



*Check Point™ FireWall-1®/VPN-1™ Next Generation &
Radware FireProof™ 2.3 Setup Guide*

Document Information

Version – 1.01

Revision History

<i>Author</i>	<i>Title</i>	<i>Contact Email</i>	<i>Rev. Version</i>	<i>Rev. Date</i>
Aaron Glass	Sr. Engineer Bus. Dev.	aarong@radware.com	1.00	September 13, 2001
Aaron Glass	Sr. Engineer Bus. Dev.	aarong@radware.com	1.01	November 19, 2001

Table Of Contents

Introduction.....	4
Implementation Concept	5
Basic Configuration of FireProof and Check Point FireWall-1/VPN-1	6
FireProof Deployment & Route Configuration.....	6
Interface Configuration – Check Point and FireProof Devices	7
Routing Table Setup – FireWall-1/VPN-1 Enforcement Modules and FireProofs.....	8
IP Forwarding.....	10
FireProof Routing/Default Gateway Setup	12
Configuration Overview	13
FireProof Setup	13
FireWall-1/VPN-1 Network Object Configuration	16
Basic Security Policy Setup	23
Solution Traffic Flow	25
Outbound Traffic from Internal LAN.....	25
Network Address Translation.....	27
Hide NAT	27
Static NAT	28
FireProof Setup within NAT Environment.....	31
High Availability and Load Balancing Configuration	33
Basic Connectivity Checking.....	33
Full Path Health Monitoring	34
Dispatch/Load Balancing Methods.....	36
Firewall Priority	37
Warm-up/Recovery Timers	37
Firewall Grouping	38
Source Network Grouping	38
Application Grouping.....	41
Rule Combinations.....	43
Virtual Private Network Setup	45
Gateway-to-Gateway Cluster Scenario.....	45
FireProof Configuration.....	46
Policy Configuration at the Main Gateway Cluster Location	47
Policy Configuration at the Remote Standalone Gateway Location	55
SecuRemote Scenario	56

SecureClient Scenario	57
VPN Implementation Details	58
Gateway Cluster VPN Traffic Flows	58
VPN utilizing FireProof Source & Destination Grouping	60
Syslog & ELA Logging.....	62
Syslog Setup.....	62
ELA Setup.....	64

Introduction

Radware's FireProof™ is a hardware based network device that provides full fault tolerance and load balancing between multiple deployed Check Point™ FireWall-1®/VPN-1™ NG devices. The powerful traffic management and health monitoring capabilities of the FireProof ensure optimal performance and reliability of all installed firewalls within a farm ("cluster"). Using an advanced array of built-in load balancing algorithms that monitor the number of clients and real-time network load on each firewall, FireProof dynamically distributes traffic evenly between firewalls. The FireProof also takes into consideration module health, capability/capacity, as well as inbound and outbound traffic load.

This document addresses the installation of multiple Check Point FireWall-1/VPN-1 NG devices and Radware's FireProof within a single load balanced and completely fault tolerant solution. Specifically, this document will outline and discuss the following key issues:

- Setup criteria for Check Point FireWall-1/VPN-1 within a FireProof solution.
- Load balancing and FireWall-1/VPN-1 high availability configuration of the FireProof.
- FireWall-1/VPN-1 Enforcement Module NAT configuration details.
- Configuration highlights for load balanced VPN traffic (VPN-1 SecuRemote™, VPN-1 SecureClient, and Gateway-to-Gateway).

Implementation Concept

Throughout this document, references will be made to the sample diagram below (Figure-1). This example represents the fundamental network architecture for this implementation and demonstrates three actively deployed FireWall-1/VPN-1 Enforcement modules acting as security gateways between the External Network Segment (Internet), the Demilitarized Zone (DMZ), and the Internal LAN.

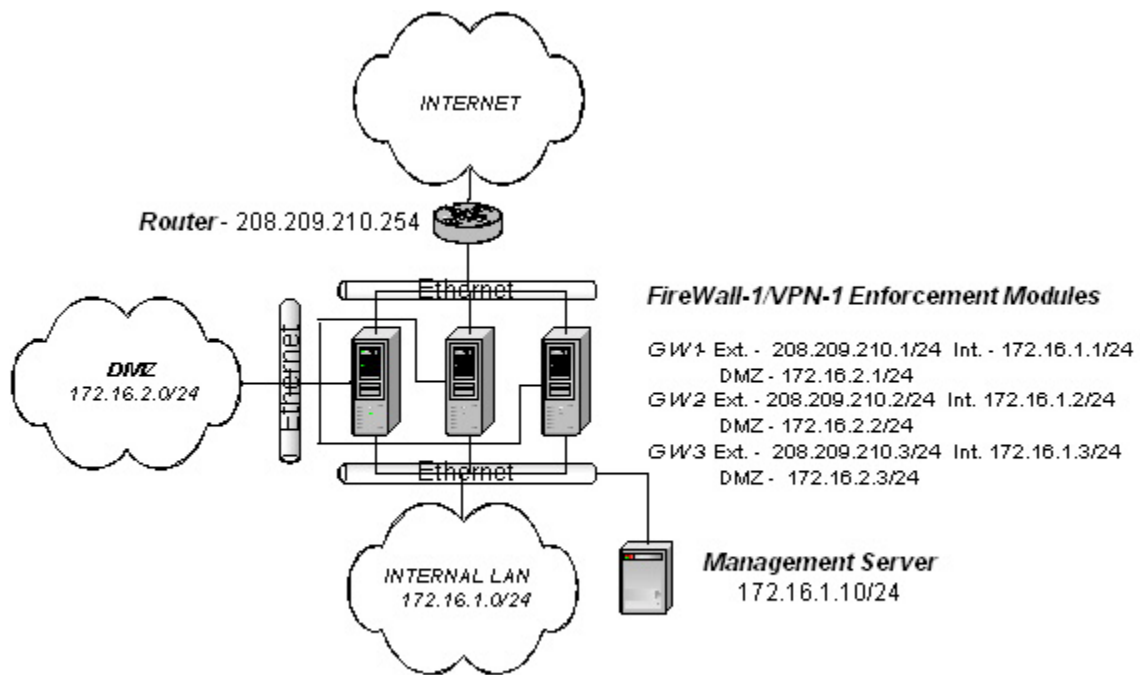


Figure 1

The above diagram shows a sample implementation utilizing three simultaneously deployed Check Point FireWall-1/VPN-1 Enforcement Modules.

Basic Configuration of FireProof and Check Point FireWall-1/VPN-1

This section will describe the basic setup of each component (i.e. network objects, gateway clusters, etc.) For detailed installation instructions, please refer to the *Check Point Next Generation Getting Started Guide*, or the *Radware FireProof User's Manual*.

Very little modification is necessary within the Check Point NG solution for integration of the Radware FireProof. Primarily, only route statements must be added to both the FireWall-1/VPN-1 Enforcement Modules and FireProof to ensure proper traffic flow through the solution.

NOTE Within this document, configuration of Radware devices is performed via ConfigWare. ConfigWare is Radware's java based SNMP management software. For information regarding CLI and other configuration methods, refer to the *Radware FireProof User's Manual*.

FireProof Deployment & Route Configuration

Typically, the FireProof is deployed similar to a router within a network. In this way, the FireProof is logically and transparently within the path of traffic allowing for a relatively non-intrusive network architecture.

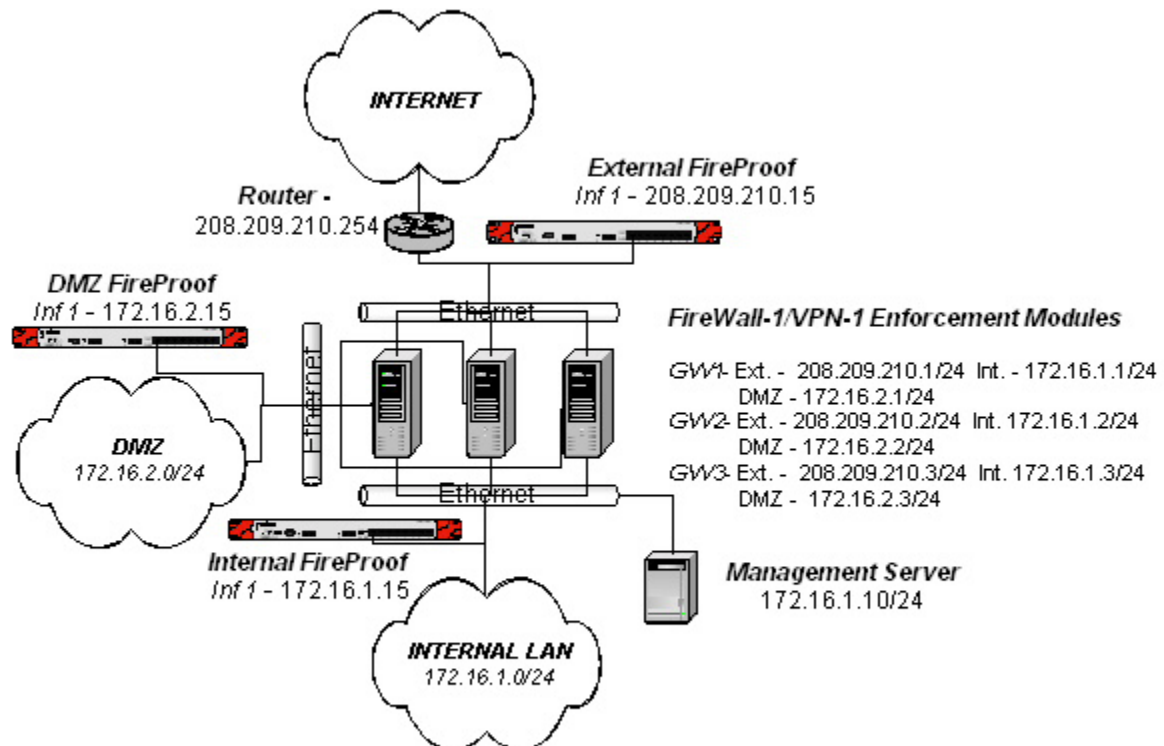


Figure 2

The above diagram outlines the implementation of FireWall-1/VPN-1 modules from Figure-1, but includes the placement of FireProof devices configured in a single-leg architecture.

Figure 2 demonstrates where the FireProof units will be deployed within the network. The FireProof can be deployed in many different ways. This example however, will utilize a “single-leg” configuration in which one physical interface is used on each FireProof. Each interface is configured with an appropriate network address for that segment and will generally serve as the default gateway for that network. Additional information regarding this setup follows.

NOTE While the FireProof is most often implemented redundantly (in pairs), this document focuses strictly on the setup and functionality of the Check Point FireWall-1/VPN-1 environment with a single FireProof device on each segment. For more information regarding the configuration of FireProof device redundancy, consult the FireProof User Manual provided with the product.

Interface Configuration – Check Point and FireProof Devices

Using Figure-2, we can first establish the basic setup and connectivity criteria for all devices within this solution.

- FireWall-1/VPN-1 Enforcement Modules
 - GW-1 – External Interface - 208.209.210.1/24
 - Internal Interface – 172.16.1.1/24
 - DMZ Interface – 172.16.2.1/24
 - GW-2 – External Interface – 208.209.210.2/24
 - Internal Interface – 172.16.1.2/24
 - DMZ Interface – 172.16.2.2/24
 - GW-3 – External Interface – 208.209.210.3/24
 - Internal Interface – 172.16.1.3/24
 - DMZ Interface – 172.16.2.3/24
- Check Point Management Server 172.16.1.10
- FireProof Devices
 - External – Interface 1 – 208.209.210.15/24
 - Internal – Interface 1 – 172.16.1.15/24
 - DMZ – Interface 1 – 172.16.2.15/24

NOTE This type of single-leg, single-network implementation is considered non-intrusive since no IP interface information must be altered (accept for routes) on each of the FireWall-1/VPN-1 modules. Instead, the FireProof is deployed on each existing network and routing is reconfigured to forward all network traffic from an interface of an FW-1/VPN-1 module through the local FireProof. Some implementers prefer to physically segment the firewall interfaces from the internal and DMZ networks. For instance, the FireProof can be implemented in a two-leg environment in which the internal LAN is a separate physical segment/subnet from the internal FW-1/VPN-1 network. This type of implementation can be slightly more complicated than the one demonstrated in this document because of the existence of additional networks but may simplify the route altering process (See next section).

Routing Table Setup – FireWall-1/VPN-1 Enforcement Modules and FireProofs

As with any security implementation, routing is extremely important to ensure proper traffic flow through security devices. The tables below represent the additional routes necessary to ensure a basic working solution.

Enforcement Modules

Routing Table for GW-1

<i>Destination</i>	<i>Mask</i>	<i>Gateway</i>	<i>Interface</i>
0.0.0.0	0.0.0.0	208.209.210.15	208.209.210.1
172.16.1.0	255.255.255.0	172.16.1.15	172.16.1.1
172.16.2.0	255.255.255.0	172.16.2.15	172.16.2.1

Routing Table for GW-2

<i>Destination</i>	<i>Mask</i>	<i>Gateway</i>	<i>Interface</i>
0.0.0.0	0.0.0.0	208.209.210.15	208.209.210.2
172.16.1.0	255.255.255.0	172.16.1.15	172.16.1.2
172.16.2.0	255.255.255.0	172.16.2.15	172.16.2.2

Routing Table for GW-3

<i>Destination</i>	<i>Mask</i>	<i>Gateway</i>	<i>Interface</i>
0.0.0.0	0.0.0.0	208.209.210.15	208.209.210.3
172.16.1.0	255.255.255.0	172.16.1.15	172.16.1.3
172.16.2.0	255.255.255.0	172.16.2.15	172.16.2.3

Enforcement Module Routing Overview

- The Default Gateway of each enforcement module is the external FireProof interface.
- The local network routes for each of the Internal and DMZ interfaces have been changed to force traffic destined to those local networks through the FireProof interface on that network first. Additional explanation is included below.

The following example demonstrates how to properly change the local route information for NT/2000 servers. The table below represents the default routing table for GW-1 before any modifications have been added.

<i>Destination</i>	<i>Mask</i>	<i>Gateway</i>	<i>Interface</i>
172.16.1.0	255.255.255.0	172.16.1.1	172.16.1.1
172.16.2.0	255.255.255.0	172.16.2.1	172.16.2.1

The existing local routes must first be removed from the table with the following command:

```
route delete 172.16.1.0 mask 255.255.255.0
route delete 172.16.2.0 mask 255.255.255.0
```

Next, these routes must be replaced so that the gateway will be the Internal and DMZ interfaces of the FireProof respectively:

```
route add -p 172.16.1.0 mask 255.255.255.0 172.16.1.15
```

```
route add -p 172.16.2.0 mask 255.255.255.0 172.16.2.15
```

The “-p” switch ensures that the routes will remain persistent within the routing table of the Enforcement Module, even after reboot.

The default route is added to each Enforcement Module as well and should be directed to the interface of the External FireProof.

```
route add -p 0.0.0.0 mask 0.0.0.0 208.209.210.15
```

The routes should now reflect those given in the above routing table configuration section.

NOTE *The local routes of the enforcement modules are altered to ensure that the FireProof is in the logical path of all traffic, both inbound and outbound. This is important for multiple reasons. First, the FireProof must be able to track inbound sessions. For example, if a session originates from the Internet destined to a server in the DMZ, the FireWall-1/VPN-1 module would forward that packet to the FireProof on the DMZ instead of forwarding it directly to the server. This allows the FireProof to create a Client Table entry for this session (Dest IP, Source IP, Originating Firewall) and guarantees that response traffic from the server (which is using the FireProof as a Default Gateway) will be forwarded through the correct FireWall-1 DMZ interface from which it came.*

Recommended Implementation

Described above are the basic route adjustments that are required for this type of single-leg, single-network configuration. There can be a problem however, when the FW-1/VPN-1 modules are rebooted. Unfortunately, routes can be added persistently (as demonstrated above), but not deleted that way. This means that the default network routes that were deleted will reappear upon reboot and the solution will not work. There are two possible solutions for this problem.

- 1) Create an automated “batch” file that starts upon boot-time. This batch file can simply input the commands displayed above so that the routes will be appropriately deleted and new ones added upon reboot. Within the NT Resource Kit, an AutoExecNT utility exists to aid in implementing this type of configuration.
- 2) Alternatively, additional persistent routes can be created that are more specific than the default interface network route. Since specific routes take higher precedence over default network routes, no route deletion will be required. Using the Internal network as an example, instead of deleting the existing 172.16.1.0/24 route, simply add two more specific routes that will cover the entire range of addresses.

```
route add -p 172.16.1.0 mask 255.255.255.128 172.16.1.15  
route add -p 172.16.1.128 mask 255.255.255.128 172.16.1.15
```

These two route statements provide the exact same routing as the default interface route. However, it is more specific, and therefore preferred over the default interface route, thus allowing you to define these routes persistently without having to create an automated batch file.

IP Forwarding

FireWall-1 NG includes major Boot Security enhancements.

In `cpconfig`, the questions whether or not to install a default filter or to disable IP forwarding no longer appear. Instead, a Default Filter is installed that accepts outgoing packets (and some Check Point protocols). IP forwarding is always disabled until a policy is loaded.

It is strongly recommended to have FireWall-1 start automatically on reboot. In UNIX `cpconfig`, answer yes to this question. For Windows NT/2000 this is the default. Starting FireWall-1 on reboot allow the Initial Policy that replaces the Default Filter to load.

The default filter protects the FireWall-1 machine from the time that the FireWall-1 machine boots up until the Security Policy is installed.

If a Security Policy has not yet been installed, an Initial Policy protects the FireWall-1 machine from the time that the FireWall services come up until a Security Policy is installed. The Initial Policy allows GUI clients to connect to the FireWall-1 machine. Until the Initial Policy loads, the Default Filter does not allow access to the FireWall-1 machine.

To change the default settings use the `fwboot bootconf` commands.
To allow IP forwarding:

```
/etc/fw.boot/fwboot bootconf set_ipf 0/1
```

To disable the default filter:

```
/etc/fw.boot/fwboot bootconf set_def 'fullpath' or (0 or no parameter) to disable the default filter.
```

After fresh installation, to remove the Default Filter you may need also to run 'fw unload localhost' . You can verify whether or not the Default Filter policy is installed by running 'fw stat'.

To check the current boot security configuration, use the `fwboot bootconf` command line commands with the "get" option.

To check the current IP forwarding setting use the commands:

```
/etc/fw.boot/fwboot bootconf get_ipf
```

To check the current Default filter setting use the command:

```
/etc/fw.boot/fwboot bootconf get_def
```

To remove or install both initial policy and default filter at once, from one command line

```
# control_bootsec
```

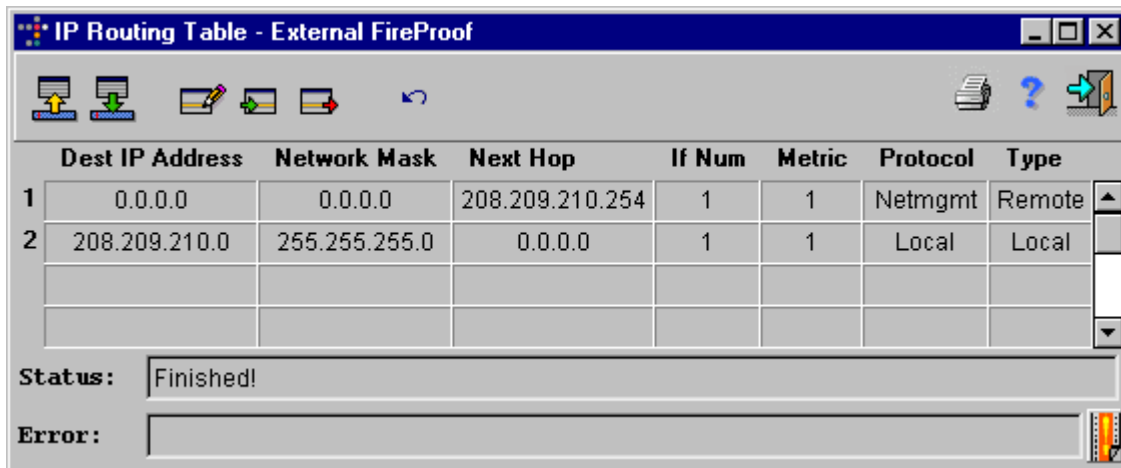
NOTE On UNIX platforms the boot settings are in the `$FW_BOOT_DIR/boot.conf` file.

