for complex protocols, but the user may disable them if they are not required.

FireWall-1’s ability to look inside a packet enables it to allow certain commands within an application while disallowing others. For example, FireWall-1 can allow an ICMP ping while disallowing redirects, or allow SNMP gets while disallowing sets, and so on. FireWall-1 can store and retrieve values in tables (providing dynamic context) and perform logical or arithmetic operations on data in any part of the packet. In addition to the operations compiled from the Security Policy, the user can write his or her own expressions.

Packets that the Security Policy does not explicitly accept are simply dropped, in keeping with the principle “That Which Is Not Expressly Permitted is Prohibited.”

A rule can specify a host, an interface and a direction to which a rule applies. Packets can be inspected in any of three directions:

- eitherbound — packets are checked and verified both on entering and exiting the FireWalled system
- inbound — packets are checked and verified on entering the FireWalled system
- outbound — packets are checked and verified on exiting the FireWalled system

On gateways, the usual practice is to inspect packets on entering, though it is also possible to inspect on exiting.

Authentication

User Authentication

FireWall-1 enables the user to define a Security Policy on a per-user basis, where not only a packet’s source, destination and service are verified, but in addition, individual users of interactive sessions (TELNET, RLOGIN, HTTP and FTP) are authenticated.

When User Authentication is installed, FireWall-1 replaces the standard daemons on the FireWalled gateway with special FireWall-1 Security Servers. FireWall-1 (on the gateway) intercepts the user’s attempt to start an authenticated session on the server and “folds” the connection into the appropriate Security Server on the gateway, which performs the authentication.

When the FireWall-1 Security Server — running at the application layer — receives a connection request, it initiates an authentication procedure in accordance with the User Authentication scheme specified for the user. For example, if SecurId authentication is specified, the FireWall-1 daemon requests authentication from the local ACE client. The authenticated user must then specify the host on which to start the interactive session. FireWall-1 then verifies that the specified host is an “allowed destination” for that user.

Even after the user has been authenticated, FireWall-1 does not allow the user to open an interactive session directly on the specified host. Instead, the FireWall-1 Security Server on the gateway starts a secured interactive session on the host. The interactive session’s packets are inspected by the FireWall-1 Inspection Module as they enter the gateway, passed up to the FireWall-1 Security Server at the Application layer, and then passed down again to the FireWall-1 Inspection Module to be inspected once again before they continue on to the host. At each point, packets can be logged and alerts can be issued. In this way, the interactive session is mediated and secured by the FireWall-1 Security Server, but the user is unaware of that daemon and — except for the initial login — has the illusion of working directly on the host.

Although the Security Policy is implemented in different protocol layers, there is still only one Security Policy, one Rule Base, and one centralized logging and alerting mechanism for all layers. [4]

HTTP Security Server

The FireWall-1 HTTP Security Server provides a mechanism for authenticating users of HTTP services. The FireWall-1 HTTP Security Server can protect any number of HTTP servers behind the gateway, and control local users' access to outside HTTP servers.

Client Authentication

The User Authentication feature authenticates users of TELNET, FTP and HTTP. In contrast, the Client Authentication feature provides a mechanism for authenticating any application, standard or custom. There is no need to modify an application, either on the client or server side. The administrator can determine how each individual should be authenticated, which servers and applications will be accessible, at what times and days, and how many sessions are permitted. Moreover, after an authenticated application session begins, FireWall-1's Stateful Inspection technology provides the highest network efficiency possible. No proxies are involved, so very high hit rates and Transactions per Second are achieved through a standard workstation, with complete integration into the GUI Rule Base editor and Log Viewer.

Session Authentication

Session Authentication can be used to authenticate any service on a per-session basis.

The Session Authentication process is as follows:

1. The user initiates a connection directly to the server.
2. The FireWall Module connects to a Session Authentication Agent, which replies with the necessary authentication data.

   The Session Authentication Agent is a user written application that communicates with the FireWall Module using the FireWall-1 Session Authentication Agent Protocol. The Session Authentication Agent can run on any computer.

