

Rick Wanner  
Practical Project for GCFW  
For SANS New Orleans 2001  
Firewalls, Perimeter Protection and VPNs  
Version 1.5b

# Assignment 1 - Security Architecture

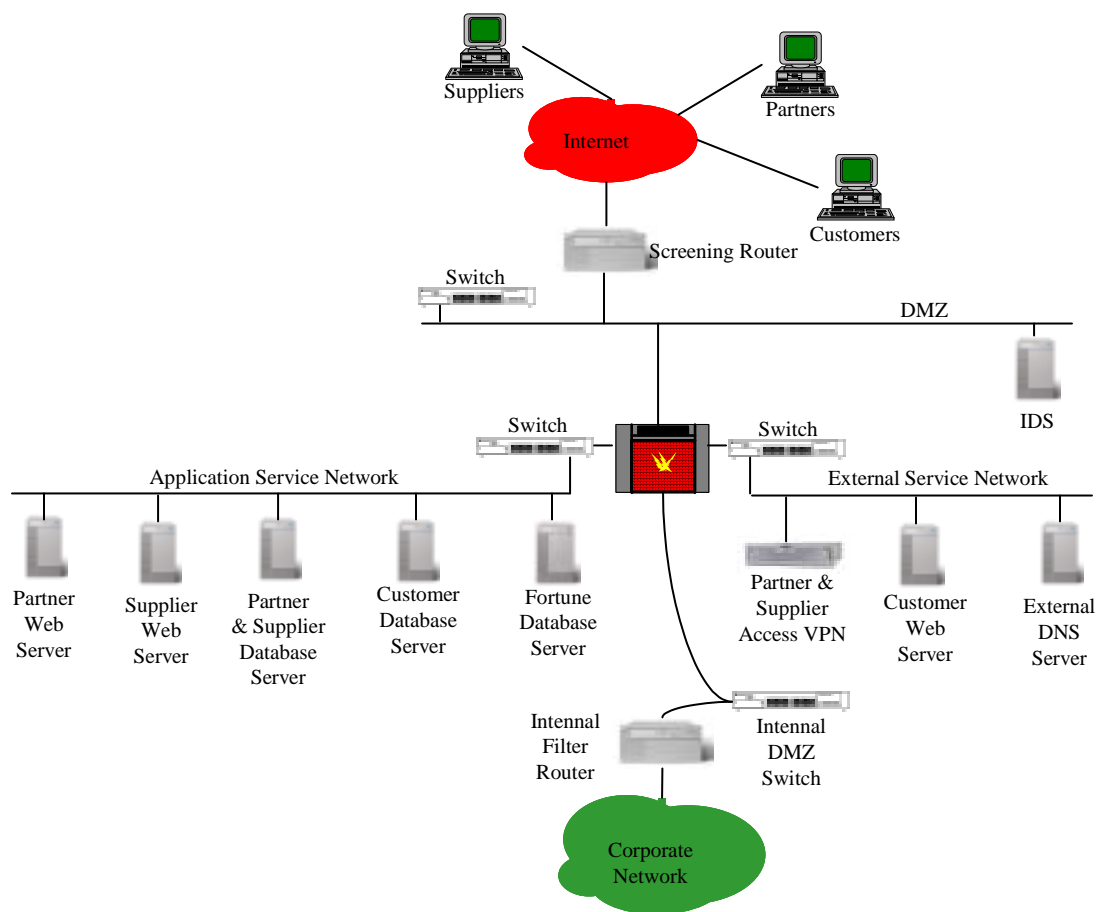
## Assignment 1 - Security Architecture (25 Points)

Define a security architecture for GIAC Enterprises, a growing Internet startup that expects to earn \$200 million per year in online sales of fortune cookie sayings, and which has just completed a merger/acquisition. Your architecture must specify filtering routers, firewalls, VPNs to partners, secure remote access, and internal firewalls. Be explicit about the brand and version of each perimeter defense component. Produce a diagram or set of diagrams with explanatory text that define how to use perimeter technologies to implement your security architecture.

You must consider and define access for:

- Customers (the companies that purchase bulk online fortunes);
- Suppliers (the authors of fortune cookie sayings that connect to supply fortunes);
- Partners (the international partners that translate and resell fortunes).

## GIAC Enterprises E-business Network Diagram



The following access is provided through this network implementation:

- Customer access via web server to purchase fortunes.
- Partner access via VPN. Once connected via VPN, partners may access an application web server to obtain fortunes for translation, and upload translated fortunes.
- Supplier access via VPN. Once connected via VPN, suppliers may access an application web server to upload fortunes.
- Employee access from the Corporate Network to perform ongoing maintenance and support on the infrastructure components.