On Win32 platforms it is in the registry at:

"HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\FW1\Parameters"

FireProof Routing/Default Gateway Setup

Displayed below are the routing tables for each FireProof. Make note of the Default Route configured for each.

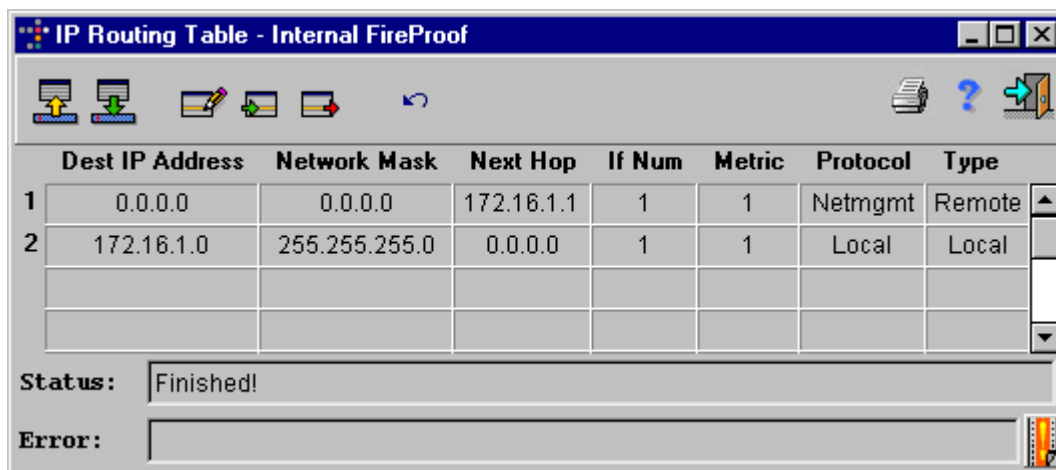


The screenshot shows the 'IP Routing Table - External FireProof' window. It contains a table with the following data:

	Dest IP Address	Network Mask	Next Hop	If Num	Metric	Protocol	Type
1	0.0.0.0	0.0.0.0	208.209.210.254	1	1	Netmgmt	Remote
2	208.209.210.0	255.255.255.0	0.0.0.0	1	1	Local	Local

Below the table, the 'Status:' field is set to 'Finished!' and the 'Error:' field is empty.

External FireProof Routing Table – Default Route configured through site Internet Router

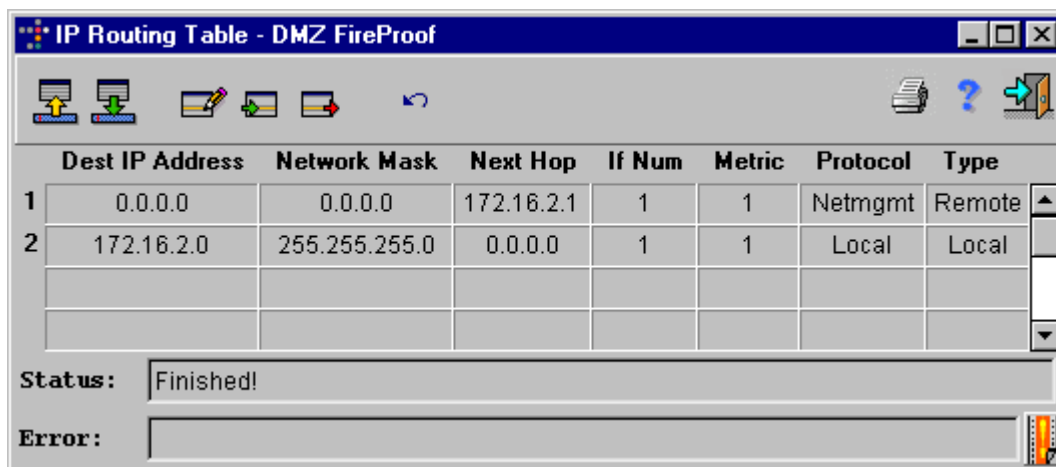


The screenshot shows the 'IP Routing Table - Internal FireProof' window. It contains a table with the following data:

	Dest IP Address	Network Mask	Next Hop	If Num	Metric	Protocol	Type
1	0.0.0.0	0.0.0.0	172.16.1.1	1	1	Netmgmt	Remote
2	172.16.1.0	255.255.255.0	0.0.0.0	1	1	Local	Local

Below the table, the 'Status:' field is set to 'Finished!' and the 'Error:' field is empty.

Internal FireProof Routing Table – Default Route configured as one of the cluster members.



The screenshot shows the 'IP Routing Table - DMZ FireProof' window. It contains a table with the following data:

	Dest IP Address	Network Mask	Next Hop	If Num	Metric	Protocol	Type
1	0.0.0.0	0.0.0.0	172.16.2.1	1	1	Netmgmt	Remote
2	172.16.2.0	255.255.255.0	0.0.0.0	1	1	Local	Local

Below the table, the 'Status:' field is set to 'Finished!' and the 'Error:' field is empty.

DMZ FireProof Routing Table - Default Route configured as one of the cluster members.

For the Internal and DMZ FireProof devices, the default gateway should be configured to one of the FireWall-1/VPN-1 modules within the cluster. Even though only one is selected, the route serves as a directional guide for load balancing decisions. Essentially, if traffic is routed from the Internal Network (whose clients will be configured to route through the interface of the Internal FireProof), the FireProof, seeing that the next hop route exists through one of the configured FireWall-1/VPN-1 modules within the cluster, will make a load balancing decision based on administrator defined criteria and module health for all of the configured devices in the cluster.

Configuration Overview


At this point, the following aspects of configuration should be completed:

- FireWall-1/VPN-1 Enforcement and Management Modules should be properly installed and licensed.
- The interfaces for all devices should be appropriately addressed.
- The routing should be properly setup on all devices.

The next steps of installation include the setup of all of the FireProof devices as well as the network object creation and policy installation.

FireProof Setup

Connect to each FireProof with ConfigWare. Once connected, proceed to the Firewall Table (from the Menu bar; *FireProof* → *Firewall Table*).



Once the Firewall Table is open, insert the modules from the cluster for each appropriate FireProof. To do this, double-click on any empty cell within the table, or click on the green "Insert" arrow located on the tool bar.

The screenshot shows the 'Firewall Table Insert - Internal FireProof' dialog box. It contains the following fields and settings:

- Firewall Address: 172.16.1.1
- Firewall Name: GW-1
- Admin Status: Enabled
- Firewall Priority: 1
- Kbits limit: 0
- Inbound Kbits limit: 0
- Outbound Kbits limit: 0
- Firewall Mode: Regular
- Connection Limit: 0

Firewall Table Insert Window

As the figure above shows, there are many settings available for each firewall module configured within the FireProof. While all that is required is the *Firewall Address* and *Firewall Name*, the other settings can be configured to tweak or control the flow of traffic through each individual FireWall-1/VPN-1 module (i.e. priority, inbound and outbound bandwidth, etc.). These additional features are beyond the scope of this document, but more information can be found within the *Radware FireProof User's Manual*.

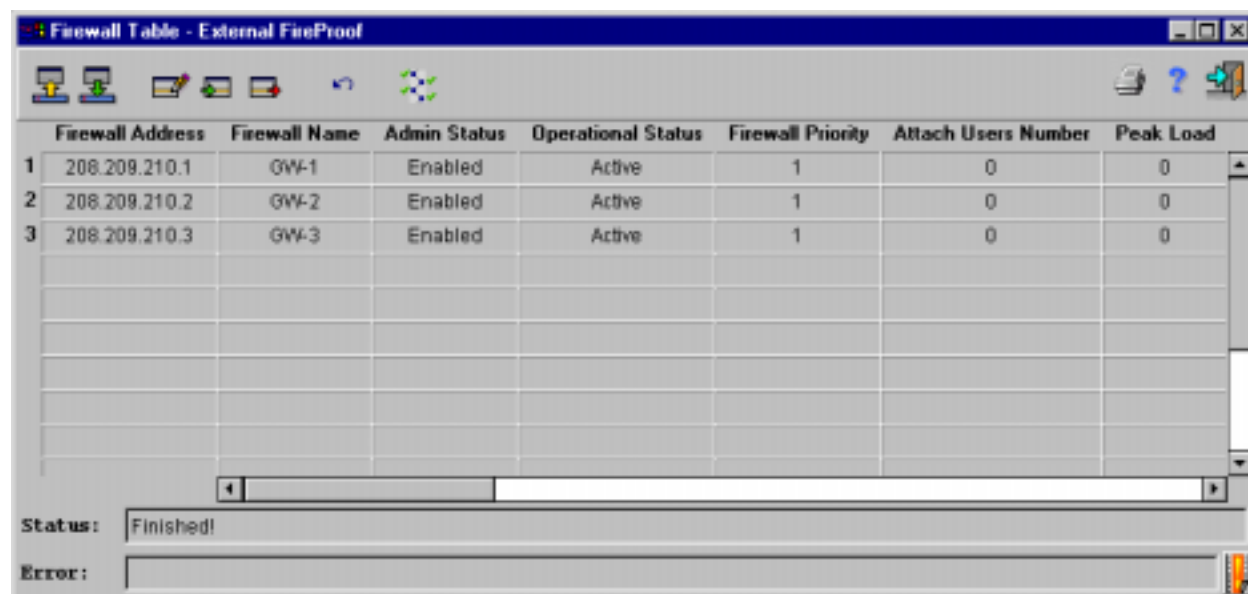
Remember to install the correct firewall modules on the appropriate FireProof. For example, if configuring the External FireProof, insert only the external interface of each Enforcement Module within the cluster. The following pictures demonstrate what the Firewall Table of each FireProof should look like given the discussed configuration.

The screenshot shows the 'Firewall Table - Internal FireProof' window. The table below represents the data shown in the window:

	Firewall Address	Firewall Name	Admin Status	Operational Status	Firewall Priority	Attach Users Number	Peak Load
1	172.16.1.1	GW-1	Enabled	Active	1	0	0
2	172.16.1.2	GW-2	Enabled	Active	1	0	0
3	172.16.1.3	GW-3	Enabled	Active	1	0	0

At the bottom of the window, the Status is 'Finished' and the Error field is empty.

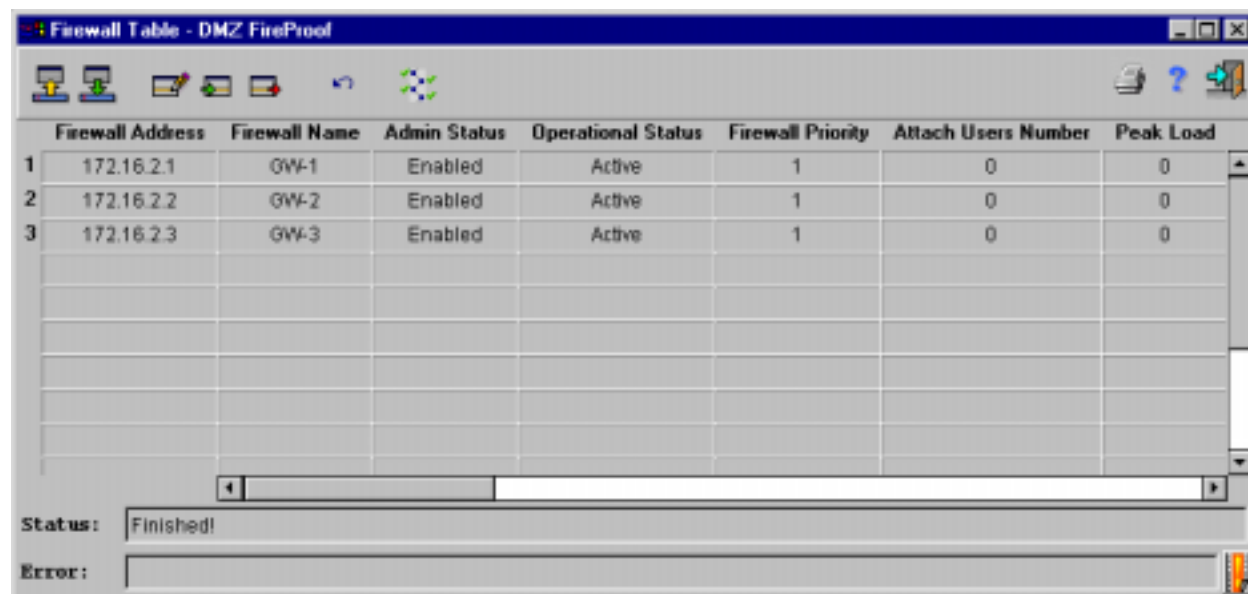
Firewall Table of the Internal FireProof



	Firewall Address	Firewall Name	Admin Status	Operational Status	Firewall Priority	Attach Users Number	Peak Load
1	208.209.210.1	GW-1	Enabled	Active	1	0	0
2	208.209.210.2	GW-2	Enabled	Active	1	0	0
3	208.209.210.3	GW-3	Enabled	Active	1	0	0

Status: Finished!

Error:

Firewall Table of External FireProof

	Firewall Address	Firewall Name	Admin Status	Operational Status	Firewall Priority	Attach Users Number	Peak Load
1	172.16.2.1	GW-1	Enabled	Active	1	0	0
2	172.16.2.2	GW-2	Enabled	Active	1	0	0
3	172.16.2.3	GW-3	Enabled	Active	1	0	0

Status: Finished!

Error:

Firewall Table of the DMZ FireProof


Once the interface, routing, and Firewall Table information is successfully updated on each FireProof, you can now move on to configuring Network Objects and a Security Policy for the cluster of FireWall-1/VPN-1 modules.

FireWall-1/VPN-1 Network Object Configuration

Connect to the management server with the Policy Editor GUI. Once connected a few basic objects must be configured.

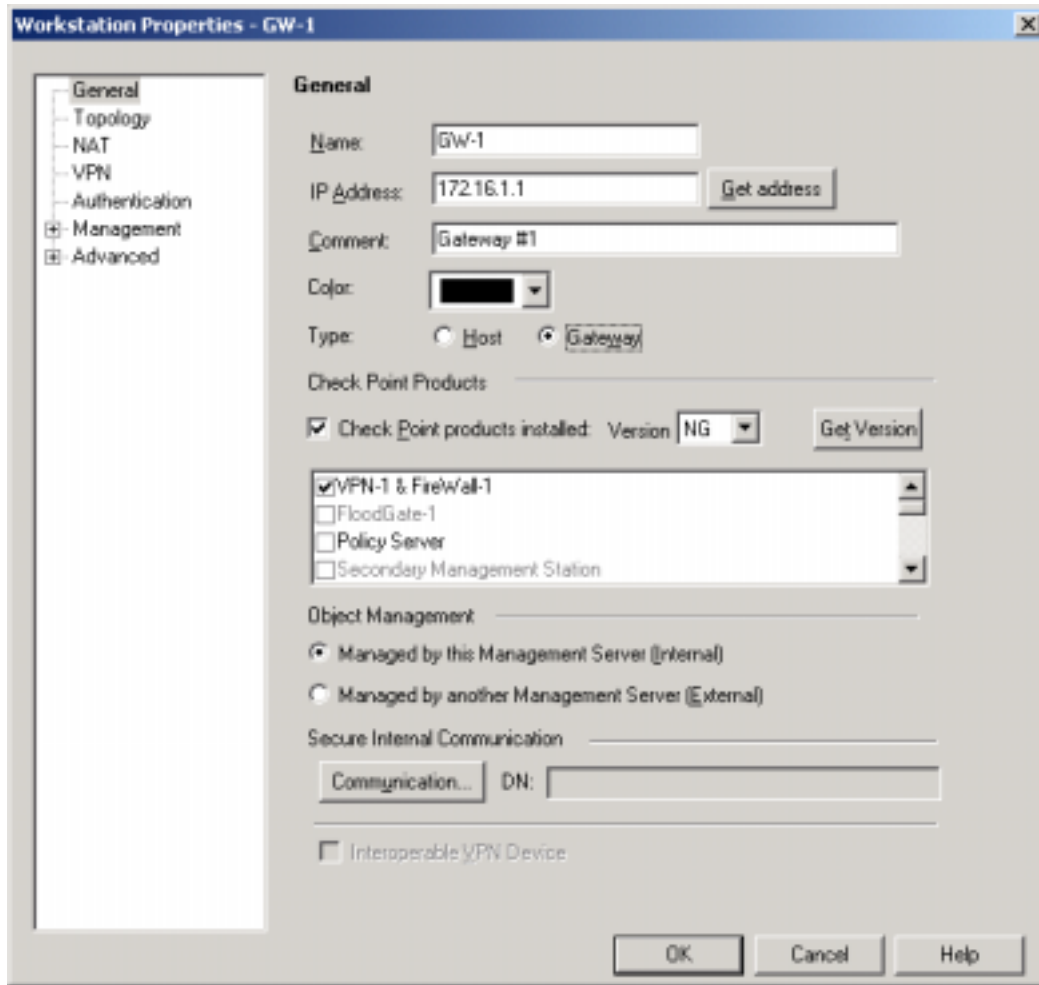
- Workstation (Gateway) Objects for each Enforcement Module
- Gateway Cluster Object which will include Workstation (Gateway) Objects
- Network Objects for Internal and DMZ networks

To begin configuration of Network Objects, open the Network Objects window from the Menu bar (*Manage* → *Network Objects*).



Network Objects Window

Click on the "New..." button and select "Workstation".

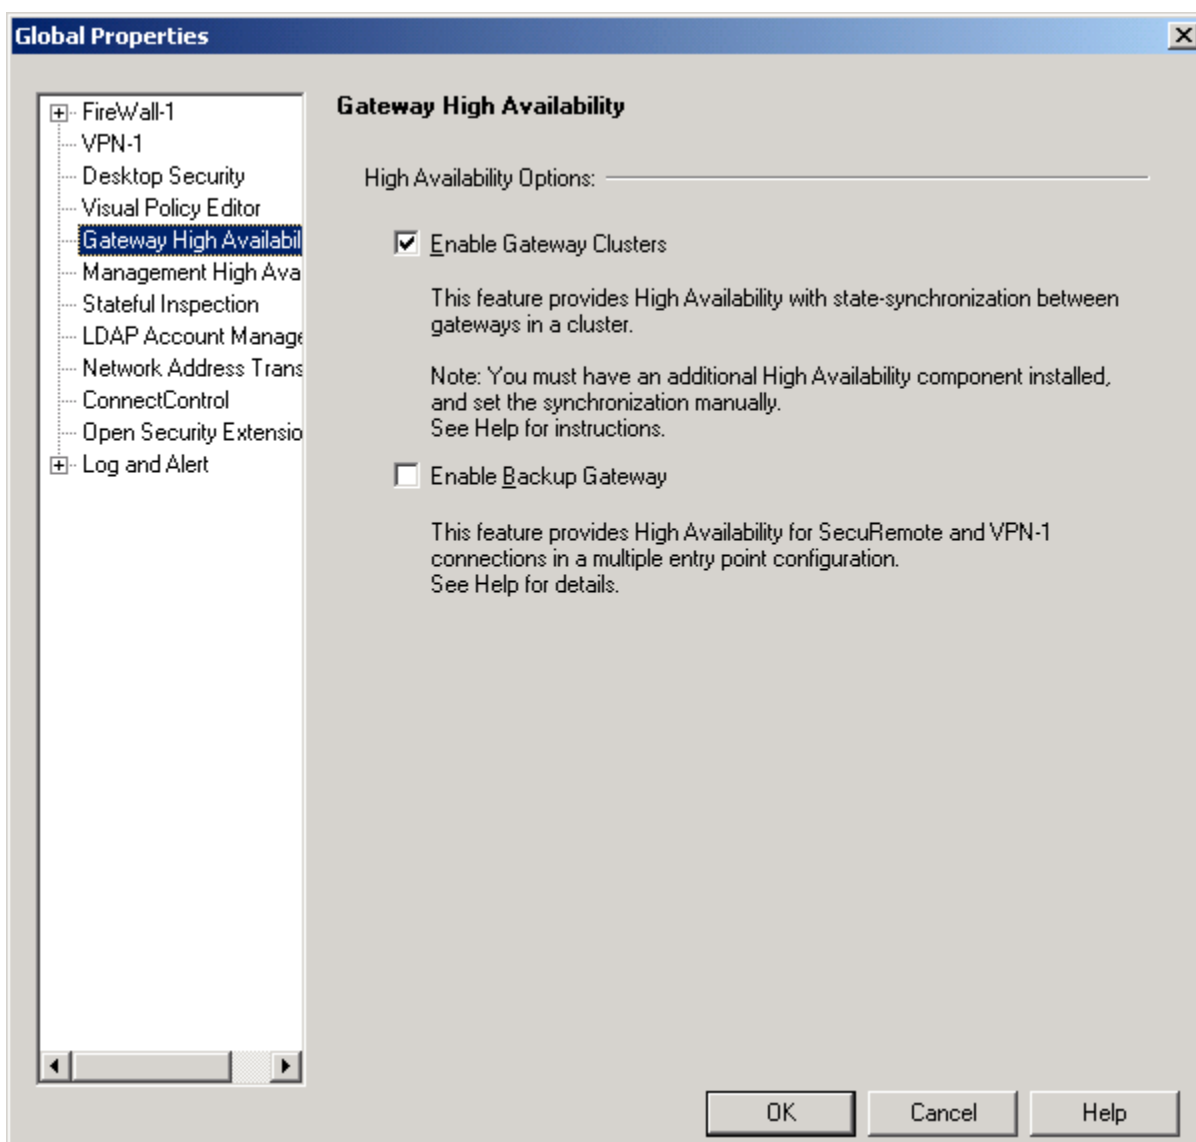


Workstation Object Insert Window

Once open, insert all relevant information within the Workstation Object Insert Window as shown in the figure above. Select the Type as “Gateway” and insure that the “VPN-1 & FireWall-1” box is selected as installed. Configure Secure Internal Communication and any necessary topology information for this solution. For additional installation instructions, refer to *Check Point Next Generation Getting Started Guide*.