[4] In addition to a single integrated Security Policy, the system administrator can, if required, maintain different Rule Bases to be implemented, for example, at different times of day.

In Figure 26, the Session Authentication Agent is shown running on the client, but it can run on any computer.

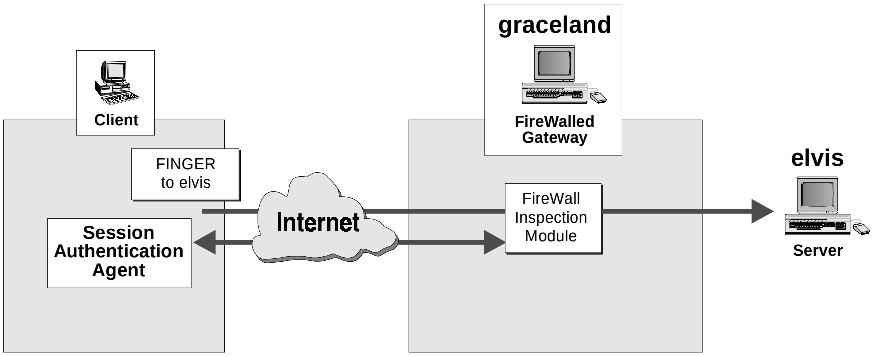

Figure 26 Session Authentication

3. If the authentication is successful, then the FireWall module allows the connection to the server to pass through the gateway.

Encryption

The security risks involved in communicating over the Internet have deterred enterprises from taking full advantage of Virtual Private Networks. Doing business over the Internet —including transferring funds, obtaining and verifying credit information, selling and even delivering products — requires a reliable and effective security solution.

Check Point FireWall-1 Solution

Check Point FireWall-1 provides secure bi-directional communication to the Internet and *transparent encryption* for full data integrity and confidentiality when establishing Virtual Private Networks.

FireWall-1 addresses all the security needs of an enterprise, ensuring:

privacy — no eavesdropping on communications

authenticity — no impersonation of network computers

integrity — no tampering with data as it passes through the network

FireWall-1 employs a single integrated Security Policy to establish enterprise-wide security, authenticate clients and encrypt communications. Offering multiple encryption schemes and integrated key management, FireWall-1 enables an enterprise to make full use of the Internet for all its business and connectivity needs.

FireWall-1 supports encryption speeds greater than 10 Mbps through a standard desktop workstation. Integrated key management is provided as well as the ability to communicate with a certificate authority. Management is simple and completely integrated into FireWall-1's GUI Rule Base editor and Log Viewer.

FireWall-1 supports three encryption schemes:

1. FWZ, a proprietary FireWall-1 encryption scheme
2. Manual IPSec, an encryption and authentication scheme that uses fixed keys
3. SKIP (Simple Key-Management for Internet Protocols), developed by Sun Microsystems, that adds improved keys and key management to IPSec

The relationship between the components of the encryption schemes, as implemented in FireWall-1, is illustrated in Table 5.

Table 5 Encryption Schemes

encryption scheme	FWZ	Manual IPSec	SKIP
authentication algorithm	MD5	MD5, SHA-1, CBC DES, MAC	MD5, SHA-1, CBC DES, MAC
encryption algorithm	DES, FWZ1	DES, RC4, FWZ1	DES, RC4, FWZ1
encryption is ...	in-place	encapsulated	encapsulated

DES, FWZ1 and RC4 are all encryption algorithms used to encrypt the data portion of a packet.

FireWall-1 provides completely transparent selective encryption for a wide range of services. Encryption, decryption, and key management are all seamlessly integrated with other FireWall-1 features.

Configuration Example

Figure 27 depicts a corporate internetwork where two private networks are connected via the Internet through FireWalled gateways. HQ is the Management Station for itself and London. The private networks (HQ-net, London-net and DMZ-net) are protected by the FireWalled gateways, but the public part of the network — the Internet — must be considered insecure.