This project only covers the E-business implementation for GIAC Enterprises. It describes the components required to do business with customers, partners, and suppliers, and the corresponding maintenance activities that must occur to support this infrastructure. It is assumed that employee related services such as Internet access, email, web access and employee remote access are provided through a different network implementation. The only employee access that will be considered will be the access requirements for supporting and maintaining the Ebusiness architecture.

This implementation was designed with the intent to compartmentalize the network into functions based on the nature of the traffic. The DMZ is used as a cushion from the Internet. Although the e-business managers fought to place web servers on the DMZ in order to reduce latency, for the moment the security analysts have won and none of the application services are being offered on the DMZ. This means that there are at least two lines of filtering between the Internet and any servers.

The external service network is where all of the Internet facing services will be offered. For example, the web server which customers access to buy fortunes is on this network.

The application service network provides the support services for the Internet facing services. For example this is where you will find the database servers that contain the fortunes and customer data used by the web servers

The internal DMZ (for lack of a better term) is used to keep unnecessary and unwanted noise from our corporate network from getting to our external networks. It has been our experience that enough noise is generated on our corporate network (especially from netbios traffic), that it can have an impact on the performance of the firewall. The internal filter router is used to filter this noise.

The following sections provide more on each hardware component used in our e-business network.

## ***Infrastructure Components***

### **Border Router**

First of all it must be pointed out that we are a Nortel Networks shop, most of our networking hardware will reflect this. The border router is a Nortel Networks BLN router running BayRS version 14.11. It is our first line of defense against unauthorized access. It is configured with a “deny all except what is explicitly permitted” policy. This policy not only creates a strong network security model, but it also permits control over what servers are placed on our external networks. Since any new Internet visible services require a change to the border router the security analysts are aware of all servers that are placed on these networks, and what services these servers are providing. This control can allow the security analysts to ensure that all servers are hardened and scanned before they are put into production.

### **Switches**

There are switches on each of our external networks. All of them are Nortel Networks BayStack 450s running software version 3.1. They are shown because they are an often-overlooked part of the security of the networks. A separate switch is used on each network, and virtual LANs are not permitted on any of these switches. This mitigates the consequences of attacks against these switches by tools such as dnsiff<sup>10</sup>. Also, all ports on the switches that are not in use are disabled. This is another check that servers cannot be added to the network without the knowledge of the security analysts.

### **Intrusion Detection Server**

Intrusion detection is provided with ISS RealSecure Network Sensor version 5.0 running on a Sun Ultra 10 with Solaris 2.6. The intrusion detection server is attached to the network via a spanning port on the switch. Using the switch spanning port seems to be a relatively robust solution to the problems involved in

doing intrusion detection in a switched network. However, this configuration has not been tested under high throughputs. I have some concern of the ability of the spanning port to keep up once network traffic increases.

## Firewall

As I said earlier, we are a Nortel Networks shop. The firewall is a Checkpoint Firewall-1 version 4.0 integrated firewall. It is integrated in a Nortel Networks Contivity Extranet Switch 4500 running Contivity OS version 2.61. This is a cheap way to provide a firewall on existing equipment. I do have a couple of concerns about this however:

- Because this is an integrated firewall module, we do not have the full capabilities of Firewall-1. Although the basic stateful packet filtering functionality of Firewall-1 is present, several of the advanced features are missing. Some of these features can be used to enhance the security of the firewall.
- This is not the latest version of Firewall-1 available. We are losing out on some of the features, improvements, and bug fixes in subsequent versions.
- The integrated module does not permit patching. This version of the integrated module is patched along with the Contivity OS when new releases. This results in delays in receiving firewall patches until the next version of the Contivity OS is available. Often these releases don't come until several weeks after the patch has been released by Checkpoint.

The default policy on the firewall is also “deny all except what is explicitly permitted”. We count on the screening and filter routers to eliminate spoofed packets and noise. Because of this the firewall runs at a substantially lower CPU occupancy than they did before the filter and screening router filters were implemented.

## Partner and Supplier Access VPN Server

The VPN Server used for supplier and partner access is a Nortel Networks Contivity Extranet Switch 2600 running Contivity OS version 2.61. Suppliers and partners are provided a copy of Contivity Extranet client version 2.62.

It would have been nice to separate suppliers and partners onto different VPN servers, unfortunately due to budget constraints that is not possible. Instead different access groups are defined for each partner and supplier company. By using this group information, each individual supplier or partner company can be allocated IP addresses out of a unique DHCP address pool. When a supplier or partner connects using the IPSec client they receive an IP out of their company's DHCP pool. By placing all the supplier pools in a contiguous range, and the partner pools in a different contiguous range, we have the ability to restrict traffic through the firewall based on the IP the user received when they connected. This configuration also gives us the flexibility to define firewall rules based on an individual company if we wish to do that at some point in the future.