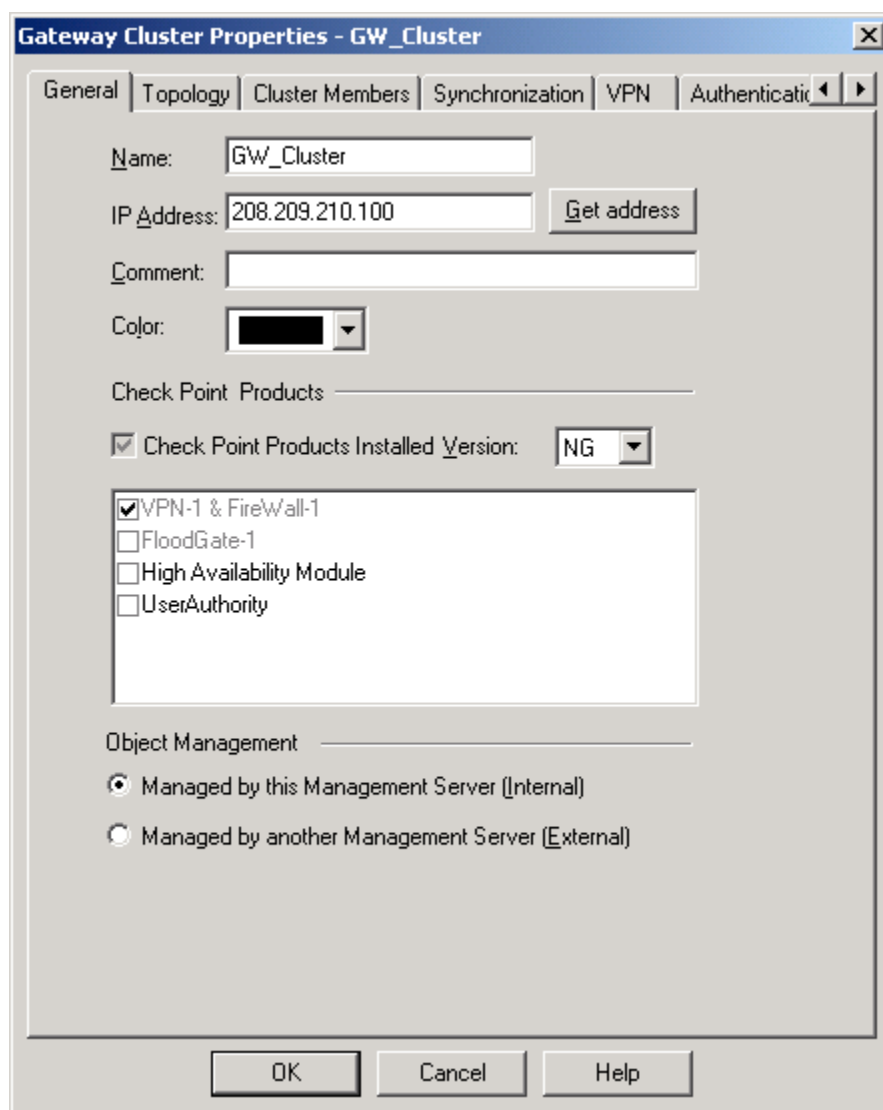
Configure a Workstation (Gateway) object for each Enforcement Module within the solution and proceed to configuring a Gateway Cluster Object.

Before a Gateway Cluster can be created, the Gateway High Availability option must first be enabled. This is found within the Global Properties window.



Global Properties – Gateway High Availability Window

Once the “Enable Gateway Cluster” box has been checked within the Global Properties window, proceed again to the Network Objects window. From this window, once again click on “New...” and now select Gateway Cluster. The following window will be displayed.




The screenshot shows a dialog box titled "Gateway Cluster Properties - GW_Cluster". It has several tabs: "General", "Topology", "Cluster Members", "Synchronization", "VPN", and "Authentication". The "General" tab is selected. The "Name" field contains "GW_Cluster". The "IP Address" field contains "208.209.210.100" and has a "Get address" button next to it. There is an empty "Comment" field and a "Color" dropdown menu. Under "Check Point Products", there is a checkbox for "Check Point Products Installed Version" which is checked, and a dropdown menu showing "NG". Below this is a list of products with checkboxes: "VPN-1 & FireWall-1" (checked), "FloodGate-1", "High Availability Module", and "UserAuthority". Under "Object Management", there are two radio buttons: "Managed by this Management Server (Internal)" (selected) and "Managed by another Management Server (External)". At the bottom are "OK", "Cancel", and "Help" buttons.

Gateway Cluster Insert Window

From the Gateway Cluster Insert Window, specify a name and Cluster IP Address. This address should be a valid and available address from the public network. For this example, the address 208.209.210.100 has been selected and inserted into the window. This address will be important for the VPN implementation covered later in this document.

When the Name and IP Address has been configured, proceed to the “Cluster Members” tab of this window.

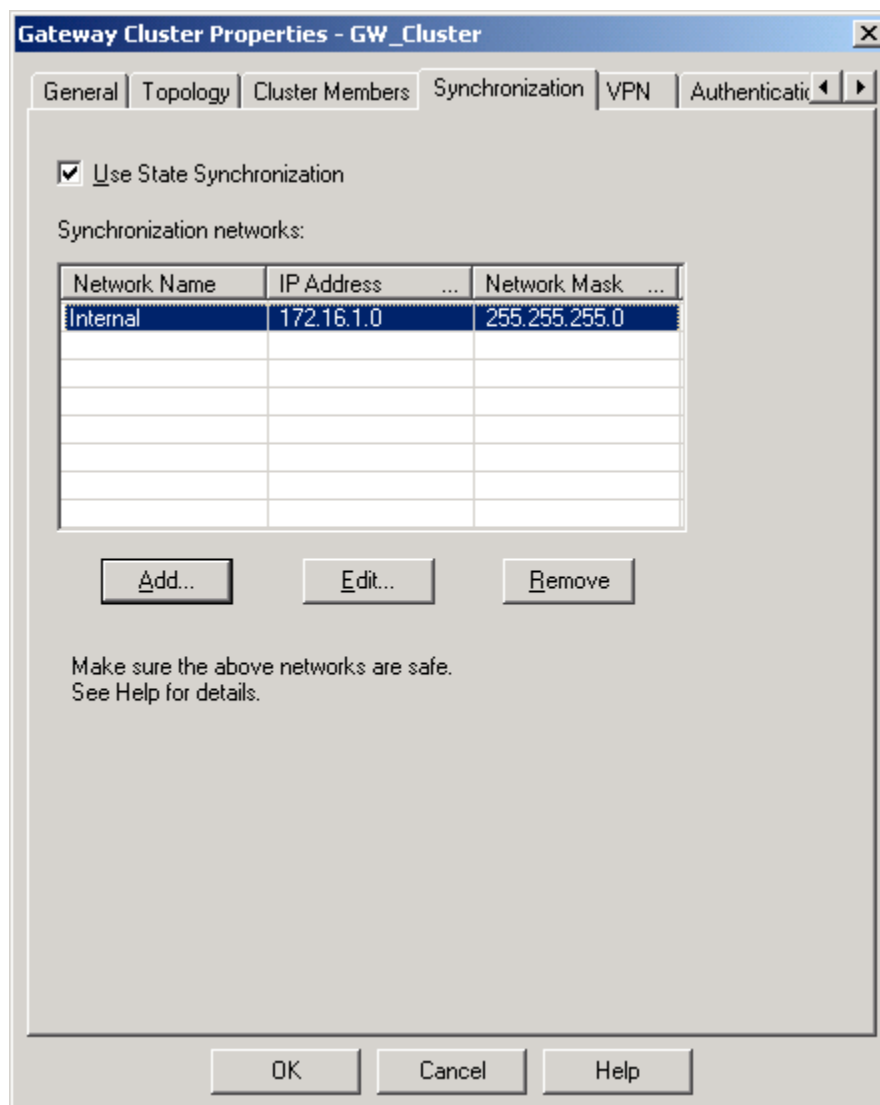


Cluster Members Window

From this window, click "Add..." and insert each of the configured Gateway Objects for this site into the cluster.

NOTE *Mixing FireWall-1/VPN-1 4.1 Enforcement Modules in an NG Cluster is not supported.*

Once complete, proceed to the "Synchronization" tab of this window.

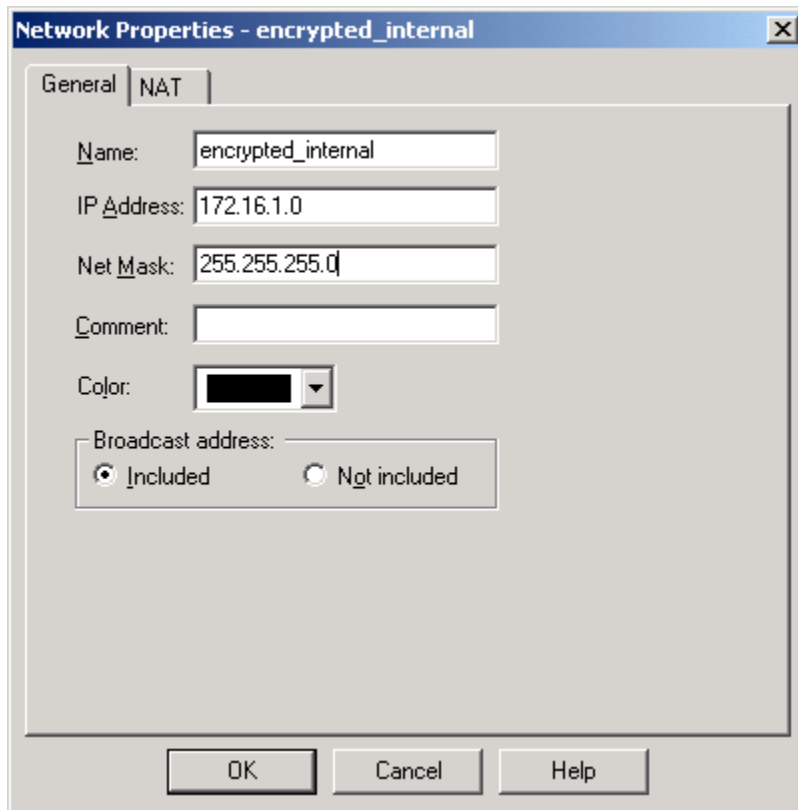


Synchronization Setup Window

This window allows the administrator to enable and define what network all Gateway Objects within the Gateway Cluster will use for synchronization traffic. It is possible to configure a dedicated network for synchronization and management, for instance, if a fourth interface existed.

NOTE For the Synchronization feature to work properly, the option must be enabled on each Enforcement Module as well. This setting is configured within the Check Point Configuration Tool on each module. Additionally, to insure that the feature is setup and functioning properly, run the "fw tab -t connections" command from each Enforcement Module. This should verify that connections have been synchronized.

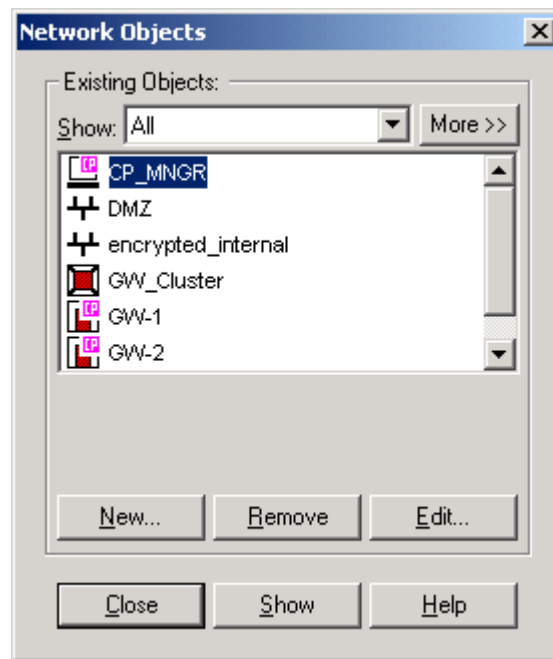
Network objects should be created for each network within a given implementation. These are necessary when creating general security and VPN policies. Within this example, two Network Objects will be created, "encrypted_internal" and "DMZ".



Network Object Insert Window

This figure demonstrates the basic setup for the Internal Network object. It has been named "encrypted_internal" because traffic destined to and from this network will participate in a Virtual Private Network. VPN setup and configuration will be covered in detail later in this document.

When all Workstation (Gateway), Gateway Cluster, and Network objects have been created, the Network Objects window should look as follows:



Network Objects Window

Basic Security Policy Setup

In order for the FireWall-1/VPN-1 modules to forward traffic between interfaces, a Security Policy must be created.

From the Menu bar, select *Rules* → *Add Rules* → *Top*

NO.	SOURCE	DESTINATION	SERVICE	ACTION	TRACK	INSTALL ON
1	★ Any	★ Any	★ Any	● drop	- None	Gateways

Standard Security Policy Table

A default “drop” rule will be automatically configured. Continue to add and modify rules that apply to your security implementation, one rule at a time.

An example of a very basic rule base is described below.

NO.	SOURCE	DESTINATION	SERVICE	ACTION	TRACK	INSTALL ON
1	encrypted_internal	Any	http ftp icmp-protol	accept	- None	Gateways
2	Any	DMZ	http	accept	- None	Gateways
3	Any	Any	Any	drop	Log	Gateways

Standard Security Policy Table

This table contains three basic rules.

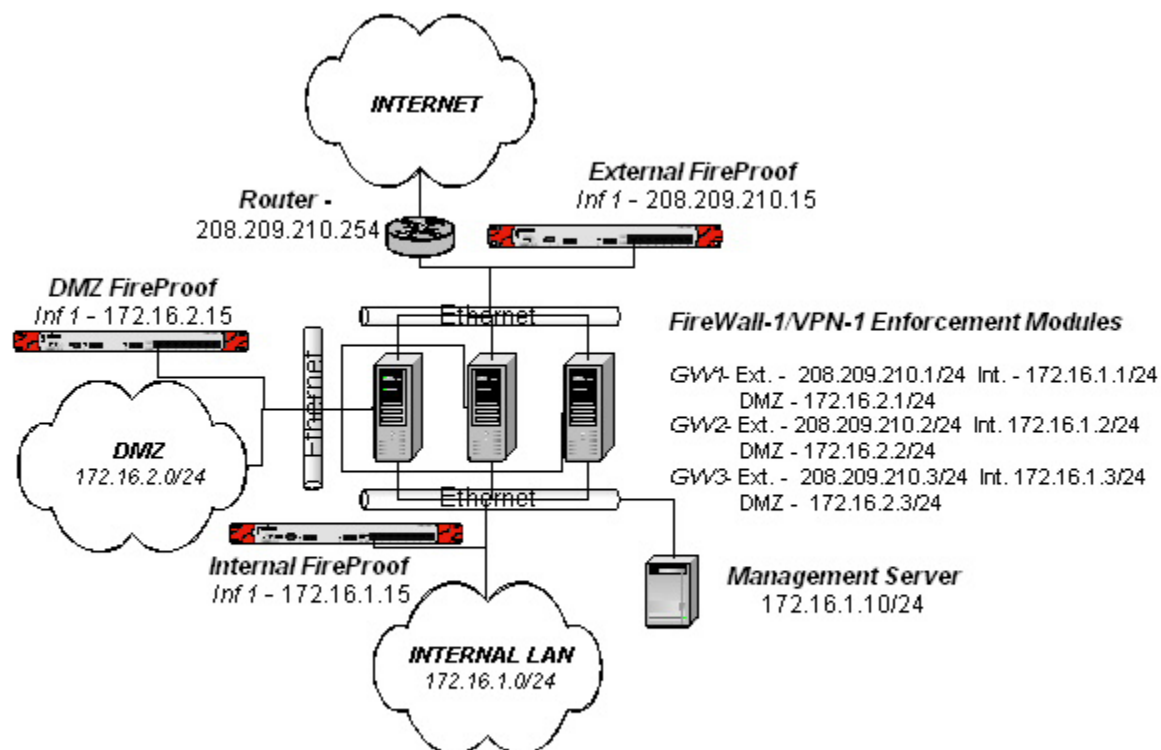
1. Allows originating HTTP, FTP, and ICMP traffic from the Internal Network (172.16.1.0) to any destination.
2. Allows HTTP access from any source into the DMZ network (172.16.2.0).
3. This rule is the “none of the above” or “cleanup” rule; it rejects and logs all other communications.

To install the Security Policy on the Gateway Cluster, choose *Policy* → *Install* from the Menu bar.

NOTE When installing the Security Policy, it is recommended that the administrator select “Install on all the members of the selected Gateway Clusters, if it fails don’t install at all” option from the Install Policy window.

Solution Traffic Flow

Once installed, the Security Policy will enforce the specified rules for all traffic traversing the firewalls. To provide a better understanding of how traffic will traverse this solution, this section will outline a sample traffic flow utilizing the Security Policy defined in the preceding section as well as the network diagram below.



Sample Network Diagram (Same as Figure 2)

Outbound Traffic from Internal LAN

Situation

- Internal LAN client 172.16.1.80 initiates an HTTP connection to 178.34.67.100

Outbound Traffic Flow

- The initial packet is forwarded to the interface MAC address of the Internal FireProof (172.16.1.15) because it is the Internal LAN's default gateway.
- The packet is analyzed by the FireProof to determine if this is a new session by looking through its existing Client Table entries. If no record exists of this session, a load balancing decision is made. *More detail regarding load-balancing algorithms can be found in the High Availability and Load Balancing Configuration section of this document.*

- The FireProof chooses among the available FireWall-1/VPN-1 modules within its Firewall Table, forwards the packet through the selected module (GW-1) and makes a TO Client Table entry as follows:

<i>Client Addr</i>	<i>Dst Addr</i>	<i>Firewall Addr</i>	<i>Src Port</i>	<i>Dst Port</i>	<i>Direction</i>
172.16.1.80	178.34.67.100	172.16.1.1	1234	80	TO

Internal FireProof Client Table

This entry is used for subsequent packets from this session so that the FireProof will be aware of the session and not inadvertently "spray" the packets from this session among multiple FireWall-1/VPN-1 modules.

- Next, the selected FireWall-1/VPN-1 module receives the packet, which is analyzed and applied to the existing Security Policy rule base. Since this packet matches rule #1 (HTTP from Internal Network), the FireWall-1/VPN-1 module will forward this packet through it's next hop router (the External FireProof – 208.209.210.15).
- The FireProof receives this packet, recognizes that it is being forwarded by one of the FireWall-1/VPN-1 modules from its Firewall Table and makes a FROM entry in its Client Table and forwards the packet to its default gateway (208.209.210.254).