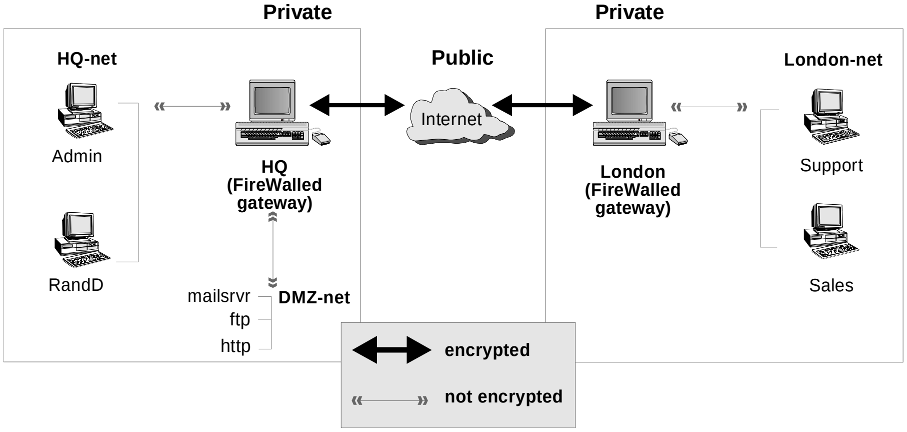

Figure 27 Two Gateway Network Configuration

Encryption Domains

A gateway performs encryption on behalf of its encryption domain: the local area network (LAN) or group of networks that the gateway protects. Behind the gateway, in the internal networks, packets are not encrypted. HQ's encryption domain consists of HQ-net and DMZ-net, and London's encryption domain is London-net.

If the system administrator has specified that communications between HQ and London are to be encrypted, FireWall-1 encrypts packets traveling between them on the Internet. A packet traveling from Sales to Admin is encrypted on the London-HQ segment of its trip, but not encrypted on the Sales-London and HQ-Admin segments.

In this way, encryption can be implemented for a heterogeneous network without the need to install and configure encryption on every host in the network.

Key Management

Key management is based on the Diffie-Hellman and RSA schemes. For each gateway it manages, a Management Station:

- generates the gateway's Diffie-Hellman key pair
- acts as the gateway's Certificate Authority
- distributes the key to the other gateways in a secure manner

At the start of an encrypted session, the gateways compute a secret key from their Diffie-Hellman keys and then begin encrypted communications.

Step-by-Step

To specify encryption, you must answer three questions:

1. Who will encrypt?

   Define the encrypting gateways and their encryption domains.

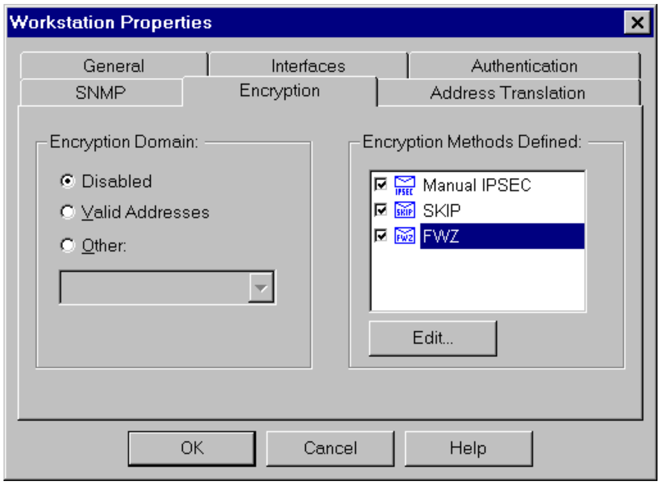

Encryption Domain:
● Disabled
○ Valid Addresses
○ Other:

Encryption Methods Defined:
☑ Manual IPSEC
☑ SKIP
☑ FWZ

Figure 28 Workstation Properties window, showing a gateway's encryption domain

Define a network object group (named “Enterprise”) that consists of all the networks that will participate in the encryption (HQ-net, DMZ-net, and London-net).