## Filter Router

The internal filter router is a Nortel Networks BLN router running BayRS version 14.11. Its primary function is to eliminate the corporate network noise that would otherwise get to the firewall and be dropped. This substantially reduces the effort required by the firewall, thus preserving CPU cycles for serving the e-business implementation. The secondary function is to protect the e-business environment from unauthorized access from the corporate network. The filter router is configured with a “deny all except what is explicitly permitted” policy. Filters on this router are not as strict as on the screening router, but are limited to only the services that are required to perform diagnostics, and provide ongoing maintenance and support for the e-business infrastructure.

## **External DNS Server**

The external DNS server is an HP C3600 workstation running HP-UX 10.20 and Bind version 8.2.3. the external DNS server provides DNS information for the GIAC enterprises e-business infrastructure. This machine does not accept zone transfers.

## **Access Scenarios**

### **Customer Access**

Customer access to purchase fortunes is via http and https to the web server on the external service network. The web server contains no useful data, and must query the database servers on the application service network for fortunes and customer data.

### **Partner Access**

Partner access is provided via IPSec VPN to the partner/supplier VPN server on the external service network. As well as individual logins using userid and password authenticated by radius, group userid and passwords are also used. The group logins permit the VPN server to assign each partner to a unique group and to assign them an IP address out of a pool explicitly defined for their company. This IP is used to limit access through the firewall to only the partner application web server on the application service network. This partner application web server contains web applications that permit the partners to obtain fortunes for translation and to upload translated fortunes and perform other functions related to their business with GIAC Enterprises. The application web server contains no useful data, and must query the database servers, also on the application service network, for fortunes and partner data.

### **Supplier Access**

Supplier access is provided via IPSec VPN to the partner/supplier VPN server on the external service network. As well as individual logins using userid and password authenticated by radius, group userid and passwords are also used. The group logins permit the VPN server to assign each supplier to a unique group and to assign them an IP address out of a pool explicitly defined for their company. This IP is used to limit access through the firewall to only the supplier application web server on the application service network. This supplier application web server contains web applications that permit the suppliers to upload fortunes and perform other functions related to their business with GIAC Enterprises. The application web server contains no useful data, and must query the database servers, also on the application service network, for fortunes and supplier data.

### **Employee Access for Maintenance and Support**

Employee access for support of the e-business infrastructure is only available for troubleshooting and ongoing maintenance and support of the components. This means that web access, and IPSec access are permitted from the corporate network to the e-business networks for testing and troubleshooting of components. Also ssh access is permitted for access to all of the e-business components. Because some of the components don't support ssh access, such as the switches, telnet access is also permitted. All other access is dropped at the filter router and the firewall. Other services that may be required for special maintenance, such as tftp for upgrading the routers and switches, must be approved by the security analysts and are only permitted for a limited period of time.

### **Assignment 2 - Security Policy (25 Points)**

Based on the security architecture that you defined in Assignment 1, provide a security policy for AT LEAST the following three components:

- Border Router
- Primary Firewall
- VPN

You may also wish to include one or more internal firewalls used to implement defense in depth or to separate business functions.

By 'security policy' we mean the specific ACLs, firewall ruleset, IPSec policy, etc. (as appropriate) for the specific component used in your architecture. For each component, be sure to consider internal business operations, customers, suppliers and partners. Keep in mind you are an E-Business with customers, suppliers, and partners - you MAY NOT simply block everything!

(Special note VPNs: since IPSec VPNs are still a bit flaky when it comes to implementation, that component will be graded more loosely than the border router and primary firewall. However, be sure to define whether split-horizon is implemented, key exchange parameters, the choice of AH or ESP and why. PPP-based VPNs are also fully acceptable as long as they are well defined.)

For each security policy, write a tutorial on how to implement each ACL, rule, or policy measure on your specific component. Please use screen shots, network traffic traces, firewall log information, and/or URLs to find further information as appropriate. Be certain to include the following:

1. The service or protocol addressed by the ACL or rule, and the reason these services might be considered a vulnerability.
2. Any relevant information about the behavior of the service or protocol on the network.
3. The syntax of the ACL, filter, rule, etc.
4. A description of each of the parts of the filter.
5. An explanation of how to apply the filter.
6. If the filter is order-dependent, list any rules that should precede and/or follow this filter, and why this order is important. (Note: instead of explaining order dependencies for each individual rule, you may wish to create a separate section of your practical that describes the order in which ALL of the rules should be applied, and why.)
7. Explain how to test the ACL/filter/rule.

Be certain to point out any tips, tricks, or "gotchas".