<i>Client Addr</i>	<i>Dst Addr</i>	<i>Firewall Addr</i>	<i>Src Port</i>	<i>Dst Port</i>	<i>Direction</i>
178.34.67.100	172.16.1.80	208.209.210.1	80	1234	FROM

External FireProof Client Table

In this FROM entry, you'll notice that the Client and Dst addresses have been reversed. This feature, known as Session Tracking, will insure that returning traffic for this session is forwarded through the appropriate firewall, thus alleviating the possibility of asymmetric routing through the farm.

Return/Response Traffic Flow

- Return traffic will be routed by the site's ISP router to the External FireProof on its way back to the original client.
- The packet will first be compared to the External FireProof Client Table for an existing entry. Since one does exist, the FireProof knows to forward this packet through GW-1.
- Once received by GW-1, the packet is then analyzed and shown to belong to an existing session that has originated through this FireWall-1/VPN-1 module, and is then forwarded back to the Internal Network. Remember, that the network route has been modified so that all traffic destined to this local network will first be forwarded through the Internal FireProof.
- The packet is forwarded to the Internal FireProof, and then on to 172.16.1.80.

This example should properly demonstrate the importance of correct routing throughout this entire environment to insure proper load balancing and session tracking, thus avoiding asymmetric routing problems.

Network Address Translation

Most implementations like the one discussed in this document will utilize “invalid”, or privately addressed networks, behind the FireWall-1/VPN-1 cluster. Because of this, it is important to understand how to enable and properly configure NAT policies within a FireProof implementation.

This section will cover:

- The two types of NAT configuration – Hide & Static
- FireWall-1/VPN-1 Setup Instructions
- FireProof and upstream router modifications to allow NAT on Enforcement Modules

There are two methods of translating IP addresses. One method, *hiding*, is used to hide all the invalid addresses behind the gateway’s valid address. Using this method, it is impossible to initiate connections to hosts in the invalid/privately-addressed networks “behind” the FireWall-1/VPN-1 modules.

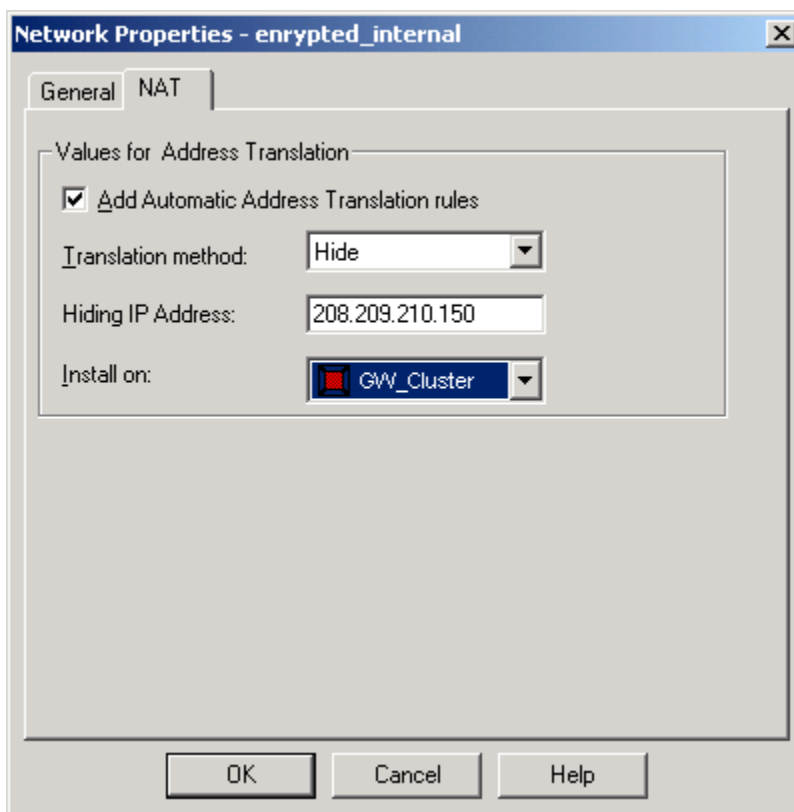
The second method, *static translation*, allows the administrator to assign or “map” valid addresses to hosts on the privately addressed networks on a one-to-one basis. This method will allow outside hosts to initiate connections to hosts on the Internal or DMZ networks (i.e. inbound mail services, web, ftp, etc.).

Hide NAT

The Hide NAT feature is used to translate internally private addressed hosts/networks to a single valid address. This is often referred to as a many-to-one NAT, as there are many clients utilizing one or few valid NAT addresses.

For this implementation we will NAT the outbound traffic from the Internal Network to the Internet with a single valid address. Because this is a Gateway Cluster an address not utilized by any single entity will be used, 208.209.210.150.

This address will be used for all outbound traffic. The following picture displays how to configure the Hide NAT address for the internal_encrypted network object created earlier in this document.



Configuring Hide NAT on a Network Object

Select "Add Automatic Address Translation rules".

Select "Hide" as the Translation method and specify the Hiding Address. This will be the IP address that traffic originating from the network object will translate to.

Finally, select the GW_Cluster object to install this NAT policy on and click OK.

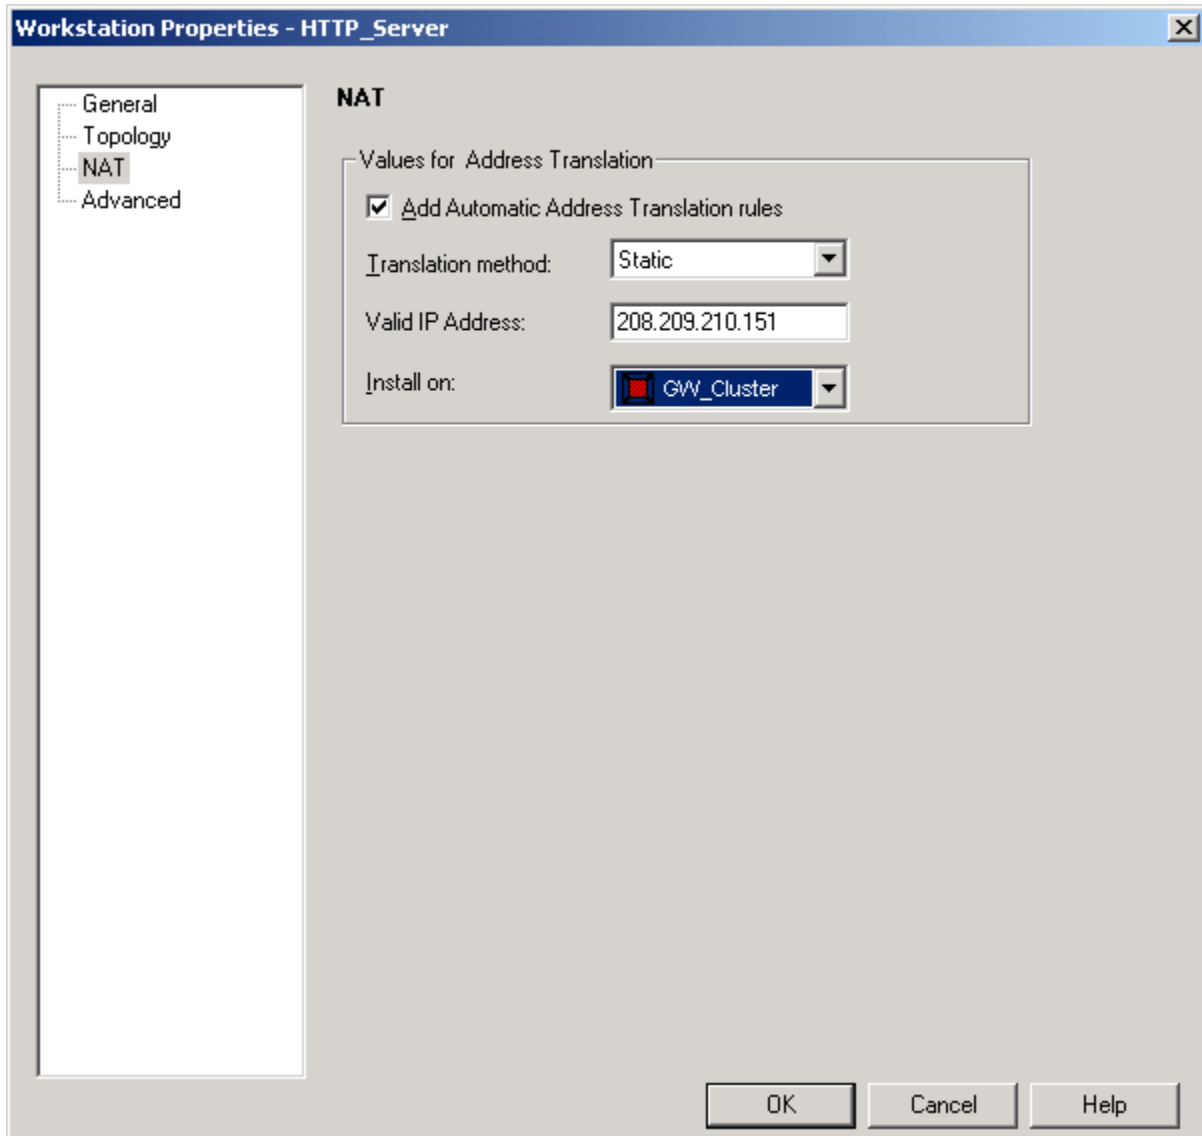
Remember, once a NAT rule has been set, the Policy must be reinstalled on the Gateway Cluster.

Static NAT

The Static NAT address is used to create one-to-one NAT relationships with internal privately addressed network hosts. The most common implementation of Static NAT is the configuration of web or mail services. This feature will allow services that exist on the Internal or DMZ network to be represented on the External side of the FireWall-1/VPN-1 modules by a single NAT IP address. For this example, a Static NAT will be created for an HTTP Server in the DMZ.

First, create an object that will require a Static NAT rule. For this example, an HTTP_Server “Workstation/Host” object was created. This object resides on the DMZ.

Next, within the objects Workstation Properties window select NAT, as illustrated in the picture below.



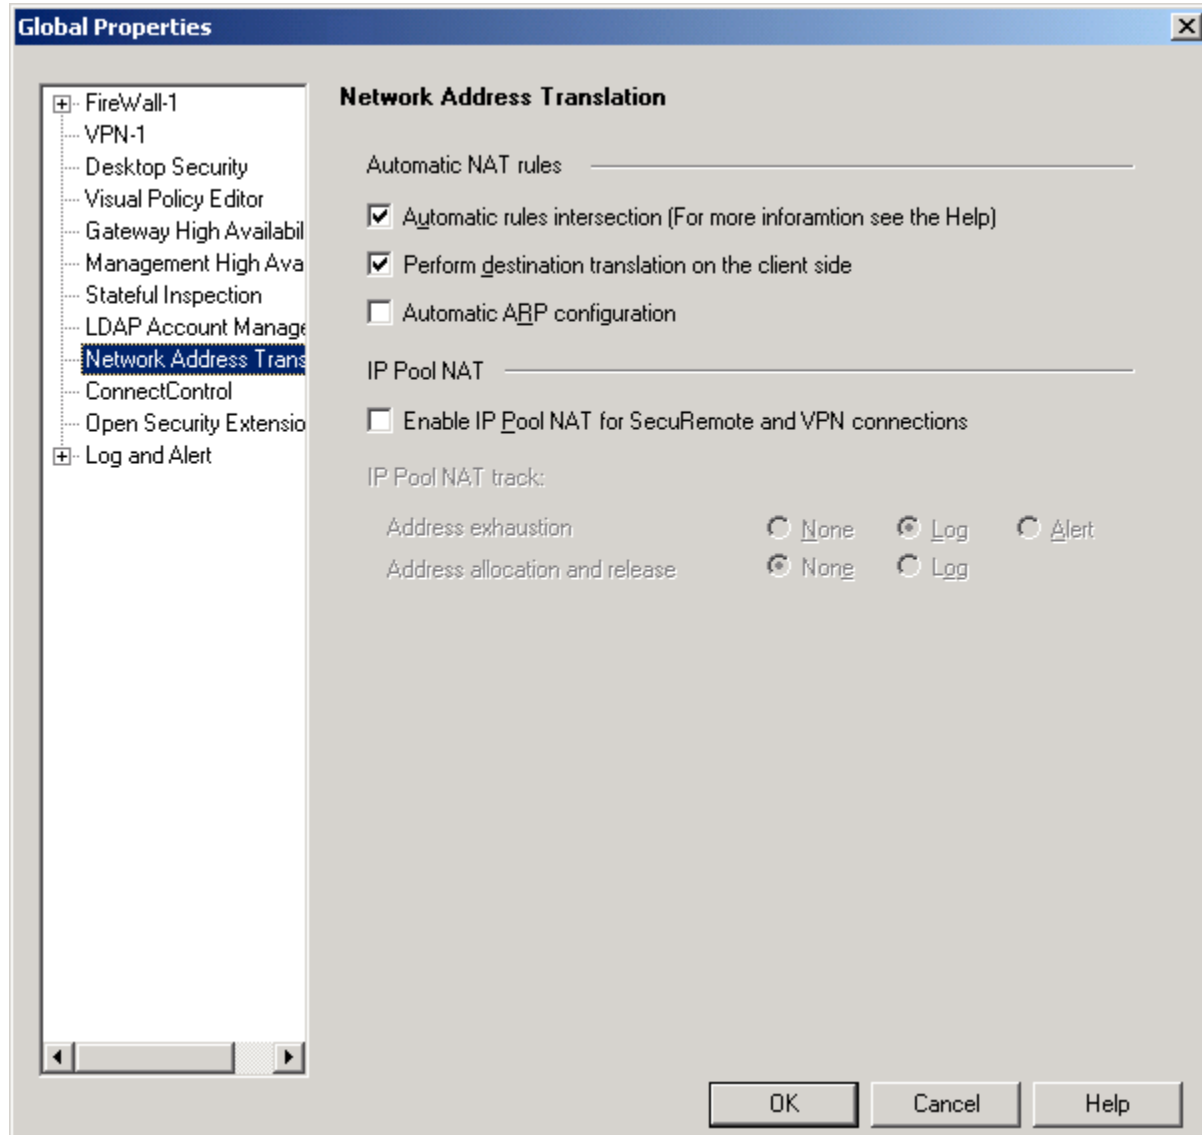
Workstation Properties – NAT Configuration

Select “Add Automatic Address Translation rules”.

Select “Static” as the Translation method and specify the Valid Address. This will be the IP address that traffic destined to, or originating from the HTTP Server will be translated to.

Finally, select the GW_Cluster object to install this NAT policy on and click OK.

NOTE The “Automatic ARP Configuration” feature must be disabled in the FireWall-1/VPN-1 Security Policy if working within a FireProof environment. It is enabled by default and can be modified within the Policy → Global Properties → Network Address Translation window illustrated below.



Global Properties – Network Address Translation Setup

FireProof Setup within NAT Environment

Once NAT addresses are added to the Gateway Cluster Security Policy, it is important to understand how NAT takes place, and how these packets are treated by the FireProof. Since NAT will be taking place on the External Interfaces of the FireWall-1/VPN-1 modules, the External FireProof is of primary concern for configuration.

For every valid NAT address that is configured within the Security Policy, a static route is required on the FireProof and Upstream router. The reason for this is simple. Since there is a single IP address shared by multiple physical devices (the enforcement modules) it is important that the flow of traffic be strictly controlled through the FireProof to insure the proper delivery of packets. This is particularly important since no enforcement module will send ARP response packets to any queries for the NAT addresses. If any one of the modules did, this could create asymmetric routing issues that would disrupt the traffic flow. This configuration utilizes the FireProofs routing table and Radware's Session Tracking functionality.

Using the NAT addresses configured within the above sections (Hide and Static) the appropriate configuration changes will be made to the External FireProof and upstream router. The traffic flow will then be explained in detail.

Two NAT addresses have been configured; 208.209.210.150 and 208.209.210.151. No enforcement modules will ARP for these addresses (see above section on disabling "Automatic ARP Configuration").

Upstream Router Setup

On the upstream router, a static "host" route must be created for each of these NAT addresses which forwards this traffic directly to the External FireProof.

The routing table should look as follows:

<i>Destination</i>	<i>Mask</i>	<i>Gateway</i>	<i>Interface</i>
208.209.210.150	255.255.255.255	208.209.210.15	208.209.210.254
208.209.210.151	255.255.255.255	208.209.210.15	208.209.210.254

Upstream Router Routing Table

As displayed in the table above, static routes have been configured directing any traffic from the Internet destined for the NAT addresses to be forwarded to the FireProof.

NOTE *When the FireWall-1/VPN-1 module first forwards the packet from its External Interface destined to a host on the Internet, the NAT'd packet is forwarded through the modules default gateway, the External FireProof. When this occurs, a FROM Client Table Entry is created so that subsequent packets for this session will be properly forwarded through the correct FireWall-1/VPN-1 module. If however, the above-mentioned static routes were not inserted into the upstream router, the router would attempt to ARP for the NAT address. Since this is disabled, it would receive no response and drop the packet. However, with the existence of the static route, the packet will be forwarded directly to the FireProof instead, allowing it to forward the packet to the appropriate module.*

FireProof Setup

Next, the FireProof must have similar static “host” routes configured to insure proper packet forwarding.

On the FireProof, insert a host route, similar to the one’s made on the upstream router, but these will have a configured next hop router of one of the enforcement modules within the farm.

<i>Destination</i>	<i>Mask</i>	<i>Gateway</i>	<i>Interface</i>
208.209.210.150	255.255.255.255	208.209.210.1	1
208.209.210.151	255.255.255.255	208.209.210.1	1

External FireProof Routing Table

In the above table, the gateway for these static host routes has been set as the external interface of GW-1 (208.209.210.1). This does not mean that every packet destined to one of these addresses will be forwarded through GW-1. This route is necessary to force the FireProof to search the Client Table for a matching session entry.

NOTE *Without the static route entry for these NAT addresses, a packet forwarded from the upstream router would be received by the FireProof, which would in turn try and ARP for the NAT address that will inevitably not be responded for. Instead, since a host route statement exists within the Routing Table of the FireProof, the FireProof will bypass the ARP resolution phase and proceed to searching the Client Table for a matching entry. If a Client Table entry is found, the FireProof then forwards the packet to the appropriate FireWall-1/VPN-1 module. However, if no existing entry exists within the Client Table a load balancing decision will be made and an a new entry will be created for subsequent session traffic.*

High Availability and Load Balancing Configuration

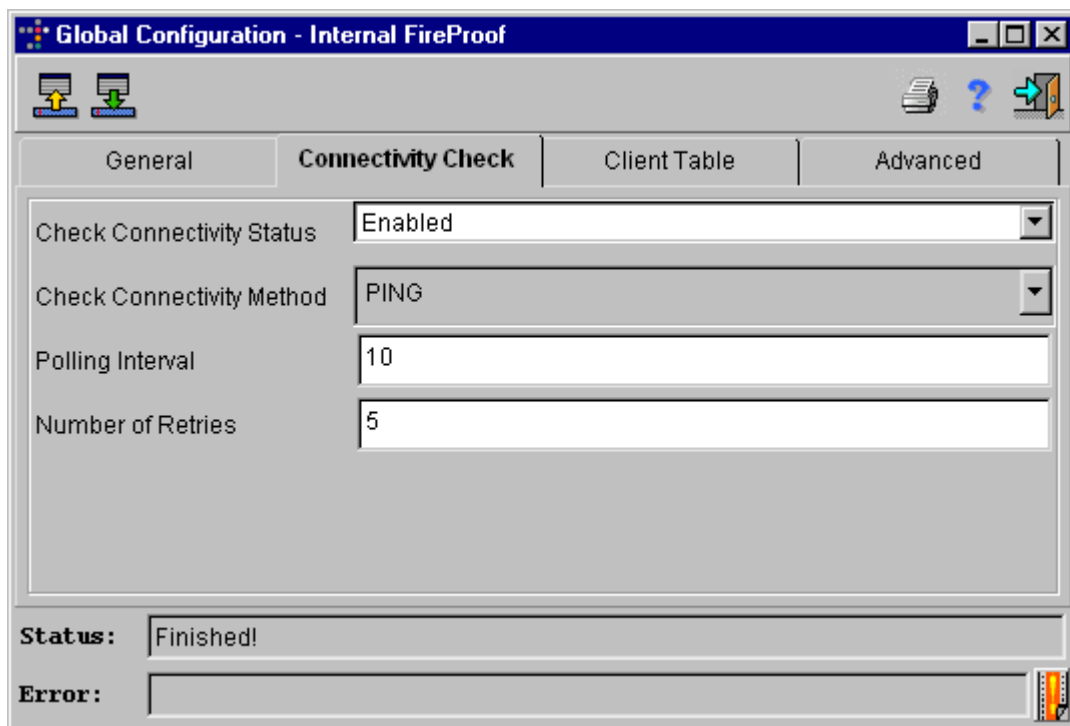
Until this section, this document has focused on the basic operating setup requirements for a Check Point FireWall-1/VPN-1 cluster within a FireProof environment. This section will address specific FireProof settings that determine and control health monitoring and load distribution of each FireWall-1/VPN-1 module within the farm as well as methods of directing traffic among these modules.