2. What are the encryption keys?

   Define the encryption keys in the Key Management window, and generate a Diffie-Hellman key pair for the gateway.

3. What will be encrypted?

   Define a rule in the Rule Base that specifies Encrypt as the action for communications between Enterprise hosts.

For additional information about FireWall-1’s Encryption Feature, refer to the publication “Privacy in Public Networks Using Check Point FireWall-1.”

FireWall-1 SecuRemote

FireWall-1 SecuRemote enables mobile and remote Microsoft Windows 95 and NT users to connect to their enterprise networks via dial-up Internet connections —either directly to the server or through Internet Service Providers — and communicate sensitive corporate data as safely and securely as from behind the corporate Internet FireWall. FireWall-1 SecuRemote extends the Virtual Private Network to the desktop and laptop.

FireWall-1 SecuRemote is based on a technology called Client Encryption. Because FireWall-1 SecuRemote encrypts data before it leaves the laptop, it offers a completely secure solution for remote access user-to-FireWall connections.

FireWall-1 SecuRemote can transparently encrypt any TCP/IP communication. There is no need to change any of the existing network applications on the user’s PC. FireWall-1 SecuRemote can interface with any existing adapter and

TCP/IP stack. A PC on which FireWall-1 SecuRemote is running can be connected to several different sites that use VPN.

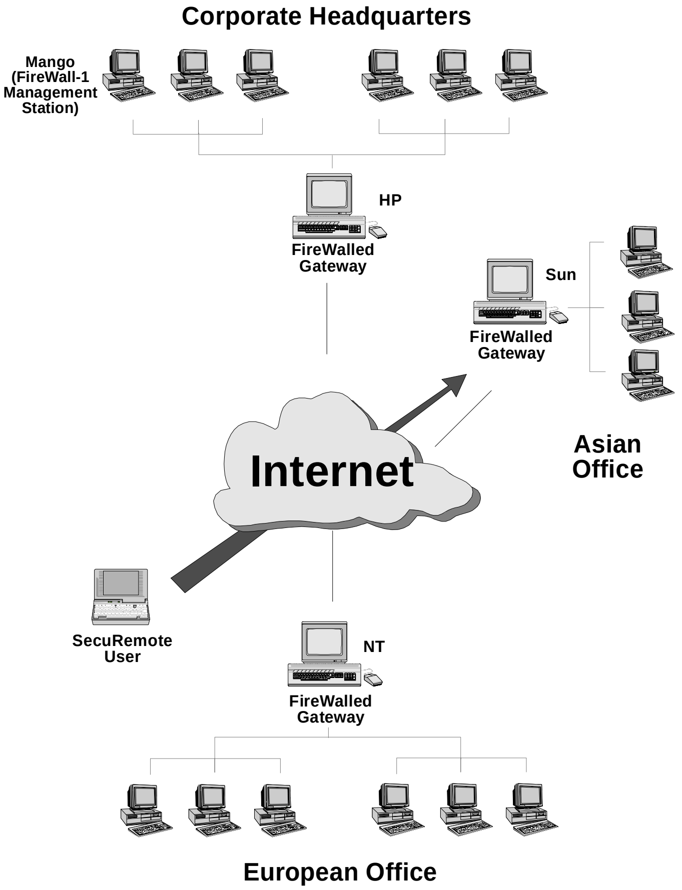

Figure 29 Virtual Private Network with Nomadic Users

A FireWall-1 security manager can enable access for FireWall-1 SecuRemote users with the standard FireWall-1 Rule Base editor. After a FireWall-1 SecuRemote user is authenticated, a completely transparent secured connection is established and the user is treated just as any user in the Virtual Private Network. The network administrator can enforce FireWall-1 security features, including authentication servers, logging and alerts, on FireWall-1 SecuRemote connections (just as with any other connection).