## **General Security Policy**

Because of the important business provided by this network, the following general policies have been applied:

- All Internet facing services must reside on the external service network. No traffic is permitted from the Internet directly to the application service network, or the corporate network.
- All servers must be scanned and hardened before being placed in service. Monthly scans are performed to ensure the servers remain properly hardened. Good documents on hardening Solaris and NT servers can be found at <http://www.enteract.com/~lspitz/armoring.html><sup>13</sup> and <http://www.enteract.com/~lspitz/nt.html><sup>14</sup>.
- Where possible superuser accounts have been eliminated, or renamed, or their logins disabled. Access to superuser privileges is through sudo-like functionality or external security manager only.
- Access to all servers is with ssh or IPSec where supported. Where ssh or IPSec access is not possible, telnet is restricted to designated nodes where supported.
- SNMP is used for network monitoring of servers on these networks. On all network devices SNMP read and write community strings have been set to non-default values.<sup>5</sup> Some devices support management and/or configuration through SNMP. On all devices where another method of management and/or configuration is available SNMP write has been disabled if possible.
- All devices that support NTP are synchronized to an Internet NTP server. This ensures that all logs are synchronized for log correlation and possible investigative purposes.

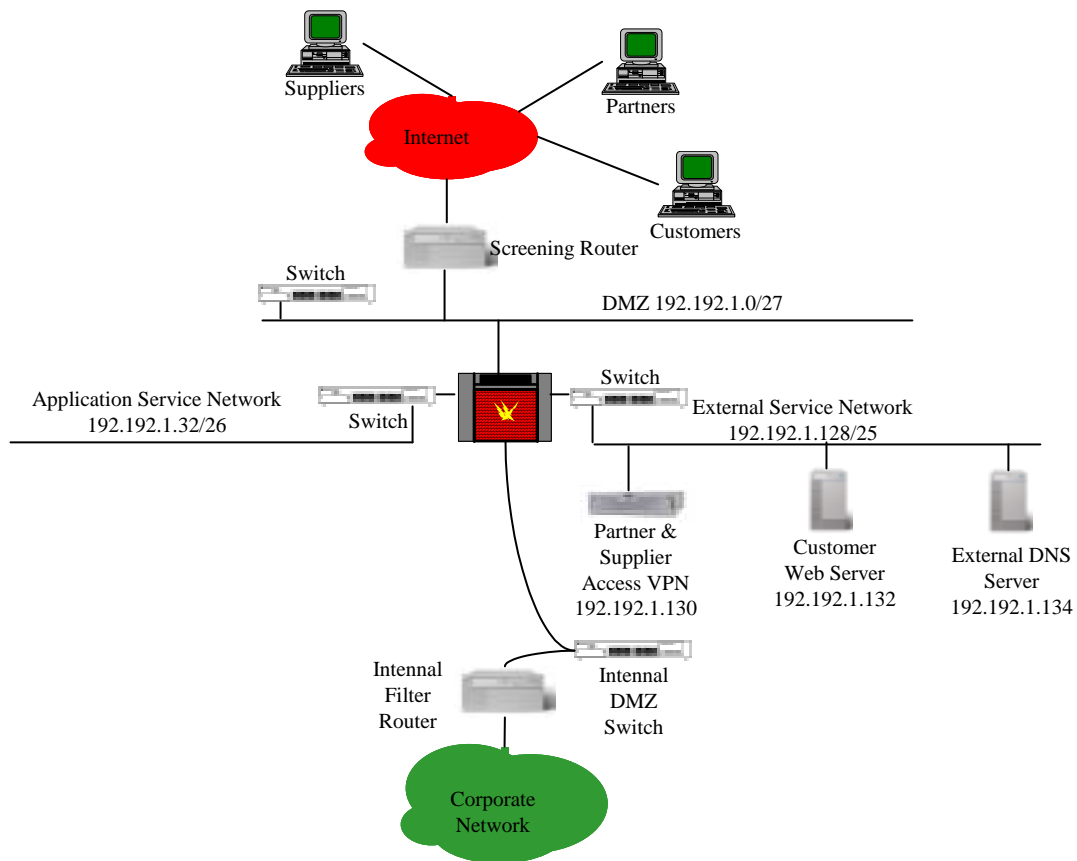
- All devices that support syslog are logged to a syslog server on the application service net. Logs are written off to disk and then stored on a write once media
- Where possible the login banner on all devices has been changed to the following message<sup>5</sup>:

“Access to this network and the information on it are lawfully available only for approved purposes by employees of GIAC Enterprises and other users authorized by GIAC Enterprises. If you are not an employee of GIAC Enterprises or an authorized user, **do not attempt to log on**. Other than where prohibited by law and subject to legal requirements, GIAC Enterprises reserves the right to review any information in any form on this network at any time.”