This sections of the document is broken into several sub-sections:

- Basic Connectivity Checking of FireWall-1/VPN-1 Modules
- Full Path Health Monitoring – Ability to check “through” each module
- Dispatch Methods – “Load Balancing Algorithms”
- Firewall Priority – “Server Weighting”

Basic Connectivity Checking

Basic connectivity checking is configured within the *FireProof* → *Global Parameters* window within the Connectivity Check tab.



Global Configuration – Connectivity Check

The following global parameters are configurable within this window.

Check Connectivity Status – Enabled or Disabled Globally

Check Connectivity Method (PING/TCP) – PING firewall interface or specified TCP port (FireProof initiates a TCP handshake to verify connectivity). For TCP port checking, simply input the port you wish to check in the Check Connectivity Method window.

Polling Interval – Time between tests (in seconds).

Number of Retries – Number of polling retries before device is marked “NotInService” and existing traffic is redirected to remaining firewalls.

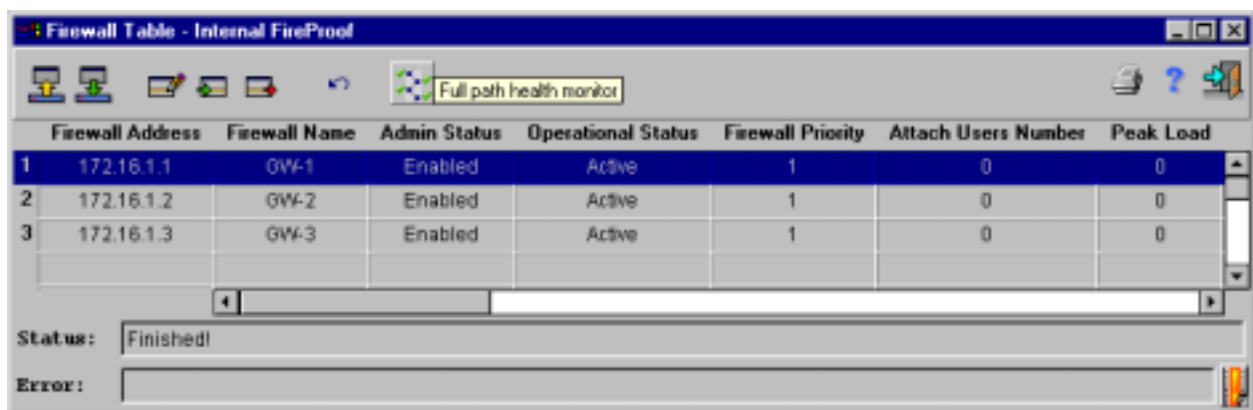
Check Connectivity Method (SNMP) - The value SNMP means the device sends SNMP GET requests to the firewall for a configured SNMP OID (Object Identifier), using a configured Community Name. The device checks the SNMP reply, if the reply does not arrive on time or if the reply indicates an error, the firewall is considered Not in Service.

Full Path Health Monitoring

The traditional connectivity checking configured in the Global Parameters window will verify connectivity to the interfaces in the Firewall Table only. Because the load balanced firewalls are responsible for routing both inbound and outbound traffic, it is important to test the firewalls ability to forward traffic through each interface as well. This type of connectivity checking is referred to as “Full Path Health Monitoring” within the FireProof and is performed in addition to the default testing of the firewall interface.


Essentially, this feature is configured on a per firewall basis and allows the administrator to define up to 10 network nodes to test for each firewall. These nodes are tested through each specified firewall with the Connectivity Check Method (PING or TCP Port) configured in the Global Parameters. If any node within the list fails a connectivity test for the specified amount of retries, the firewall is disabled within the configuration until the FireProof can successfully verify health to each test node.

To configure Full Path Health Monitoring, highlight the specified firewall from within the Firewall Table and click on the “Full path health monitor” button on the tool bar as shown below.



Firewall Table – Full Path Health Monitor Table

Clicking on the “Full Path Health Monitoring” button will open the FPHM window for the specified firewall as shown below.



Full Path Health Monitor Window

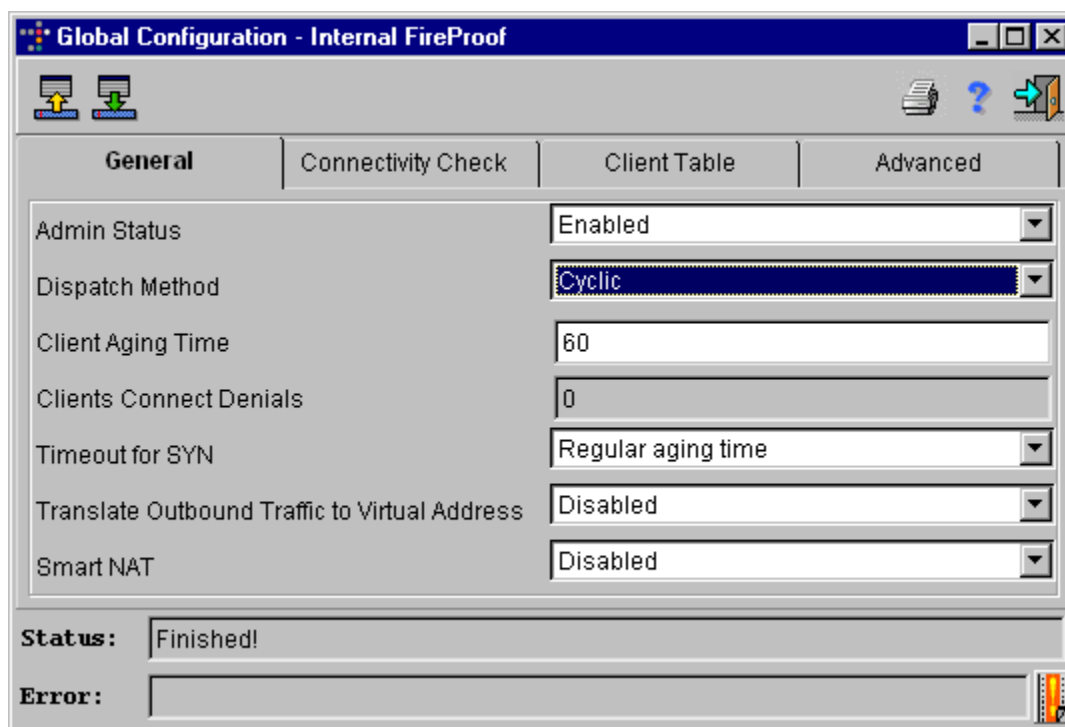
The above figure displays the FPHM table for firewall 172.16.1.1. Within this table, two additional IP addresses have been configured for testing. It is recommended for the FireProof to verify connectivity to all interfaces of the firewall (DMZ, Internal, etc.) as well as additional nodes beyond this (i.e. upstream routers, other FireProof devices, etc.).

FireProof-to-FireProof testing through the configured firewalls is highly recommended for verifying forwarding capabilities of the firewalls. For example, the External FireProof can be configured with a special Connectivity Virtual Address. The Internal and DMZ FireProof units are then configured to verify connectivity to this address through each configured firewall, thus testing the full path of traffic through the firewalls.

NOTE *It is recommended for most implementations utilizing Full Path Health Monitoring that ICMP be configured as the connectivity method within the Global Configuration. This is important since every node in the FPHM table will be tested via this method. Additionally, if the FireWall-1/VPN-1 modules have been configured with a Security Policy that restricts ICMP traffic forwarding, special rules may need to be created to allow ICMP messages sent from the FireProof to be forwarded and received through the Gateway Cluster modules.*

Dispatch/Load Balancing Methods

When a packet for a new session (one not in the Client Table already) is forwarded by the FireProof, a new load balancing decision is made. This section will describe the different metrics for distributing traffic among modules within the Firewall Table. This option is controlled within the *FireProof* → *Global Configuration* and is referred to as the “Dispatch Method”.



Global Configuration – Dispatch Method

Only one method is selectable globally. The following section outlines the available methods.

Cyclic – Unweighted round-robin distribution among active firewalls.

Least Amount of Traffic – FireProof sends new session to firewall with least real-time packets per second traffic.

Fewest Number of Users – FireProof sends new session to firewall with least concurrent sessions based on Client Table information.

Fewest Bytes Number – Traffic is distributed to firewall utilizing least amount of real-time bandwidth based on bytes per second.

Private-1/Private-2 – The Private-1 and Private-2 load balancing methods are user-customizable metrics based on SNMP information polled from the firewall modules by the FireProof. Each Private algorithm allows the administrator to configure up to two simultaneous SNMP OIDs to be polled by the FireProof. The returned variables will be used to base load balancing decisions. A common use of this load balancing technique is to poll processor and memory utilization and allow the FireProof to load balance based on

the returned information as well as concurrent connections per firewall. Remember, most firewalls disable SNMP connectivity by default requiring the addition of special security policy rules that allow the FireProof device to poll specific information.

Firewall Priority

Each firewall can be configured with a specific priority. Firewall Priority setting is often referred to as "Server Weighting". Essentially, many firewall farms may not include firewalls of the same capacity and performance abilities. In environments such as this, priorities can be established per firewall (via the Firewall Table) that affect the distribution of traffic among the farm. The priority is a numerical value, 1-99 (the higher the number, the higher the priority). Priorities set on firewalls will be used when making load-balancing decisions in all Dispatch Methods except for Cyclic. The higher the priority number configured, the higher the priority for traffic direction.

Warm-up/Recovery Timers

You can control traffic to specific firewalls that have been recently booted. This means that newly booted firewalls won't be overrun by incoming traffic in accordance with the load-balancing algorithm. You use the Firewalls Advanced Configuration window to control traffic to firewalls.

To access the Firewalls Advanced Configuration window:

From the *FireProof* menu, select *FireProof* → *Firewalls Advanced Configuration* → *Firewalls*. The Firewalls Advanced Configuration window is displayed. The Firewalls Advanced Configuration window includes the following fields:

Firewall Address - The firewall IP address.

Recovery Time - The time, in seconds, during which no data will be sent to this firewall. The time begins from the moment the first firewall is active, usually after the firewall boots.

Warm Up Time - The time, in seconds, beginning after the Recovery Time ends. During this time, clients are sent to this firewall at an increasing rate, so that the firewall can slowly reach its capacity. This option will not function in the cyclic load balancing algorithm

Firewall Grouping

This section will provide only basic Firewall Grouping information and outline several basic implementation examples. Firewall Grouping is a very comprehensive and powerful feature set and is presented within this document to provide a basic working knowledge for the VPN implementation covered later in this document. For more detailed information, please refer to the *Radware FireProof User's Manual* or the Radware document *Firewall Grouping In The FireProof* available from Radware Technical Support.

Firewall Grouping allows logical control and segregation of traffic flows through the firewall farm. The value of this feature is found in farms where multiple firewalls have separate "specialized" configurations.

There are three types of grouping rules that can be created.

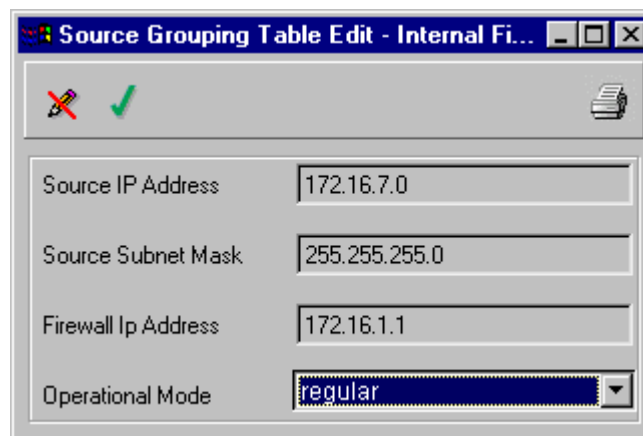
- *Destination Network* – firewalls eligible for carrying traffic to a specific destination network, as defined by an IP/mask combination.
- *Application Port* – firewalls eligible for carrying specific types of IP traffic (defined by destination TCP/UDP port number)
- *Source Network* – firewalls eligible for carrying traffic coming from a specific source network, as defined by an IP/mask combination.

The FireProof still has one global table of firewalls. However, firewall groups can be created according to the above criteria. Within each group, a subset of the global firewall table is configured. Essentially, each group contains the firewalls eligible for carrying the traffic defined by the rule. The firewalls contained within each group can also be configured as main/back-up, with this operational mode enforced inside the group.

Source Network Grouping

The following example demonstrates Source Network Grouping. From ConfigWare, this option is set within the FireProof → Firewalls Advanced Configuration → Grouping → Source Grouping.

From the Source Grouping Table, double click on an empty cell or click on the "Insert" button to open a Source Grouping Table Edit window.

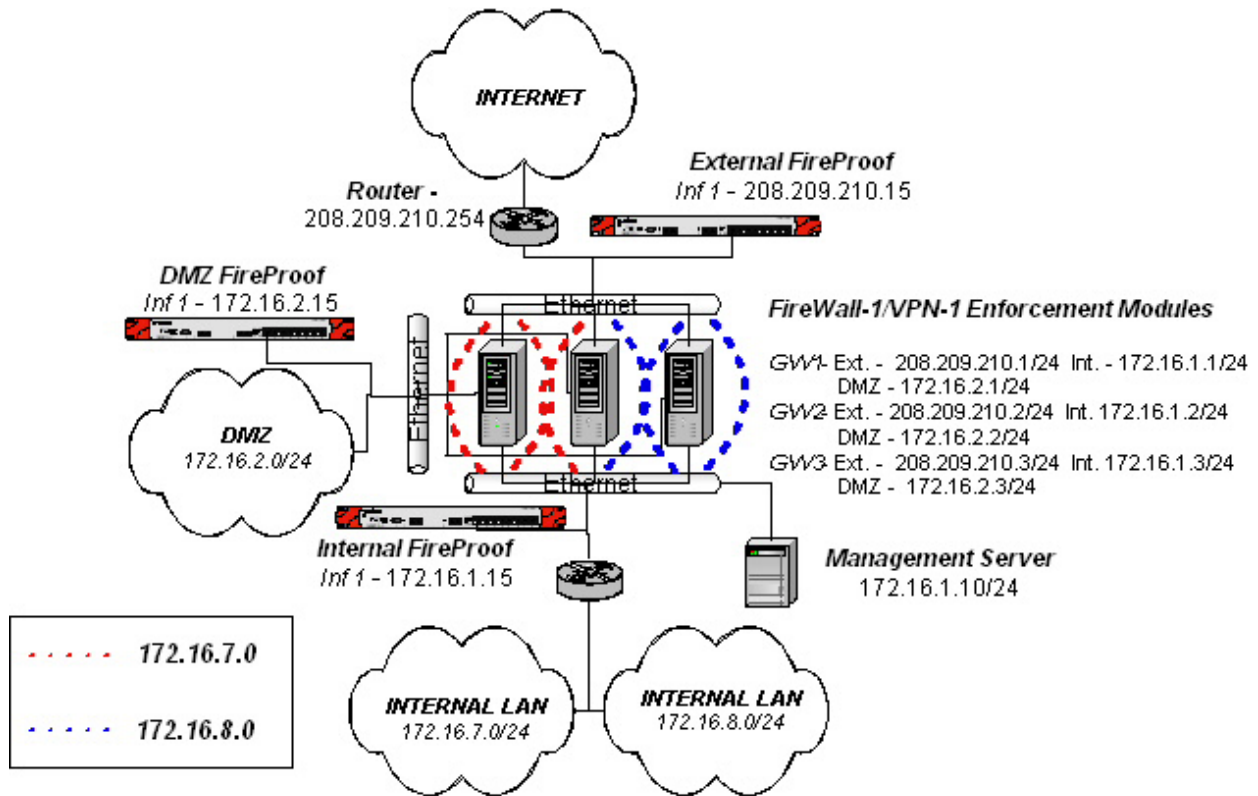


Source IP Address	172.16.7.0
Source Subnet Mask	255.255.255.0
Firewall Ip Address	172.16.1.1
Operational Mode	regular

Source Grouping Table Edit Window

Specify a source network, mask, appropriate firewall, and the firewall's mode for this classification.

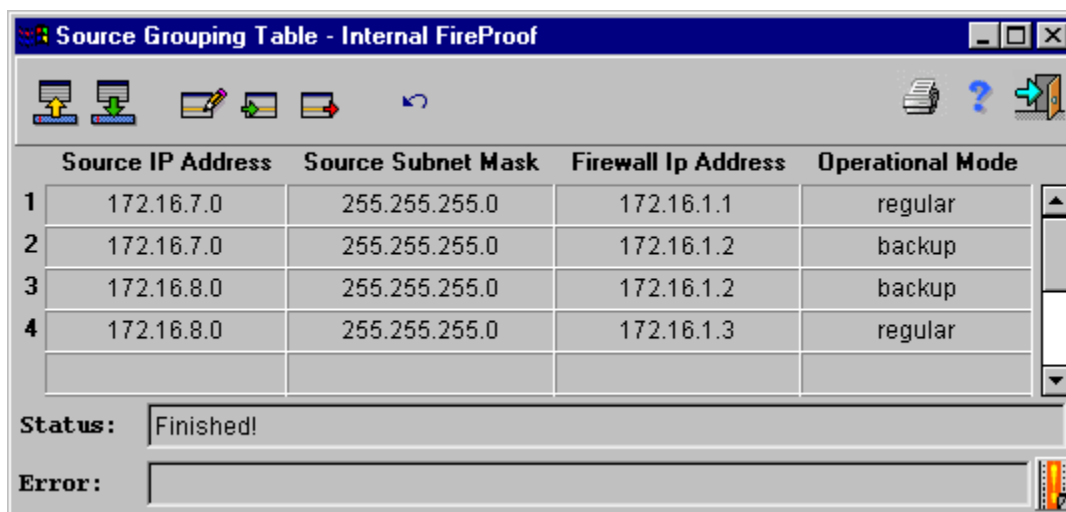
The following diagram demonstrates a very simple Source Grouping example. The goal of this implementation is to forward traffic from the 172.16.7.0 source network through GW-1 and traffic from the 172.16.8.0 network through GW-3. Additionally, GW-2 will be configured as backup for both types of traffic.



Source Grouping Implementation

The above example demonstrates the segregation of traffic based on Source IP/Mask information between firewalls within the Global Firewall Table.

The following table demonstrates the rules established for this solution.



	Source IP Address	Source Subnet Mask	Firewall Ip Address	Operational Mode
1	172.16.7.0	255.255.255.0	172.16.1.1	regular
2	172.16.7.0	255.255.255.0	172.16.1.2	backup
3	172.16.8.0	255.255.255.0	172.16.1.2	backup
4	172.16.8.0	255.255.255.0	172.16.1.3	regular

Status: Finished!