FireWall-1 SecuRemote includes the following features:

- support for dynamic IP addressing, which is necessary for dial-up communication
- FireWall-1 SecuRemote includes strong user authentication using Diffie-Hellman and RSA algorithms
- strong encryption using FWZ1 or DES

Address Translation

The need for IP address translation — replacing one IP address in a packet by another IP address — arises in two cases:

1. The network administrator wishes to conceal the network's internal IP addresses from the Internet.

   The administrator may reason that there is nothing to be gained, from a security point of view, by making a network's internal addresses public knowledge.

2. An internal network's IP addresses are invalid Internet addresses (that is, as far as the Internet is concerned, these addresses belong to another network).

   This situation may have arisen for historical reasons: an internal network was originally not connected to the Internet and its IP addresses were chosen without regard to Internet conventions. If such a network is then connected to the Internet, its long-established internal IP addresses cannot be used externally. Changing these addresses may be impractical or unfeasible.

In either case, the internal IP addresses cannot be used on the Internet. However, Internet access must still be provided for the internal hosts with the invalid or secret IP addresses.

Application gateways (proxies) have historically served as a partial solution to these problems. For example, to hide his or her internal IP addresses, a user can TELNET to a gateway and from there continue to the Internet through a proxy. FireWall-1 can be easily set up to provide and enforce such a scheme for a wide variety of services (FTP, TELNET, HTTP *etc*.). Moreover, FireWall-1 supplements this scheme by providing user authentication on the gateway.

On the other hand, proxies do have drawbacks:

- Proxies are tailored per application, so it is impossible to use applications that are not proxied, inbound or outbound.
- Proxies are not transparent, so that even authorized outbound users need to go through the application on the gateway, and impose a large overhead on the gateway host. Once a connection is accepted by a proxy, it functions as a packet forwarder at the application layer, which is an inefficient use of resources.
- It is difficult to provide good proxies for protocols other than TCP.

In contrast, FireWall-1's generic and transparent Address Translation Feature provides a complete and efficient solution. The administrator determines which internal IP addresses to conceal and which to map to a range of IP addresses visible to the Internet. At the same time, internal hosts can be configured to be accessible from the Internet even though their internal IP addresses are invalid Internet addresses. In short, FireWall-1 provides both IP address hiding and static Address Translation, achieving full Internet connectivity for internal clients at the maximum bandwidth possible through a standard workstation.

Address Translation can be used to implement "one way routing," so that there is no route from the outside to an internal network or to hosts.

For example, shows FireWall-1 translating the illegal internal addresses in the range 200.0.0.100 – 200.0.0.200 to the legal addresses 199.203.73.15 – 199.203.73.115.

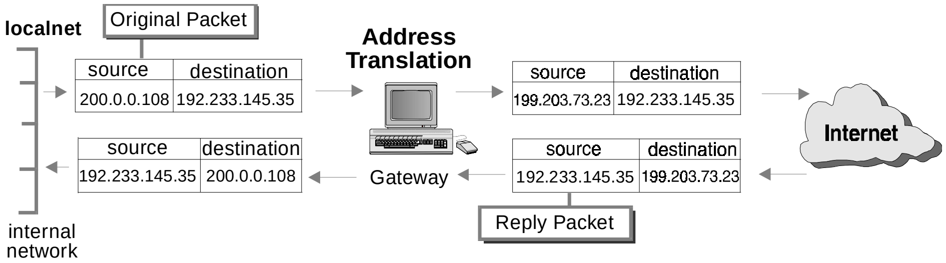

Figure 30 FireWall-1 Address Translation

A new Graphical User Interface (Figure 31) greatly simplifies specification of Address Translation rules, and allows complex rules to be more easily defined.