- Physical security of the components is important in order to limit insider and physical attacks on the e-b-business infrastructure. All components with an Internet presence are locked in a room with badge reader access. Access to the room is limited to designated support engineers only.
- All traffic passing the firewall is logged. All normal traffic is logged “short” thus containing only connection information. All denied traffic is logged “long”, to permit detailed analysis of the traffic.
- All ports on the external network switches that are not in use, are disabled. This prevents the addition of foreign devices to the networks.

## External Screening Router

The following network diagram shows the IP addresses applied to all of the external networks, and the servers on the external service network providing Internet facing services. Other servers have been removed from the diagram. Although addresses are not required for much of this practical, they are required to add clarity to the router filters and to the VPN configuration.



As you can see GIAC incorporated owns an entire class C subnet, 192.192.192.0/24. The first 32 addresses, from 0 to 31, are allocated to the DMZ. The next 64 addresses, from 32 to 95, are allocated to

the Application Service Network. The range from 128 to 255 is allocated to the External Service Network. The first 32 addresses of that range, from 128 to 159, are reserved for servers on the External Service Network. The remaining addresses from 160 to 255 are reserved for the supplier and partner VPN IP pools, with 160 to 191 being allocated to suppliers, and 192 to 255 being allocated to partners. Each company will reserve 4 IP addresses from the pool, thus we can support up to 8 suppliers and 16 partners in this configuration. This configuration does not leave a lot of room for future expansion, and will probably result in renumbering the External Service Network at some point in the future.

## ***Preliminary Setup***

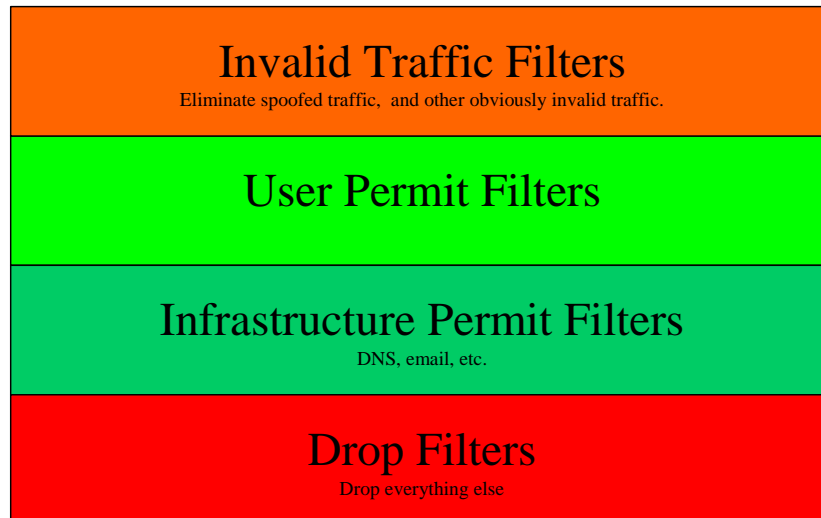
### **Preventing Smurf Attacks**

Starting in BayRS 13.01 ICMP directed broadcasts are disabled by default. We should be ok since we are using 14.11. If you are the paranoid type and want to be sure enter the following at a command prompt:

```
[1:1]$bcc  
bcc> config  
hostname# ip  
ip# directed-bcast disabled  
ip# exit
```

### ***Router Filters***

The Nortel Networks BLN is a typical packet filtering router. Filters are considered according to their precedence, and the first filter that matches is applied. The precedence of the filters is important, for efficiency reasons, security reasons, and to preserve the administrator's intent. A typical ordering scheme appears below:



The spoofed traffic filters are placed at the top because this traffic is never valid and should be disposed immediately. This traffic is either as a result of nefarious purposes of the originator, or more often than not is just the usual invalid noise that happens on any network. In our case we count on the external screening router to do this filtering, thus reducing the overhead on the firewall.

Following the spoofed traffic filters are the user permit filters. These should be close to the top because logically they should be the most often hit filters. In order to best maximize the performance of the router, by minimizing the amount of processing required, it is a good idea to optimize the user filters so that the traffic that needs to execute the fastest gets precedence within the user permit filter block. In our case we will give precedence to customer traffic, then partner traffic, then supplier traffic.

Following the user permit filters are the infrastructure permit filters. This is where filters to permit DNS, email, NTP and other infrastructure services would be permitted. These services are generally not latency critical, and can usually take lower precedence than user permit filters. The one potential pitfall here is DNS. If the volume of DNS queries is high they should probably be optimized higher in the filters in order to improve the efficiency of the router.

The remaining block contains the drop filter. If network traffic makes it this far in the rule set it will be denied.