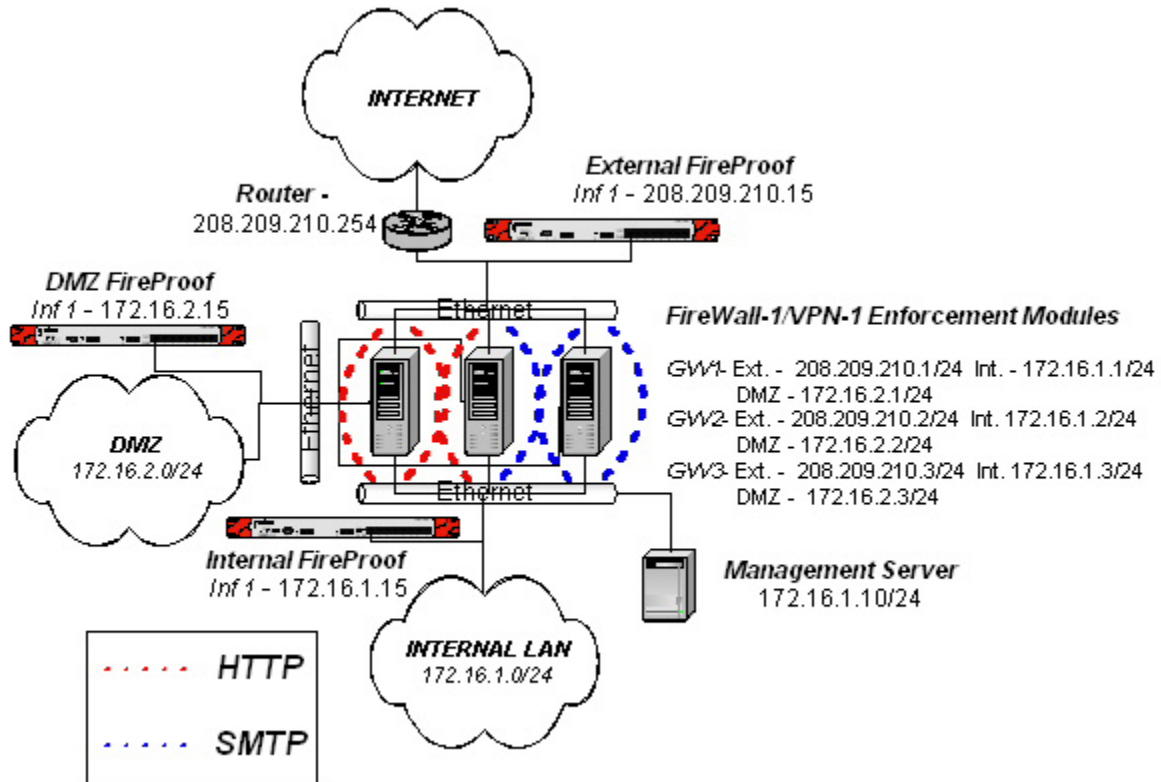
Error:

Source Grouping Table

All traffic destined through the FireProof will be compared to the rules listed in the Grouping tables. Traditionally, multiple "regular" firewalls are established for each rule, for instance traffic from the source network 172.16.7.0 could have both GW-1 and GW-2 specified as "regular". If a packet matches this rule, the best firewall is chosen among the valid firewalls. However, if the packet does not match any rule within the table, it is simply forwarded to the best firewall within the Global Firewall Table (any firewall becomes eligible).


Application Grouping

The Application Grouping feature allows the administrator to determine the flow of traffic based on an application instead of only source or destination IP information. This feature is setup similar to the Source IP Grouping example in the previous section. Utilizing a similar network architecture as the preceding example, the following diagram will demonstrate the segregation of traffic flows based on Application Port information.

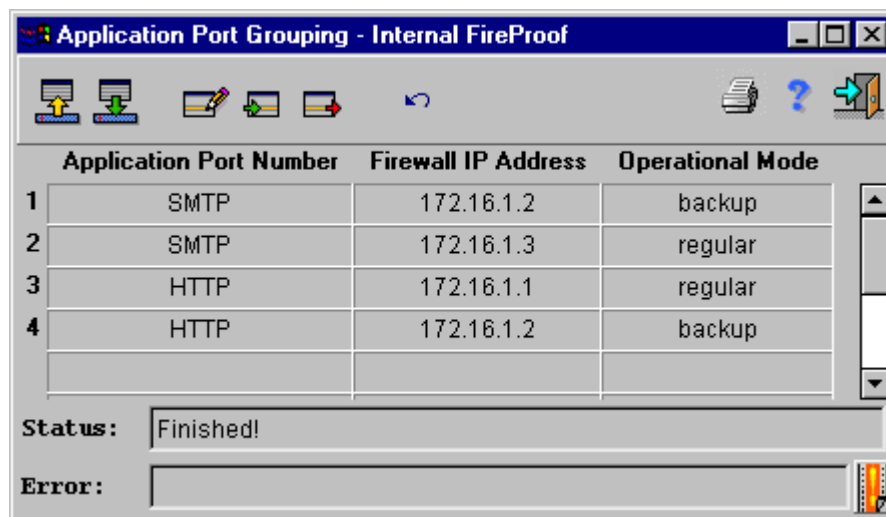


Application Port Grouping Implementation

This feature, like Source and Destination Grouping, must be configured by assigning application ports to specific firewall(s) within the Global Firewall Table. From ConfigWare, the Application Port Grouping Table must be opened via FireProof → Firewalls Advanced Configuration → Grouping → Application Port Grouping. From this table, double click on an empty cell or click the “Insert” button to display the Application Port Grouping Edit window as shown below.



The implementation diagram above outlines that HTTP traffic be configured to use GW-1 as regular and GW-2 as a backup. Additionally, SMTP traffic will be configured to use GW-3 as regular with GW-2 as a backup. Once all entries for this configuration have been made within the Application Port Grouping Table, the table will look as follows:



The screenshot shows a window titled "Application Port Grouping - Internal FireProof". It contains a table with the following data:

Application Port Number	Firewall IP Address	Operational Mode	
1	SMTP	172.16.1.2	backup
2	SMTP	172.16.1.3	regular
3	HTTP	172.16.1.1	regular
4	HTTP	172.16.1.2	backup

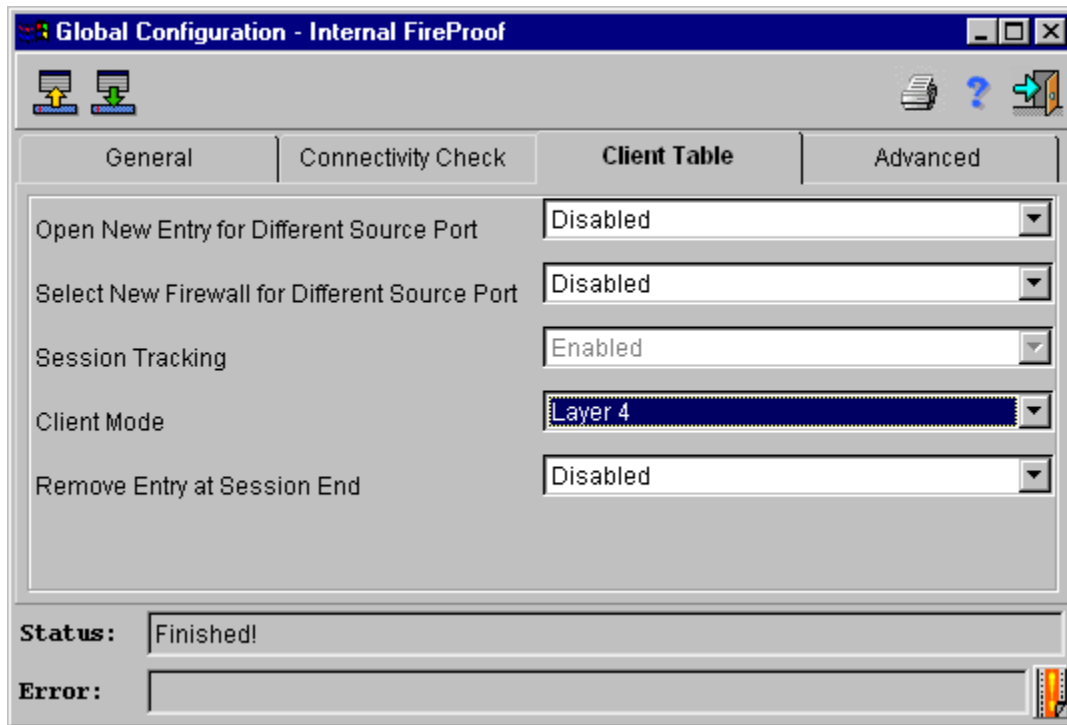
Below the table, the "Status:" field displays "Finished!". The "Error:" field is empty.

Application Port Grouping Table

This configuration now segregates HTTP traffic through GW-1 (as long as it is healthy), and SMTP traffic through GW-3 (as long as it is healthy as well). Both protocols have been configured to use GW-2 if either firewall is unavailable (backup mode). As with Source Grouping, every packet is analyzed by the FireProof and compared to the rules within the Grouping Table. If no match is found, any available firewall from the Firewall Table is eligible to receive an unmatched packet.

Client Table Mode

Application Port Grouping requires the FireProofs Client Table mode be set to Layer 4. This enables the Client Table to keep track of Layer 4 port information, which is essential to the Application Port Grouping feature. This feature is enabled from the Global Configuration window (FireProof → Global Configuration) under the Client Table tab as shown below.



Global Configuration Window – Client Table Settings

Rule Combinations

The grouping feature allows for combinations of several types of rules in order to control traffic through the firewall farm. For instance, an application grouping rule that specifies all HTTP traffic is configured to pass through GW-1 and GW-2, and a source IP grouping rule that specifies that traffic from the 172.16.1.0 network must pass through GW-1 are configured.

All traffic is analyzed and compared to the rules in the table, so if an HTTP packet is sent from a client on the 172.16.1.0 network, the FireProof compares the packet to the rules configured in the Grouping Table. The FireProof will determine that the packet fits both the Application Grouping rule and the Source IP rule and deduce that the only common firewall between the rules is GW-1 and forwards the packet appropriately.

Much of the Firewall Grouping feature's flexibility and benefits are derived from its ability to match and classify packets using multiple types of rules simultaneously. As mentioned in the beginning of this section, Firewall Grouping is extremely advanced and can be quite complicated depending on the rules being implemented. Please refer to the *Radware*

FireProof User's Manual or the Radware document *Firewall Grouping In The FireProof*, available from Radware Technical Support.

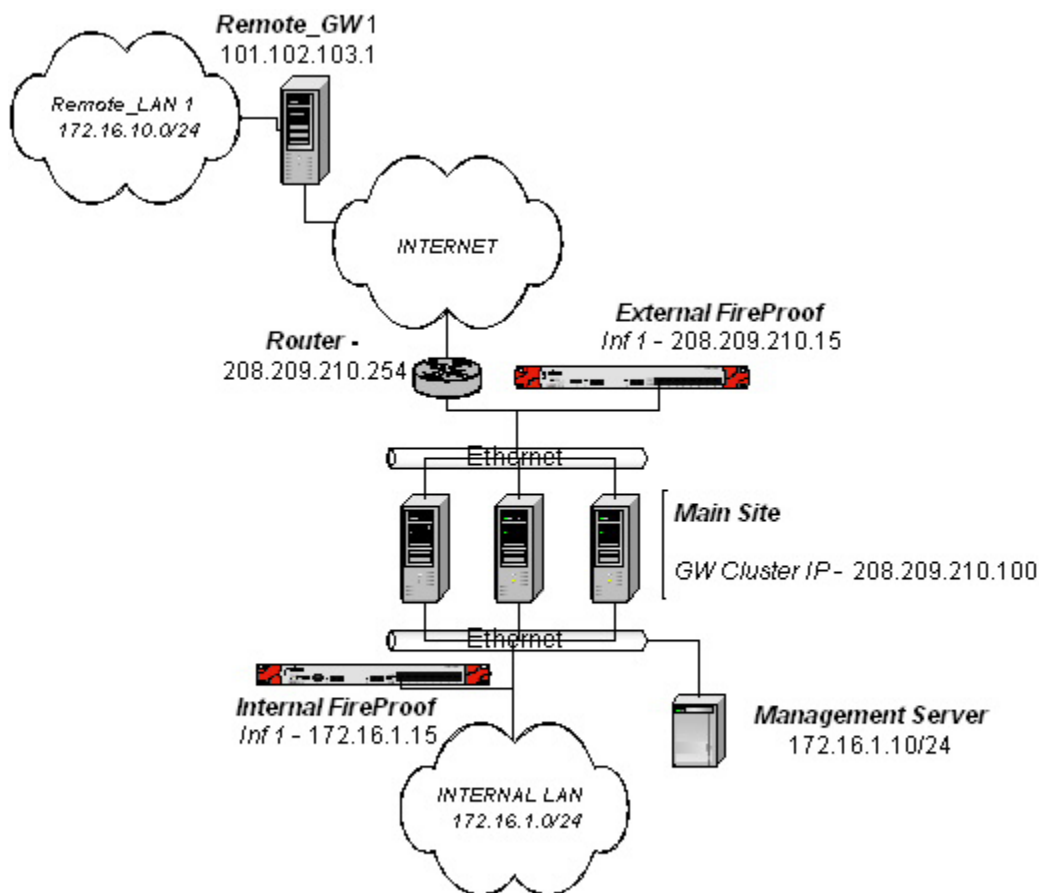
Virtual Private Network Setup

This section will outline several load balanced VPN-1 configurations utilizing the FireProof. Included in this section are setup and configuration details for:

- Gateway-to-Gateway Cluster Communication – A common VPN scenario involving a “main/central” site consisting of a FireProof load balanced cluster of VPN-1 modules and several stand-alone “remote” modules.
- SecuRemote-to-Gateway Cluster – Configuration of SecuRemote clients
- SecureClient-to-Gateway Cluster – Configuration of Policy Server and Desktop Security Policies

Gateway-to-Gateway Cluster Scenario

The following diagram will be used to demonstrate VPN setup for secure communications between hosts on Remote_LAN1 at the “remote” site, and the Internal_LAN at the “main” Gateway Cluster location. While this section focuses on the deployment and communication of a standalone remote gateway, the setup information within this section directly applies to Gateway Cluster-to-Gateway Cluster implementations as well.



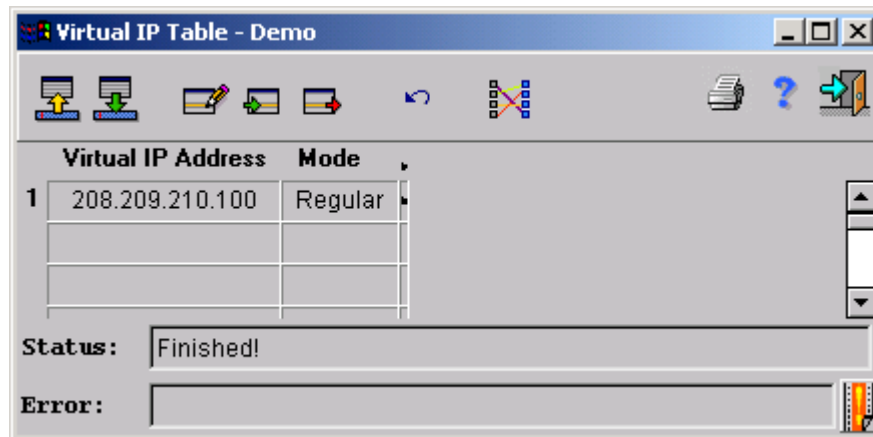
Gateway-to-Gateway Cluster VPN Scenario

FireProof Configuration

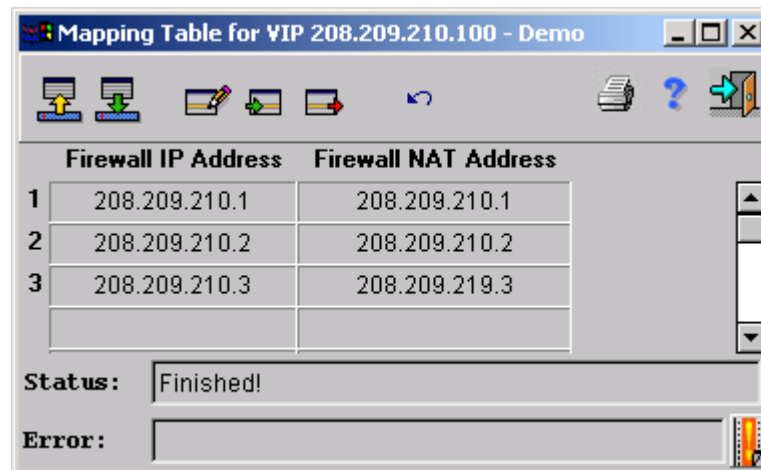
Only the External FireProof will require additional configuration. All that is required is the setup of a Virtual IP Address (VIP). This VIP must be configured with the same IP address as the gateway cluster. Additionally, this VIP is then mapped to each of the external interfaces of the FireWall-1/VPN-1 modules.

Configure Virtual IP Address

In this example, the VIP will be configured as 208.209.210.100. This is configured through ConfigWare (FireProof → Virtual IP).



Next, this VIP must be mapped to the each external interface of the configured FireWall-1/VPN-1 modules. From the Virtual IP window, the VIP just configured must be highlighted and the VIP Mapped Table button clicked.




From this table, insert each FireWall-1 interface (the NAT field should be configured with the same address as the external interface for each specific firewall).

Policy Configuration at the Main Gateway Cluster Location

Remote Network Object Creation

Create “encrypted_external” network object for the network protected by the Remote_Gateway.



Network Object Insert Window

This network object will be used in the Policy Rule Base when specifying what traffic will be encrypted with the VPN tunnel.

Remote Workstation/Gateway Object Creation

Create a Remote_Gateway object.

The screenshot shows the 'Workstation Properties - Remote_Gateway' dialog box. The 'General' tab is active in the left sidebar. The main area contains the following fields and options:

- Name:** Remote_Gateway
- IP Address:** 101.102.103.1 (with a 'Get address' button)
- Comment:** (empty text box)
- Color:** (black color selection box)
- Type:** Host Gateway
- Check Point Products:**
 - Check Point products installed: Version NG (with a 'Get Version' button)
 - VPN-1 & FireWall-1 (checked)
 - FloodGate-1 (unchecked)
 - Policy Server (unchecked)
 - Management Station (checked)
- Object Management:**
 - Managed by this Management Server (Internal)
 - Managed by another Management Server (External)
- Interoperable VPN Device

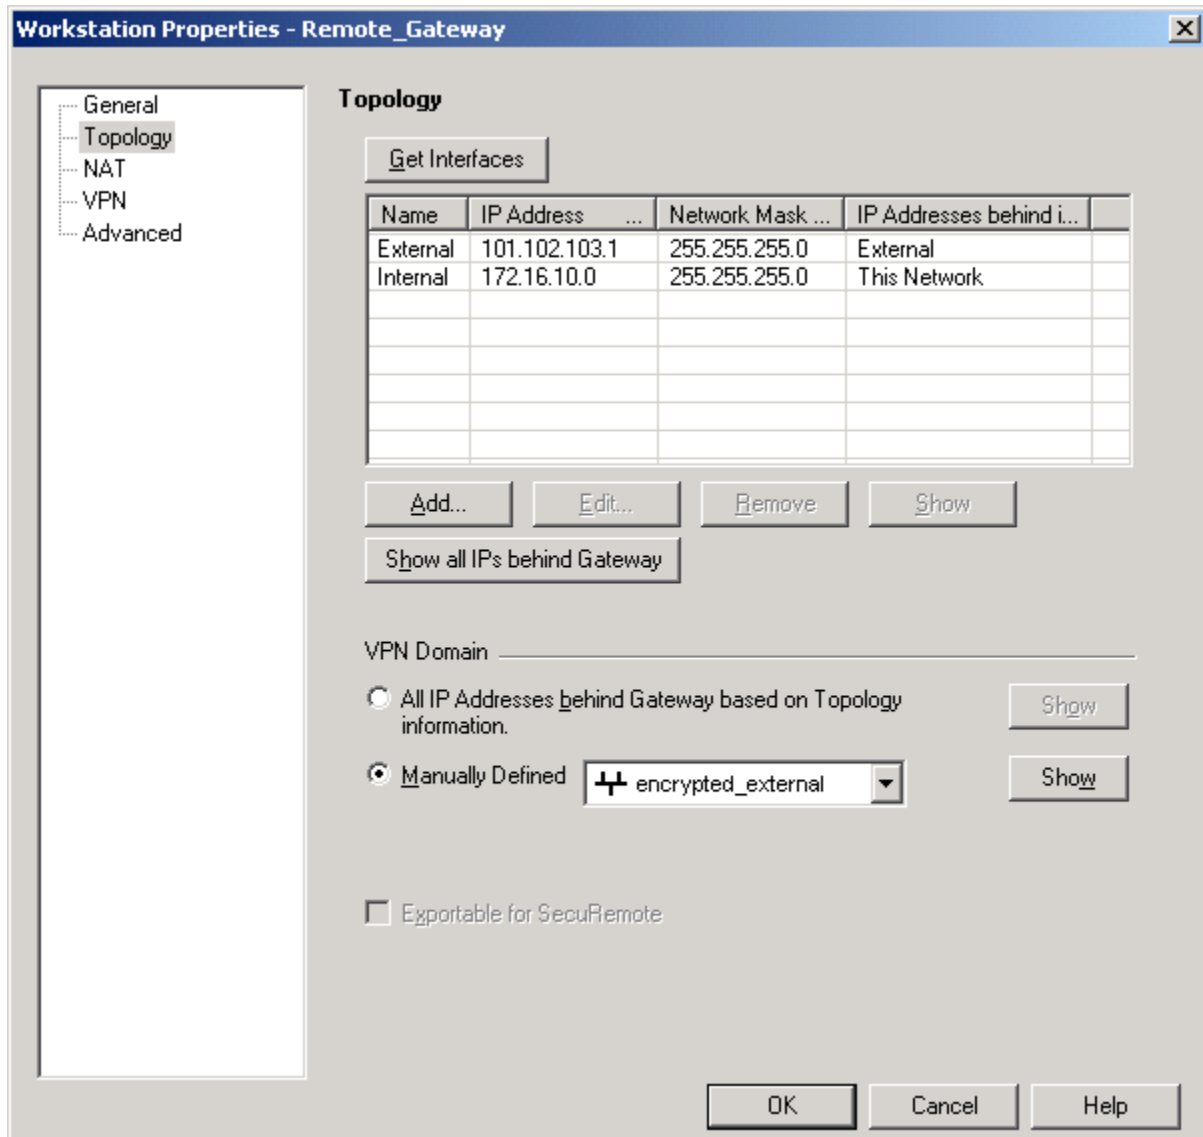
At the bottom right, there are 'OK', 'Cancel', and 'Help' buttons.