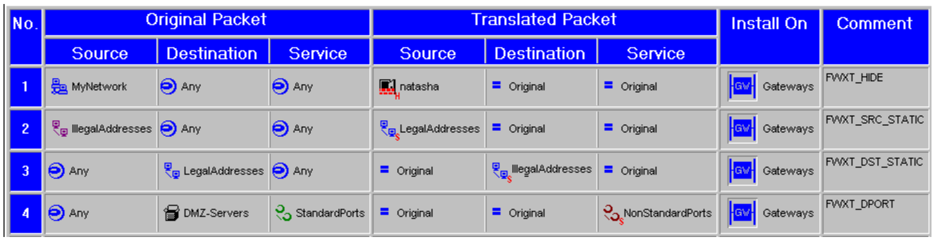

No.	Original Packet			Translated Packet			Install On	Comment
	Source	Destination	Service	Source	Destination	Service		
1	MyNetwork	Any	Any	natasha	Original	Original	Gateways	FWXT_HIDE
2	IllegalAddresses	Any	Any	LegalAddresses	Original	Original	Gateways	FWXT_SRC_STATIC
3	Any	LegalAddresses	Any	Original	IllegalAddresses	Original	Gateways	FWXT_DST_STATIC
4	Any	DMZ-Servers	StandardPorts	Original	Original	NonStandardPorts	Gateways	FWXT_DPORT

Figure 31 Address Translation Graphical User Interface

In addition, Address Translation rules can be automatically generated (Figure 32).

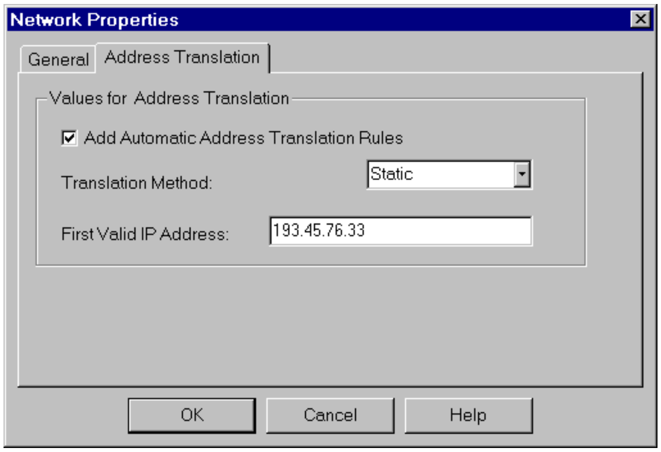

Network Properties

General | Address Translation

Values for Address Translation

☑ Add Automatic Address Translation Rules

Translation Method: Static

First Valid IP Address: 193.45.76.33

Figure 32 Automatically Generating Address Translation for a network

Outbound FTP Connections

Though the File Transfer Protocol (FTP) is one of the most basic and common TCP-based Internet protocols, securing FTP traffic is not without its difficulties.

After the client initiates an FTP session, the server establishes a new back-connection to the client. This connection goes from the server (outside the firewall boundaries) to a dynamically allocated port number on the client machine. This port number is not known in advance, and clients often open and close ports frequently. Simply opening the entire range of high-numbered ports (>1023) to incoming connections, as other firewalls do, exposes the internal network. Many successful break-ins have taken advantage of this loophole.

FireWall-1 tracks the FTP session, examining FTP application layer data. When the client requests that the server generate the back-connection (an FTP PORT command), FireWall-1 extracts the port number from the request. Both IP addresses (client and server) and both port numbers are recorded in an FTP-data pending request list. When the FTP data connection is attempted, FireWall-1 examines the list and verifies that the attempt is in response to a valid request. The list of connections is maintained dynamically, so that only the required FTP ports are opened, and only during the FTP data transfer session. As soon as the session is closed, the ports are locked again, ensuring maximum security.

Real Audio and VDOLive connections are handled in a similar manner.

UDP-based Applications

Datagram-based protocols such as UDP (User Datagram Protocol) and RPC (Remote Procedure Call) present special security problems because they are stateless, that is, without context. FireWall-1 secures these protocols by creating and dynamically maintaining context states on the basis of information in the packets and in the operating system.