A word about precedence. Precedence of filters in Nortel Networks routers is a integer value with 1 being the highest, and therefore the first executed, and the higher the number the lower the precedence, and the later it will be executed.

## Blocking Spoofed Addresses

### Methodology:

This filter is applied inbound to prevent spoofed packets. This filter drops all packets with a source of any of our e-business network addresses, RFC 1918 private addresses<sup>15</sup>, or IANA Reserved addresses<sup>9</sup>. It also drops any traffic originating from the loopback address.

### Precedence:

This filter should have a precedence of one (the highest precedence) so it is always applied first.

### Gotchas:

If any of the IANA Reserved addresses become unreserved, in other words if they become acceptable to use, this filter may need to change.

If the e-business addresses are renumbered, or a new subnet added this filter will need to change.

### Breakdown:

The filter works by matching the source network address. The first three source-network entries are the RFC 1918 private addresses, the next several are the IANA-RESERVED addresses, the second last entry is the loopback address, and the final one is the address of our e-business networks. Any packets matching any of the source-network entries are dropped.

```
traffic-filter filter-name Drop_Spoofed_Src_Packets
```

```
precedence 1
```

```
match
```

```
source-network range 10.0.0.0-10.255.255.255 (RFC 1918 Class A Private)
```

```
back
```

```
source-network range 172.16.0.0-172.31.255.255 (RFC 1918 Class B Private)
```

```
back
```

```
source-network range 192.168.0.0-192.168.255.255 (RFC 1918 Class C Private)
```

```
back
```

```
source-network range 0.0.0.1-2.255.255.255 (IANA -RESERVED)
```

back

source-network range 5.0.0.0-5.255.255.255 (IANA- RESERVED)

back

source-network range 7.0.0.0-7.255.255.255 (IANA- RESERVED)

back

source-network range 23.0.0.0-23.255.255.255 (IANA -RESERVED)

back

source-network range 27.0.0.0-27.255.255.255 (IANA -RESERVED)

back

source-network range 31.0.0.0-31.255.255.255 (IANA -RESERVED)

back

source-network range 36.0.0.0-37.255.255.255 (IANA -RESERVED)

back

source-network range 41.0.0.0-42.255.255.255 (IANA -RESERVED)

back

source-network range 58.0.0.0-60.255.255.255 (IANA -RESERVED)

back

source-network range 67.0.0.0-126.255.255.255 (IANA -RESERVED)

back

source-network range 197.0.0.0-197.255.255.255 (IANA -RESERVED)

back

source-network range 201.0.0.0-201.255.255.255 (IANA -RESERVED)

back

source-network range 219.0.0.0-223.255.255.255 (IANA -RESERVED)

back

source-network range 240.255.255.0-255.255.255.255 (IANA-RESERVED)

back

source-network range 127.0.0.0-127.255.255.255 (Loopback Address)

back

source-network range 192.192.1.0-192.192.1.255 (Corporate Network)

back

back

actions

action drop

back

back

## Permit HTTP in to the Web Server

### Methodology:

This filter is applied inbound to permit HTTP to the web server. This filter could have been combined with the HTTPS filter, but it is possible you may not want to provide both HTTP and HTTPS to a web server, so they were kept separate.

### Precedence:

This filter has a precedence of five. There is no security reason this filter must have a high precedence. It is placed where it is because due to latency and efficiency considerations we are giving precedence to customer filters.

### Gotchas:

If any new web servers are added they will need to be added into this filter.

### Breakdown:

The filter works by matching packets destined for port 80 (HTTP) that are protocol 6 (TCP), from the Internet (0.0.0.0-255.255.255.255), destined for the web server (192.192.1.132).

traffic-filter filter-name Accept\_HTTP\_IN

precedence 5

match

dest-tcp-ports {80}

protocol 6

source-network range 0.0.0.0-255.255.255.255 (INTERNET)

back

destination-network range 192.192.1.132 (WEB SERVER)

back

back

actions

back

back

## Permit HTTPS to the Web Server

### Methodology:

This filter is applied inbound to permit HTTPS to the web server. This filter could have been combined with the HTTP filter, but it is possible you may not want to provide both HTTP and HTTPS to a web server, so they were kept separate.

### Precedence:

This filter has a precedence of six. This precedence was chosen to keep this filter near the HTTP filter. There is no security reason this filters must have a high precedence. It is placed where it is because due to latency and efficiency considerations we are giving precedence to customer filters.

### Gotchas:

If any new web servers are added they will need to be added into this filter.