Workstation Properties Window

Within the Workstation Properties window for the Remote_Gateway object, insure that the object is specified as a "Gateway" in the Type field and that the correct products installed list is updated.

Topology Setup for Remote_Gateway Object

Within the Workstation Properties of the Remote_Gateway object, specify relevant interface topology information as well as the VPN Domain for this object. The domain can be manually configured as the encrypted_external network object created in the first step.

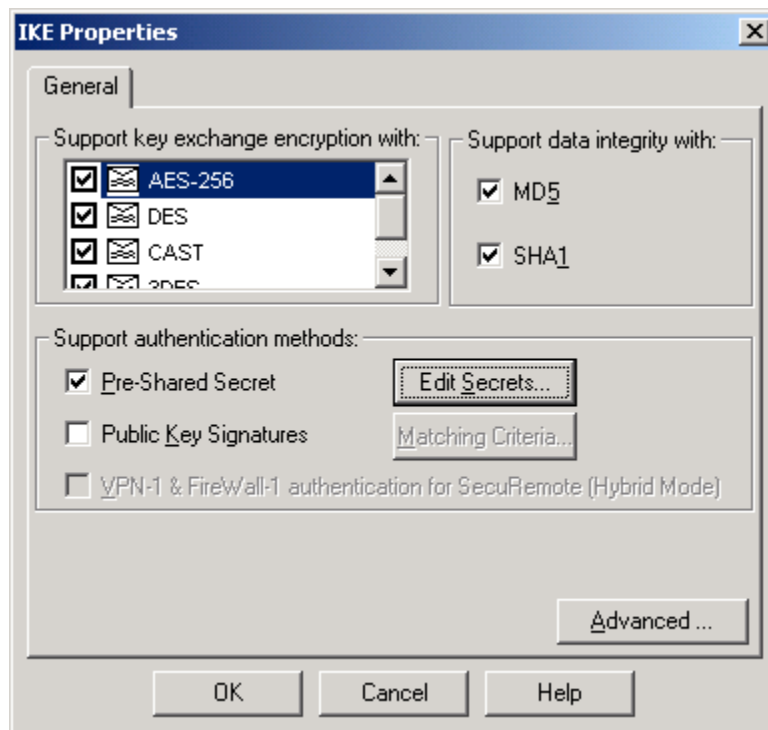


Workstation Properties – Topology Information

Encryption Schemes

From the Workstation Properties of the Remote_Gateway object, click on VPN. This will give you two Encryption Scheme options. Both IKE and FWZ are supported within a Radware FireProof configuration; however, this document will focus only on IKE.

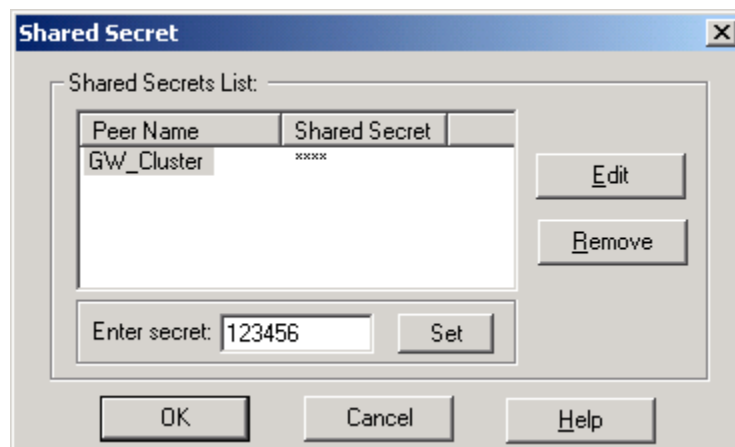
Highlight and check the IKE box and click on the "Edit..." button. This will display the IKE Properties window, as shown below.



IKE Properties

This example implementation will not utilize certificates. Instead, IKE can be configured with a Pre-Shared Secret at each location. This Pre-Shared Secret is configured at each site, and is used during the initial phases of a VPN setup between two gateways.

Check the Pre-Shared Secret box and click on the "Edit Secrets..." button. This action will display the Shared Secret window.

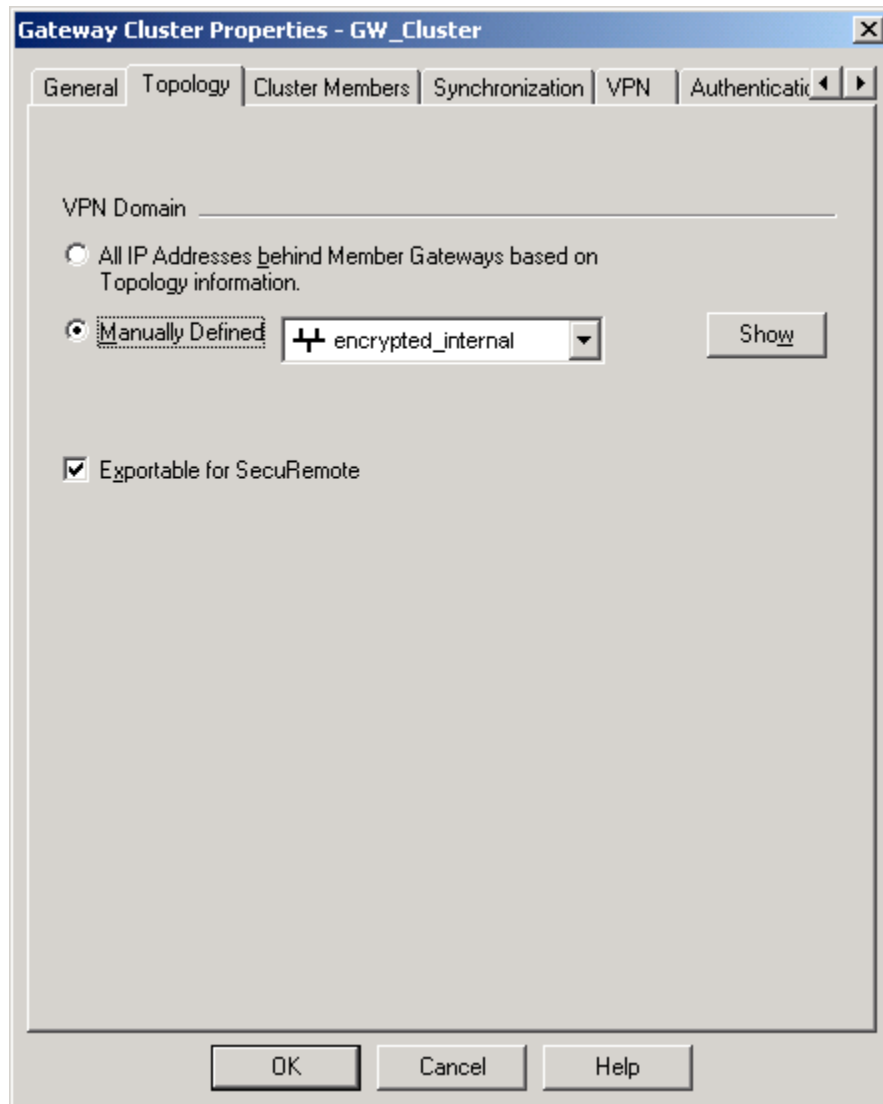


Shared Secret Window for Remote_Gateway Object

In this window, you will find the GW_Cluster peer name. Highlight this name and click on the "Edit" button. Specify the secret to be used. Remember, this secret must be at least six characters and must be configured the same at both sites.

VPN Topology Setup for GW_Cluster Object

The VPN Domain must be specified within the Topology Tab of the Workstation Properties window of the GW_Cluster object.



Gateway Cluster Properties - Topology

Define the local encrypted_internal network object as the VPN Domain for the GW_Cluster object.


IP Pools for SecuRemote and SecureClient Connections

It is recommended that IP Pools (Unique IP address ranges) be created and assigned to each individual Gateway within the cluster. IP Pools are used for incoming SR/SC connections and ensure traffic and failover persistence. For more information regarding IP Pool setup and implementation, please see CheckPoint User documentation.

Encryption Scheme Setup for GW_Cluster Object

Next, the Encryption Scheme information will need to be configured for the GW_Cluster object. To do this, highlight the GW_Cluster object from the Network Objects window. Click on the VPN tab at the top of the window. Highlight and check IKE and click on the "Edit..." button.

Again, check Pre-Shared Secret for this object and proceed to the "Edit Secrets..." button.



Shared Secret Window for GW_Cluster Object

Within the GW_Cluster object Shared Secret window, the Remote_Gateway object exists within the Peer Name column. Specify the shared secret for this object if not already populated. This should be the same secret that will be specified at each site.

Creating the Rule Base for Encrypted Traffic

Once all of the Gateway and Network objects that will be participating in a VPN implementation are setup at the main clustered site, rules must be created for each enforcement module to know how to handle specific traffic.

The rules required for this example are very simple. Essentially, there are two networks; encrypted_internal, at the main site, and encrypted_external at the remote site. During the Workstation/Gateway configuration, these network objects were defined as the VPN Domains for their respective gateways.

For this implementation, only two rules are required as demonstrated below.

NO.	SOURCE	DESTINATION	SERVICE	ACTION	TRACK	INSTALL ON
1	encrypted_internal	encrypted_external	* Any	Encrypt	Log	Gateways
2	encrypted_external	encrypted_internal	* Any	Encrypt	Log	Gateways

Rule Base for Basic Encrypted Communication Between Both Sites

1. The first rule assures all communication from the internal network to the external network is encrypted. Because the traffic is destined to a VPN Domain of a known Gateway Object, the Enforcement Module knows what secret key to use to establish encrypted communication.
2. The second rule allows encrypted traffic from the external network to be accepted by the Enforcement Modules, while at the same time insuring that the communication remains encrypted.

NOTE VPN connection information is synchronized between the firewalls using the synchronization feature built into FireWall-1/VPN-1. However, during the initialization of an IPsec connection there is a short window when the IPsec connection information has not been synchronized to the remaining FireWall-1 gateways. If the firewall handling the IPsec connection fails during this period of time, the VPN connection will failover to an active firewall. This active firewall, not having synchronized the IPsec connection information prior to the first firewall failing, may allow the information to go out unencrypted. The preceding rules will force the encryption of all information between an internal and external network. More information regarding the Synchronization feature and VPN failover is included in the VPN Implementation Details Section.

Policy Configuration at the Remote Standalone Gateway Location

Configuration at the Remote_Gateway site is very similar to the configuration at the main GW_Cluster location. Since the creation of objects has already been covered in this document, this section will easily outline what objects and rules need to be created on the Remote_Gateway device.

- Setup Network Objects – two should exist:
 - encrypted_internal – 172.16.10.0/24
 - encrypted_external – 172.16.1.0/24
- Setup Gateway Objects – two should exist:
 - Local_Gateway – 101.102.103.1
 - GW_Cluster – 208.209.210.100 – Use Cluster IP Address
- Setup VPN Domain within Workstation/Gateway Object Topology Setting
 - Local_Gateway – encrypted_internal
 - GW_Cluster – encrypted_external
- Setup Encryption Scheme – IKE Information
 - Preshared secret for each object should be same as specified at main site
- Create Encryption Rules – Should be same as main site:

encrypted_internal	to	encrypted_external	Any	Encrypt
encrypted_external	to	encrypted_internal	Any	Encrypt

SecuRemote Scenario

FireWall-1/VPN-1 Gateway Cluster Site Setup

1. Setup User Account
2. Set authentication information for this specific user
3. Add the created user to a group
4. Create a rule within the Gateway Cluster Policy that specifies the following

Source/Destination/Type/Action

User_Group/Internal_Network/Any/Client Encrypt

5. Within the GW_Cluster Object Topology – Ensure that the “VPN Domain” section is selected as Export for SecuRemote

SecuRemote Client

1. Install SecuRemote/SecureClient on client machine. After installation, the program will prompt for installation of Desktop Security. Desktop Security is only required within SecureClient solutions in which a Policy Server is utilized to issue Desktop Security Policies.
2. Once SecuRemote is successfully installed on client. You must create a new “site”. From SecuRemote, click on Sites → Create New... From this menu, insert the IP address of the sites Gateway Cluster IP (208.209.210.100).
3. Once a connection is established it will be necessary to provide authentication information via whatever means specified within the Policy configuration.

SecuRemote should now be successfully configured.

SecureClient Scenario

The use of SecureClient requires the installation and setup of a Policy Server. The Policy Server should not be configured as any member of the Gateway Cluster, and should be installed as a separate device. Support for a Policy server as an integrated cluster member is slated for Check Point NG Feature Pack 1.

Essentially, a SecureClient will first connect to the Policy Server. This policy server will authenticate the user as well as the machine by verifying configuration. Additionally, a Desktop Security Policy is configured to ensure additional security for SecureClient connections to the load-balanced site.

Once the client is authenticated and a Desktop Security Policy in place, the SecureClient will then connect to the site through the Cluster IP. This insures high availability and load balancing of established remote connections.

FireWall-1/VPN-1 Gateway Cluster Site Setup

1. Install the FireWall-1/VPN-1 Enforcement Module and Policy Server on a designated server.
2. Configure this gateway object appropriately within the policy verifying that Policy Server is included in the "Check Point Products Installed" configuration of the object.
3. The same user and rule setup should be followed as with a SecuRemote setup.
4. A Desktop Security tab exists within the Policy Editor rule base. Within this table, an access rule should be specified for the specific remote clients utilizing SecureClient. For this example a simple User_Group/Internal_Network/Any/Accept rule can be configured allowing access to any device within the Internal_Network from users within the group specified in the source.

SecureClient Setup

1. Install SecuRemote/SecureClient on client machine. After installation, the program will prompt for installation of Desktop Security. Desktop Security should be installed since it is required for SecureClient solutions in which a Policy Server is utilized to issue Desktop Security Policies as outlined in the above section.
2. Once SecureClient is successfully installed on the client machine. You must create a new "site". From SecureClient, click on Sites → Create New... From this menu, insert the IP address of the sites Gateway Cluster IP (208.209.210.100).
3. Once a connection is established it will be necessary to provide authentication information via whatever means specified within the Policy configuration.
4. Next, a message should be displayed that a Policy Server is associated with this connection and should be contacted for the appropriate Desktop Security rules.

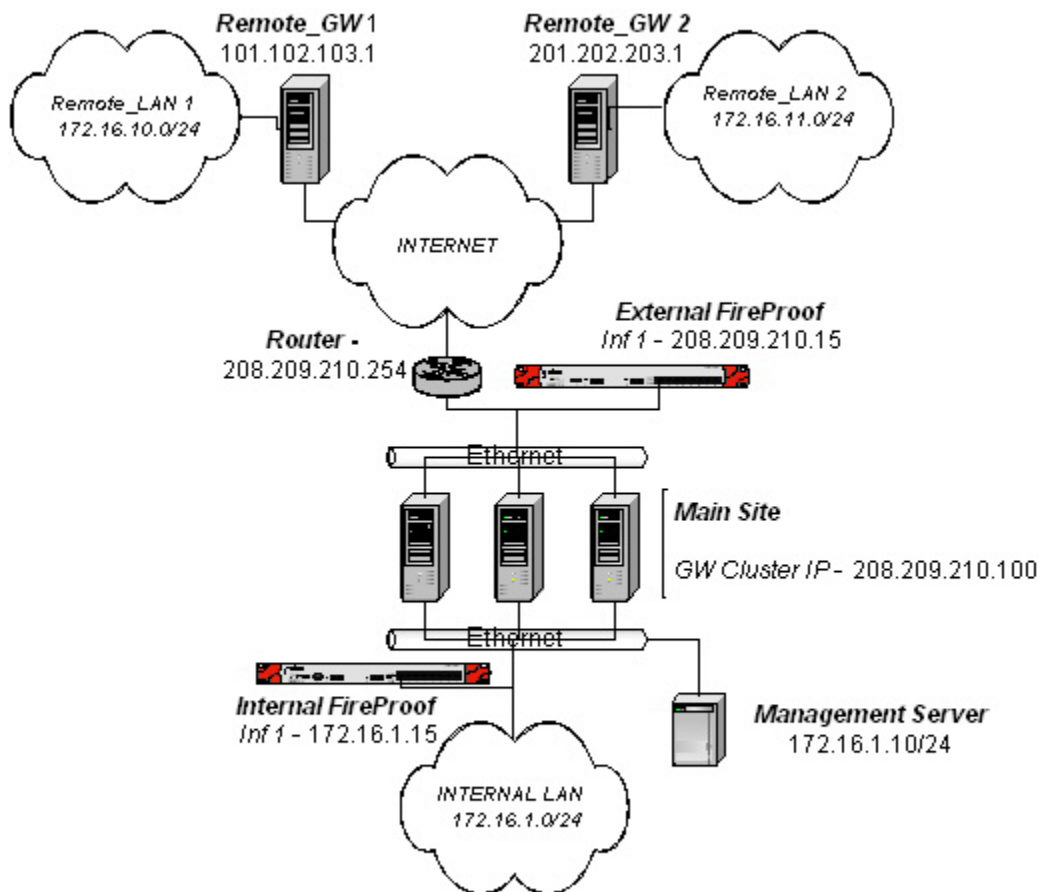
SecureClient should now be successfully configured.

VPN Implementation Details

In the VPN setup section, it was explained that enforcement modules are not always able to replicate session and IKE key information between devices quickly enough, and in some instances this can create issues in case of module failure, or asymmetric routing through different cluster members. The potential delay in key synchronization between load balanced enforcement modules and the way in which sessions are distributed among them by the FireProof can in some instances create problems regarding VPN traffic to and from the same remote sites.

This section will focus on specific issues relating to Gateway Cluster with VPN deployment. It will also outline and setup guidelines to insure a properly functioning solution.

Gateway Cluster VPN Traffic Flows



Gateway Cluster Environment with Multiple Remote VPN Sites

The diagram presented above will be used throughout this section to outline potential implementation issues as well as to define specific setup criteria for similar environments utilizing the Gateway Cluster with VPN.

Consider an outbound session originating from a host in the Internal_LAN destined to a host of the configured Remote_LAN networks.

1. A packet is forwarded through the Internal FireProof (source 172.16.1.55) destined to a host on Remote_LAN2 (172.16.11.100). The FireProof makes forwarding decisions based on the existence or non-existence of an existing entry in the Client Table. For this example, this is a new session, and thus doesn't have an existing entry. This packet will be load balanced among the available firewalls in the farm. GW-1 will be selected for this session and an entry will be created in the Client Table, insuring all subsequent packets for this session are forwarded through the same firewall.
2. The packet is forwarded to GW-1, which is configured to negotiate a VPN connection with the Gateway responsible for traffic destined to the 172.16.11.0 network, Remote_GW2. While the tunnel is established, the External FireProof has placed an entry in its Client Table tracking the traffic for this specific session, once again insuring that all traffic associated with this connection is properly forwarded through the correct firewall, GW-1.