UDP (User Datagram Protocol)

UDP-based applications (such as WAIS, Archie, and DNS) are difficult to filter with simplistic packet-filtering techniques because in UDP, there is no distinction between a request and a response. In the past, the choice has been to either eliminate UDP sessions entirely or to open a large portion of the UDP range to bi-directional communication, and thus to expose the internal network.

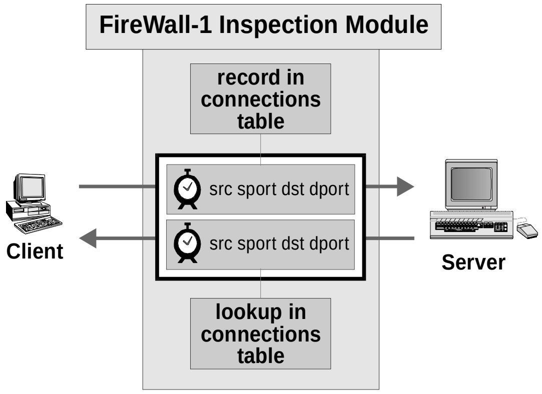
FireWall-1 Inspection Module
record in connections table
src sport dst dport
src sport dst dport
Client
Server
lookup in connections table

Figure 33 Inspecting a UDP session

FireWall-1 secures UDP-based applications by maintaining a virtual connection on top of UDP communications. FireWall-1 maintains state information for each session through the gateway. Each UDP request packet permitted to cross the FireWall is recorded, and UDP packets traveling in the opposite direction are verified against the list of pending sessions to ensure that each UDP packet is in an authorized context. A packet that is a genuine response to a request is delivered — others are dropped. If a response does not arrive within the specified time period, the connection times out. In this way, all attacks are blocked, while UDP applications can be securely used.

RPC (Remote Procedure Call)

Simple tracking of port numbers fails for RPC because RPC-based services do not use pre-defined port numbers. Port allocation is dynamic and often changes over time.

FireWall-1 dynamically and transparently tracks RPC port numbers using the port mappers in the system. FireWall-1 tracks initial `portmapper` requests and maintains a cache that maps RPC program numbers to their associated port numbers and servers.

Whenever FireWall-1 examines a rule in which an RPC-based service is involved, it consults the cache, comparing the port numbers in the packet and cache and verifying that the program number bound to the port is the one specified in the rule.

If the port number in the packet is not in the cache — this can occur when an application relies on prior knowledge of port numbers and initiates communication without first issuing a `portmapper` request — FireWall-1 issues its own request to `portmapper` and verifies the program number bound to the port.

Performance

The Inspection Module’s simple and effective design achieves optimum performance by these techniques:

- Running inside the operating-system kernel imposes negligible overhead in processing. No context switching is required, and low-latency operation is achieved.
- Advanced memory management techniques, such as caching and hash tables, are used to unify multiple object instances and to efficiently access data.
- Generic and simple inspection mechanisms are combined with a packet inspection optimizer to ensure optimal utilization of modern CPU and OS designs.

In tests, network performance degradation was too small to be measured when operating at full LAN speed (10Mb/sec) on the lowest-end SPARCstation machine; the measurement included communication latency and network bandwidth. FireWall-1 can run even at 100Mbps on standard workstations.

Performance impact on the inspecting gateway host was also negligible. These results indicate that there is almost no performance penalty for current Internet speeds (56kb or T1 - 1.5Mb/sec speeds and up to full 10Mb/sec speed) on standard desktop workstations.

Independent testing has shown consistently similar results.

FireWall-1's exceptional performance enables efficient, fully transparent outbound traffic — for example, to the Internet — while strictly enforcing powerful authentication procedures on inbound connections. FireWall-1 provides fully integrated security with no discernible performance penalty.