### Breakdown:

The filter works by matching packets destined for port 443 (HTTPS) that are protocol 6 (TCP), from the Internet (0.0.0.0-255.255.255.255), destined for the web server (192.192.1.132). Matching traffic is permitted

```
traffic-filter filter-name Accept_HTTPS_IN
```

```
precedence 6
```

```
match
```

```
dest-tcp-ports {443}
```

```
protocol 6
```

```
source-network range 0.0.0.0-255.255.255.255
```

```
back
```

```
destination-network range 192.192.1.132
```

```
back
```

```
back
```

```
actions
```

```
back
```

```
back
```

## Permit IPSec to Customer/Supplier VPN Server

### Methodology:

This filter is applied inbound to permit IPSec to the VPN server. There is a second filter below for the ESP.

### Precedence:

This filter has a precedence of 15. There is no security reason this filters must have a high precedence. It is placed where it is because due to latency and efficiency considerations we are giving highest precedence to customer traffic, than to partner and supplier filters.

### Gotchas:

If any new VPN servers are added they will need to be added into this filter.

### Breakdown:

The filter works by matching packets destined for that are protocol 50 (IPSec), from the Internet (0.0.0.0-255.255.255.255), destined for the VPN server (192.192.1.130). Matching traffic is permitted.

```
traffic-filter filter-name Permit_Ipsec_IN

precedence 15

match

protocol 50

source-network range 0.0.0.0-255.255.255.255

back

destination-network range 192.192.1.130

back

back

actions

back

back
```

## Permit ISAKMP

### Methodology:

This filter is applied inbound to permit IPSec to the VPN server.

### Precedence:

This filter has a precedence of 16. There is no security reason this filter must have a medium precedence. It is placed where it is because due to latency and efficiency considerations we are giving highest precedence to customer traffic, than to partner and supplier traffic.

### Gotchas:

If any new VPN servers are added they will need to be added into this filter.

### Breakdown:

The filter works by matching packets destined for that are protocol 17 (UDP), and source and destination ports 500. Matching traffic is permitted.

```
traffic-filter filter-name Permit_Isakmp
```

precedence 16  
match  
src-udp-ports 500  
dest-udp-ports 500  
protocol 17  
back  
actions  
back  
*back*

## **Permit DNS to DNS Server**

### **Methodology:**

This filter is applied inbound to permit domain requests to the DNS server.

### **Precedence:**

This filter has a precedence of 20. There is no security reason this filters must have a low precedence. It is placed where it is because due to latency and efficiency considerations we are giving highest precedence to customer traffic, than to partner and supplier traffic, then to infrastructure traffic.

### **Gotchas:**

If any new DNS servers are added they will need to be added into this filter.

### **Breakdown:**

The filter works by matching packets that are destined for port 53 (domain), match protocol 17 (UDP), from the Internet (0.0.0.0-255.255.255.255), destined for the DNS server (192.192.1.134). Matching traffic is permitted.

traffic-filter filter-name Accept\_DNS

precedence 20  
match  
dest-udp-ports 53  
protocol 17  
source-network range 0.0.0.0-255.255.255.255  
back  
destination-network range 192.192.1.134  
back  
back

actions

back

back

## Permit NTP OUT

### Methodology:

This filter is applied outbound to permit Network Time Protocol to the Internet from the e-business networks.

### Precedence:

This filter has a precedence of 22. There is no security reason this filter must have a low precedence. It is placed where it is because due to latency and efficiency considerations we are giving highest precedence to customer traffic, then to partner and supplier traffic, then to infrastructure traffic.

### Gotchas:

If any new e-business networks are added they will need to be added into this filter.

### Breakdown:

The filter works by matching packets that are destined for port 123 (domain), match protocol 17 (UDP), from the e-business networks (192.192.1.0-192.192.1.255) destined for the Internet (0.0.0.0-255.255.255.255). Matching traffic is permitted.

```
traffic-filter filter-name Permit_NTP
```

```
precedence 22
```

```
match
```

```
dest-udp-ports 123
```

```
protocol 17
```

```
source-network range 192.192.1.0-192.192.1.255
```

```
back
```

```
destination-network range 0.0.0.0-255.255.255.255
```

```
back
```

```
back
```

```
actions
```

```
back
```

```
back
```

## Drop\_Protocol & Drop\_All

### Methodology:

This is the deny all filter. If packet does not match any of the above filter rules, these filters drop the packet. This filter has to be implemented with the lowest precedence to ensure it executes last. Any filters with a lower precedence will never be executed.

**Precedence:**

This filter has a precedence of 31. This filter must have the lowest precedence so it will always be the last one executed.