Up to this point, communication between both sites should be working properly. For this example, problems will begin to exist when additional hosts within the Internal_LAN initiate sessions to the same destination network, Remote_LAN2. The following traffic flow explains this problem, and is based on the session information presented in steps 1 and 2.

3. Another host from the Internal_LAN network initiates a session to a host on the Remote_LAN2 network. Again, the hosts packet will route through the Internal FireProof. Because this is a session from a new host, a Client Table entry will not exist, requiring that the FireProof make a decision based on current load and availability. One of two things can occur:
 - a. The Internal FireProof selects GW-1 again and forwards the packet on. GW-1 will receive this packet and identify that a VPN tunnel should be associated with this traffic through Remote_GW2. Since a tunnel exists already, from the previous session, the packet is forwarded along and successful communication remains intact, or...
 - b. The Internal FireProof selects another firewall other than GW-1. For this example, it chooses GW-2. The packet is then forwarded to GW-2, which analyzes it and determines that a VPN tunnel should be established with Remote_GW2 for this traffic. GW-2 then begins initiating a new tunnel to Remote_GW2. Both, GW-2 and Remote_GW2 negotiate a new set of keys for this VPN tunnel. Now, VPN traffic that was previously forwarded from GW-1 to Remote_GW2 no longer functions properly since the set of keys has been changed from their original negotiation.

Essentially, if a tunnel is established through one firewall to a remote gateway, and then another firewall within the cluster attempts to establish a tunnel to the same gateway, the original tunnel that was created will no longer be valid since valuable key information would have changed at the remote gateway site.

Check Point's synchronization feature addresses this issue very well by insuring that all enforcement modules within a gateway cluster have the same key information. However, synchronization between enforcement modules is not always instantaneous and varies based on load and traffic. Since a delay does exist, if a session, as described in example 3b occurs before key information from the first tunnel can be synchronized, the same problems can occur.

VPN utilizing FireProof Source & Destination Grouping

The solution to this problem can be addressed with the FireProof's Source and Destination Grouping feature. By utilizing this feature the configuration can provide high availability for VPN based traffic while still providing load balancing to any non-VPN traffic. Essentially, rules will be configured on the Internal and External FireProof. These rules will insure that all traffic destined to and originating from the same Remote_Gateway sites will be properly forwarded through the firewalls of a Gateway Cluster. Following is an outline of the rules required:

- Destination Rules on the Internal FireProof. These rules include the destination networks secured by remote gateways, and a firewall within the farm to map this traffic to.
- Source Rules on the External FireProof. These rules correspond to rules created on the Internal FireProof, but map the actual remote gateway addresses to the firewalls within the farm.

Using the diagram presented in the beginning of this section, the following rules will be created.

Internal FireProof – Destination Rules

<i>Destination</i>	<i>Mask</i>	<i>Firewall</i>	<i>Mode</i>
172.16.10.0	255.255.255.0	GW-1	Regular
172.16.10.0	255.255.255.0	GW-2	Backup
172.16.11.0	255.255.255.0	GW-3	Regular
172.16.11.0	255.255.255.0	GW-2	Backup

External FireProof – Source Rules

<i>Source</i>	<i>Mask</i>	<i>Firewall</i>	<i>Mode</i>
101.102.103.1	255.255.255.255	GW-1	Regular
101.102.103.1	255.255.255.255	GW-2	Backup
201.202.203.1	255.255.255.255	GW-3	Regular
201.202.203.1	255.255.255.255	GW-2	Backup

The firewall within the farm specified for a destination network directly corresponds with the same firewall utilized for the source address, or gateway protecting that network.

Basically, the Destination Rules are created for each potential network that a VPN tunnel can be established to. Only one firewall should be created as "Regular" or active for a specific destination network. A backup is then designated. Next, the Source Rules is

created on the External FireProof. The Source Rule, which includes the IP Address of a remote gateway, is mapped to the same firewall as the protected network is mapped in the Destination Rules.

With this configuration, asymmetric routing of VPN bound traffic among multiple gateways within the cluster can be alleviated.

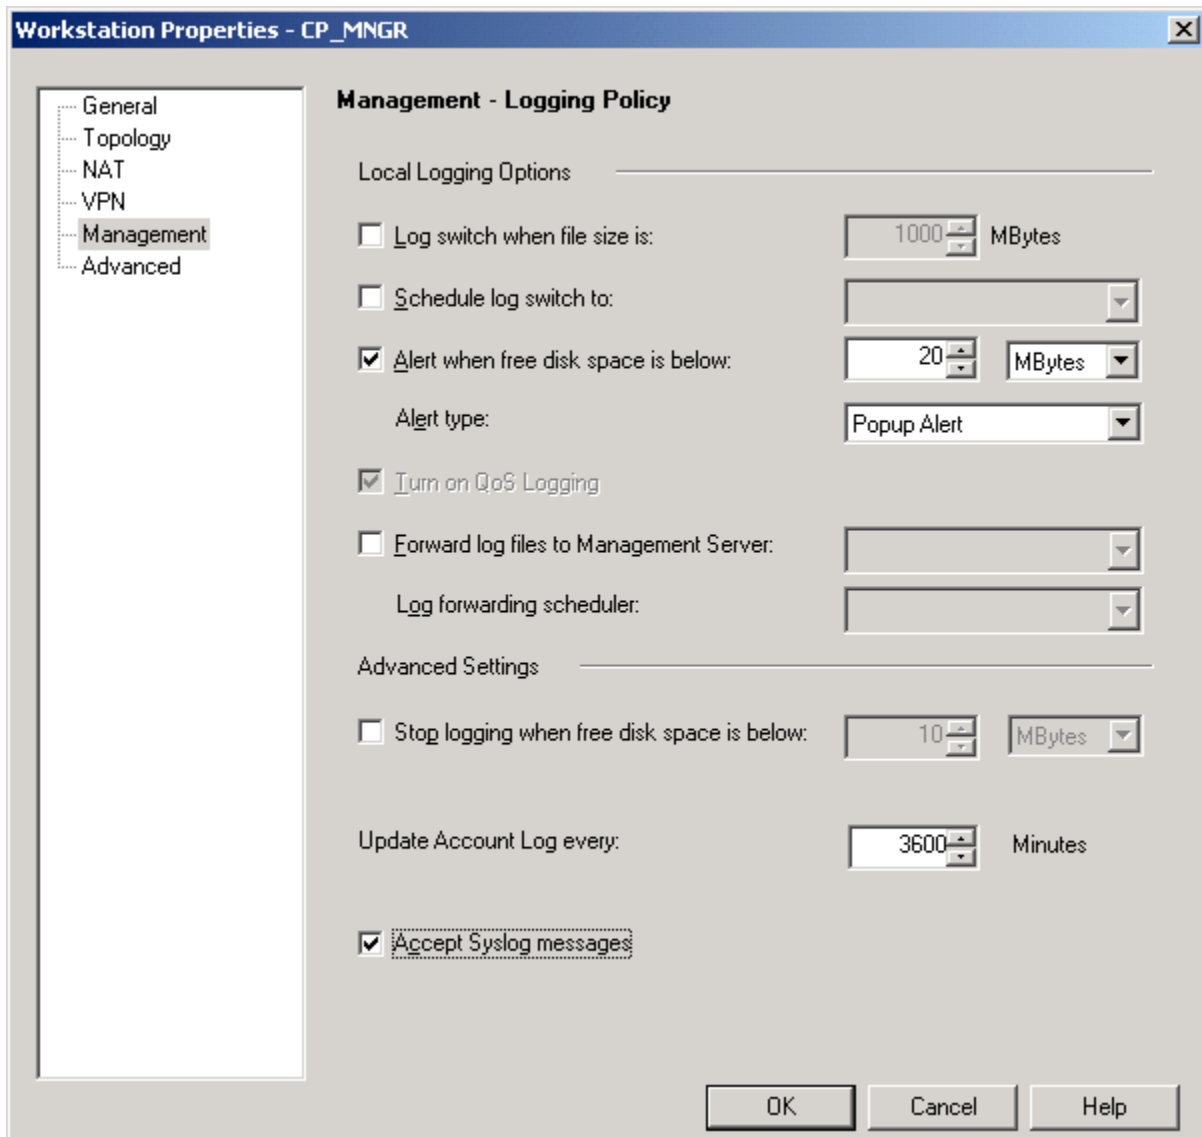
Syslog & ELA Logging

Syslog Setup

Check Point Next Generation supports the ability to allow third party devices to send logs and device status information directly to the management station via syslog. This section outlines the configuration of both the Check Point Management Server and FireProof device to enable this logging functionality.

Check Point Manager Setup

To enable the syslog daemon, you must configure the Management module to accept syslog messages. This is configured in the Workstation Properties of the Manager Server object as shown below.



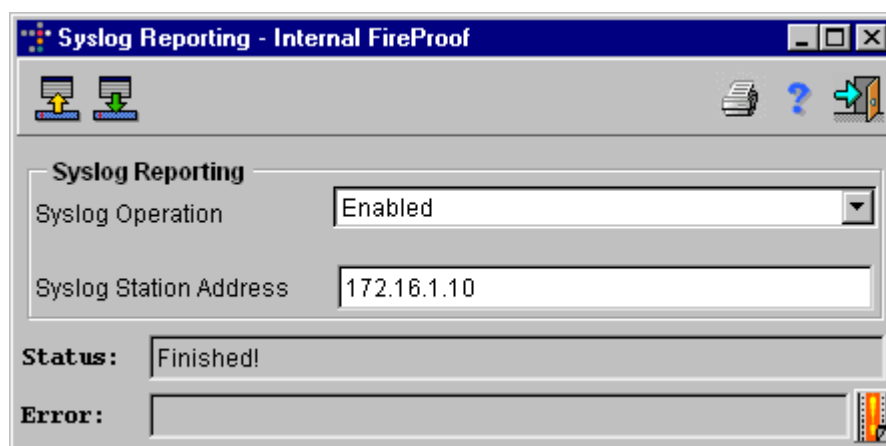
Workstation Properties – Management Window

From the Management screen of the Workstation Properties, check Accept Syslog messages. It is necessary to reboot the device that will be accepting Syslog messages after the policy has been updated.

NOTE If messages will be traversing the enforcement modules (i.e. External FireProof to internal Management Module), a special rule should be created in the Rule Base allowing syslog traffic from the configured FireProof devices to be forwarded to the Management Module.

FireProof Setup

The FireProof must be configured to log directly to the Check Point Management Server. This is configured within the Device → Syslog Reporting window.



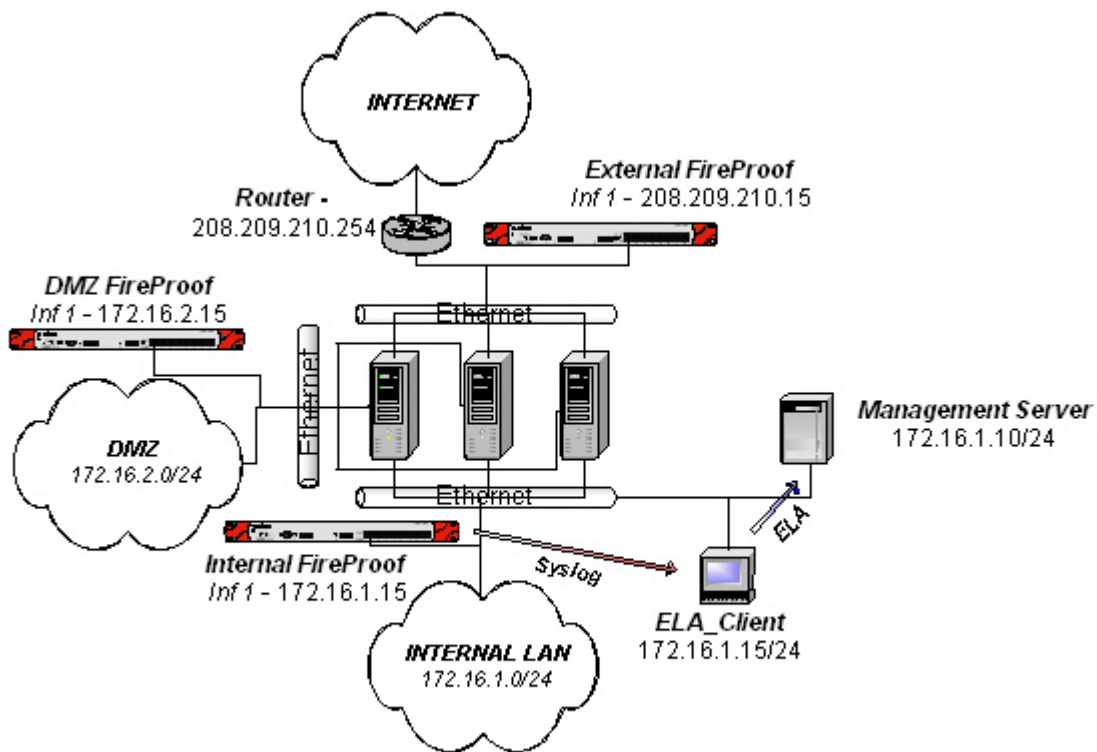
Syslog Reporting Window

From the Syslog Reporting window, simply enable Syslog Operation and specify the address of the Management Server. Once set to the FireProof, it will notify and log all trap information directly to the Management Server log.

ELA Setup

ELA is the Check Point standard mechanism for writing information to the log file. This type of logging is not directly supported by the FireProof because of its solid-state hardware architecture. However, the ability for FireProof log messages to be logged via ELA does exist through a “third-party/non-direct” logging setup.

In this scenario, an independent network host will be configured to run Radware's ELA_Client program. This program will allow all relevant FireProof devices to log messages to this host via syslog. The ELA_Client software in turn, formats the log information into ELA format and forwards the log information to the management server.



ELA_Client Implementation

Types of ELA Communication Between the ELA_Client and Management Server

There are 2 types of communication supported by Radware's ELA_Client interface:

1. Clear unauthenticated (meaning it does not verify who is sending the log) and it uses the port statement in the conf file.
2. Authenticated connection which uses the keys created by the opsec_putkey program on the ELA_Client and the fw -opsec putkey command on the Management Server to verify that the communication is coming from an authenticated source. This information is not encrypted and can be viewed on the wire, but is only allowed to be entered by an authenticated source. This connection uses the auth_port and auth_type auth_opsec statements in the conf file.

Setup

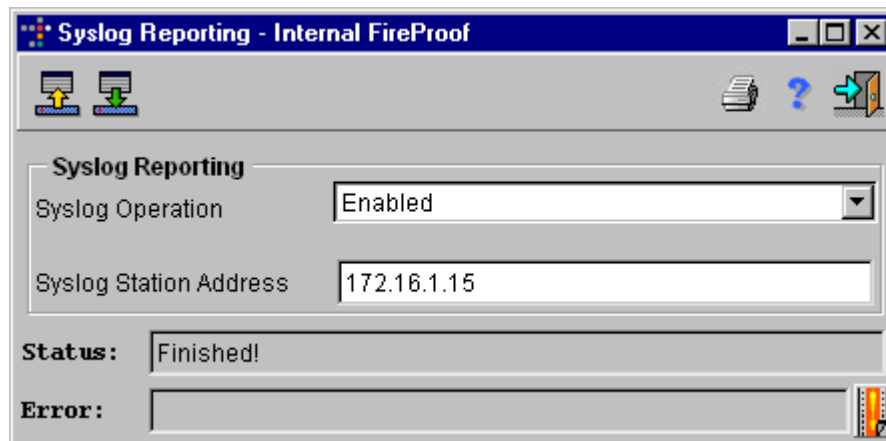
To implement ELA logging utilizing FireProof logs, the ELA_Client program must be obtained from Radware and installed on the designated ELA_Client host in the network.

There are three components that require configuration:

- FireProof(s) participating in ELA log solution
- ELA_Client
- Management Server

FireProof Configuration

The FireProof will be configured to send syslog messages to the ELA_Client host. For this example, the network host 172.16.1.15 has been designated. This host must be configured in the FireProof's Syslog Reporting window (Device → Syslog Reporting) shown below.



Syslog Reporting Window

From this window, enable Syslog Operation and set the Syslog Station Address as the ELA_Client IP Address.

ELA_Client Setup

Copy the following files into a directory.

- ela_client.exe* – Application Executable
- ela.conf* – Configuration File
- opsec_putkey.exe* – Used to set authentication keys.

NOTE The ELA_Client program currently only supports Windows NT 4.0 and Windows 2000 Operating Systems.

As outlined in the previous section, the ELA_Client can send ELA logs one of two ways; Clear or Authenticated. Each method requires a slightly different configuration, which will be covered in detail. It is important to decide which level of communication to implement and follow the exact setup rules for that method.

ELA Clear (unauthenticated) Communication Setup

In order to configure CLEAR connections on port 18187 perform the following.

1. Configure the ELA client to connect on port 18187 and give it the IP address of the NG mgmt server.

```
# Configuration file for ela_client
ela_server      ip          172.16.1.10
ela_server      port          18187
#ela_server     auth_port   18187
#ela_server     auth_type   auth_opsec
```

2. On the ELA_Client machine, start the ELA_Client program utilizing the following syntax:

```
ela_client -p <protocol, ex. TCP/UDP> -e <Port Syslog is listening, ex. 514> -i <interface ela_client is running on, ex. ela_client machine IP Address>
```

A command should look like this:

```
ela_client -p UDP -e 514 -i 172.16.1.15
```

Once running, the host will now be able to accept Syslog reporting messages from the FireProof and, utilizing the ela.conf settings, send the ELA formatted messages to the appropriate Management Server.

3. On the Management Station, issue a "cpstop" command.
4. Edit the fwopsec.conf file (\$fwdir\conf\fwopsec.conf) to force ELA Connections to listen in the clear (note by default ELA will not accept clear connections - you must edit the auth_port entry to port 0.) It will look like this when finished (note both ela_server lines are uncommented)

```
#
# lea_server  auth_port   18184
# lea_server  port        0
#
ela_server   auth_port   0
ela_server   port        18187
#
```

5. Issue "cpstart" on the NG management station.

NOTE If you wanted to have the `ela_server` listen to `auth_port` 18187 and clear connections something other than port 18187 -(for example port 9999) - the following entries would work:

```
#
ela_server      auth_port      18187
ela_server      port           9999
#
```

ELA Authenticated Communication Setup

In order to configure AUTH connections on port 18187 (for example standard `fwn1` connections in 4.1 created using `fw putkey - ela_server auth_type auth_opsec`)

1. Configure the `ela_client` to connect to `auth_port` 18187 in the `ela.conf` file.

```
# Configuration file for ela_client
ela_server      ip           172.16.1.10
#ela_server     port           18187
ela_server      auth_port    18187
ela_server      auth_type    auth_opsec
```

2. On the ELA_Client machine, start the ELA_Client program utilizing the following syntax:

```
ela_client -p <protocol, ex. TCP/UDP> -e <Port Syslog is listening, ex. 514> -i
<interface ela_client is running on, ex. ela_client machine IP Address>
```

A command should look like this:

```
ela_client -p UDP -e 514 -i 172.16.1.15
```

Once running, the host will now be able to accept Syslog reporting messages from the FireProof and, utilizing the `ela.conf` settings, send the ELA formatted messages to the appropriate Management Server.

3. Issue "cpstop" on the NG management station
4. Edit the `sic_policy.conf` file located in the CPSHARED conf directory (for example "c:\program files\checkpoint\cps\shared\5.0\conf\sic_policy.conf"). The entry to allow `fwn1` connections will be made by inserting a new rule such as the one below in the "Inbound Rules" section directly below the line that states "OPSEC configurations - place here..."

```
#
ANY; ANY; 18187; fwn1_opsec; fwn1, local_ipcheck
#
```

5. On the NG Management station perform an "fw putkey -opsec x.x.x.x" and enter the secret key - (this process is the same as 4.1 versions)

-
6. On the `ela_client` issue a `opsec_putkey y.y.y.y` and enter the same secret key (again same as previous versions)
 7. Issue "`cpstart`" on the NG management station

NOTE *There is no need to edit `fwopsec.conf` in this scenario - `sic_policy.conf` controls the authentication mechanism (i.e. `sslca`, `fwn1` etc)*