Conclusion

To summarize, FireWall-1's sophisticated technology includes a number of unique features:

Stateful Inspection Technology

By firewalling at the lowest software layer, FireWall-1 ensures that the Firewall Module intercepts and inspects all inbound and outbound packets on all FireWalled hosts. The Firewall Module can inspect all the information in the message, and understands the internal structures of the IP protocol family and applications built on top of them. FireWall-1 is able to extract data from the packet's application content and store it to provide context in those cases where the application does not provide it. FireWall-1 can perform any logical or arithmetic operation required to implement the Security Policy.

Transparency and Efficiency

The generic and flexible underlying Inspection Module is capable of learning and understanding any protocol, as well as adapting to newly defined protocols and applications. This capability is achieved by using high-level definitions, and requires no code changes. Moreover, FireWall-1 is transparent to applications and users.

FireWall-1 provides a sophisticated anti-spoofing feature which detects spoofed packets by ensuring that a packet's physical origin corresponds to its IP address.

Authentication

FireWall-1 implements strong User, Client and Session Authentication using different authentication schemes. All components of the Security Policy are fully integrated with the Rule Base, the object-oriented GUI and the Log Viewer.

Simplicity

FireWall-1 provides two compatible user interfaces — graphic and command-line — for user interaction and OS integration. In-band and out-of-band secure management is supported.

FireWall-1's rule based design enables users to define a Security Policy in terms of simple objects and relationships between them. Only the network objects used in the Rule Base9 need be defined by the user, and a comprehensive list of pre-defined services is available.

The FireWall Module is implemented as an independent compact code module, and resides inside the operating system kernel as a loadable kernel-layer module which requires no system reconfiguration or even re-boot to install. Because it is inside the operating system kernel, the FireWall Module is efficient and introduces no additional memory requirements or context switch overheads.

Manageability

Enterprise-wide security — security between different departments in the same organization and server security — as well as inter-network security can be managed by FireWall-1 with equal facility. Secure distributed management is seamlessly integrated into FireWall-1. Though different parts of the Security Policy are implemented at different layers and even on different computers, there is still only one integrated Security Policy, one Rule Base, and one centralized logging and alerting mechanism.

In addition, FireWall-1 provides an Address Translation feature which enables administrators to conceal internal IP addresses.

INSPECT Language

FireWall-1 provides the INSPECT language, an object-oriented, high level script language to facilitate debugging and enables administrators to tailor Inspection Scripts to their specialized requirements.

Security rules, application knowledge, context information, and packet data are combined into a powerful, integrated system that can implement any Security Policy based on any combination of parameters and logical expressions.

Network and services objects are grouped together and stored in hash-based access lists (static and dynamic), providing high-level definitions, efficient management and processing, and adaptive and intelligent inspection.

Check Point Software Technologies Ltd.

Platform Summary

component	available platforms
FireWall-1 client (Client/Server model)	Windows 95, Windows NT (Intel only), X/Motif
FireWall-1 server (Client/Server model)	Windows NT 3.51 and 4.0 (Intel only), SunOS 4.1.3 and higher, Solaris 2.3, 2.4 and 2.5, HP-UX 9 and 10
FireWall-1 Management Module (OpenLook GUI)	SunOS 4.1.3, Solaris 2.3, 2.4 and 2.5, HP-UX 9 and 10
FireWall-1 FireWalled host or gateway	Windows NT 3.51 and 4.0 (Intel only), SunOS 4.1.3 and higher, Solaris 2.3, 2.4 and 2.5, HP-UX 9 and 10, Bay Networks routers and Xylan switches (limited functionality)
Controlled Routers	Bay Networks, Cisco
SecuRemote	Windows 95 and Windows NT

Contacting Check Point Software Technologies Ltd.

International Headquarters:
3A Jabotinsky
Ramat Gan 52520, Israel
Tel: 972-3-613 1833
Fax: 972-3-575 9256

e-mail: info@checkpoint.com

U.S. Headquarters:
400 Seaport Court, Suite 105
Redwood City, CA 94063
Tel: 800-429-4391
415-562-0400
Fax: 415- 562-0410

HTTP://www.checkpoint.com