**Gotchas:**

None

**Breakdown:**

The filter works by matching packets for any protocol (1-255), any source network range (0.0.0.0-255.255.255.255) and any destination network range (0.0.0.0-255.255.255.255). Any packet that gets this far will match this filter. All matching traffic (any traffic) is dropped.

```
traffic-filter filter-name Drop_All_Remaining
```

```
precedence 31
```

```
match
```

```
protocol 1-255
```

```
source-network range 0.0.0.0-255.255.255.255
```

```
back
```

```
destination-network range 0.0.0.0-255.255.255.255
```

```
back
```

```
back
```

```
actions
```

```
action drop
```

```
back
```

```
back
```

A couple of excellent treatises on securing routers (if rather Cisco specific) are available on the SANS website:

<http://www.sans.org/infosecFAQ/firewall/router.htm><sup>3</sup> and

[http://www.sans.org/infosecFAQ/firewall/blocking\\_cisco.htm](http://www.sans.org/infosecFAQ/firewall/blocking_cisco.htm)<sup>4</sup>

## ***External Firewall***

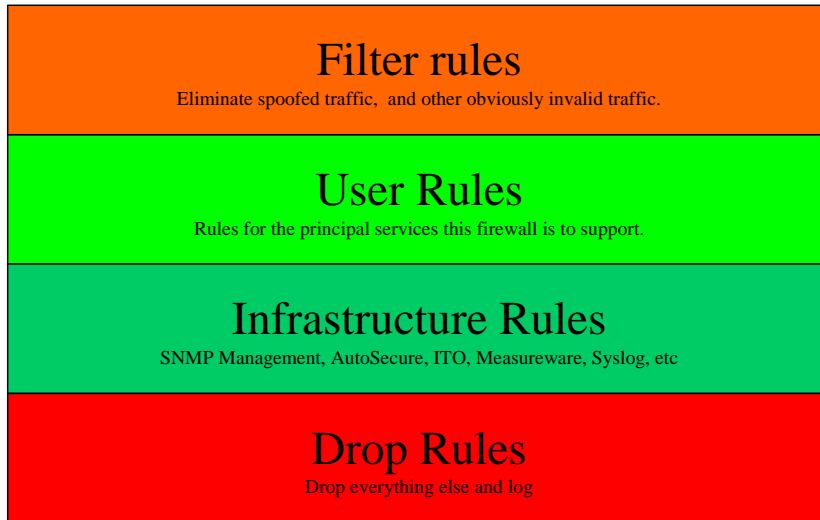
Firewall-1 is a stateful inspection firewall architecture. With Stateful Inspection, packets are intercepted at the network layer for best performance (the same as packet filters such as routers) and filtering decisions are based on context information determined from the state of the connection.

Firewall-1 rules are interpreted in a first match manner. This means that the rules are interpreted from the top down, and the first rule that matches the connection is used. Because of this it is important to pay attention to a couple of things:

- Order of rules is extremely important. Because Firewall-1 uses a first match algorithm it is important for efficiency and latency reasons that the rules that will match most frequently be placed near the beginning of the rule set when possible.

- Because the firewall is first match, in general rules should be written with the most specific near the beginning of the rule set and the least specific near the end. Improper ordering of rules can cause serious security problems. Placing a less specific rule too high in the rule set may result in a security hole caused by the less permissive rule being matched before the more specific rule.

The order in which the firewall rules appear is important, for efficiency reasons, security reasons, and to preserve the administrator's intent. A typical ordering scheme appears below:



The filters rules are placed at the top because this traffic is never valid, even if it would be allowed from a valid address and can be disposed of with little overhead. This traffic is either as a result of nefarious purposes of the originator, or more often than not is just the usual invalid noise that always seems to happen on any network. In our case we count on the screening border router to do all of this filtering for us, so we will not be doing it on the firewall, thus reducing the overhead on the firewall.

Following the filter rules are the user rules. These are close to the top because logically they should be the most often hit rules. In order to best maximize the performance of the firewall, by minimizing the amount of processing required, it is a good idea to optimize the user rules so that the traffic which needs to execute the fastest gets precedence within the user rules block. In our case we will give precedence to customer traffic, then partner traffic, then supplier traffic, and lastly maintenance traffic.

Next are connections required for management or administrative purposes. This is the connections which are required to ensure your network is healthy and happy. This would include connections for device monitoring, and NTP traffic to the Internet to ensure the clocks are all synchronized.

The remaining block contains the drop rules. If network traffic makes it this far in the rule set it will be denied. The traffic that remains generally falls into one of three categories:

- On every network there is a certain amount of noise, such as router broadcasts, which although it is annoying, is not a sign of suspicious activity. This traffic can just be dropped without logging. In our case we use the screening border router and the internal filter router to eliminate this traffic, so we don't have to write rules for this traffic.
- The second type of traffic is traffic that the security analysts consider interesting. This is traffic that shows the signature of the latest exploit, or is interesting for other reasons. This traffic is discarded and an alert generated. Again because we use the external screening

router and the internal filter router to eliminate this traffic we don't usually have rules for this traffic in the firewall.

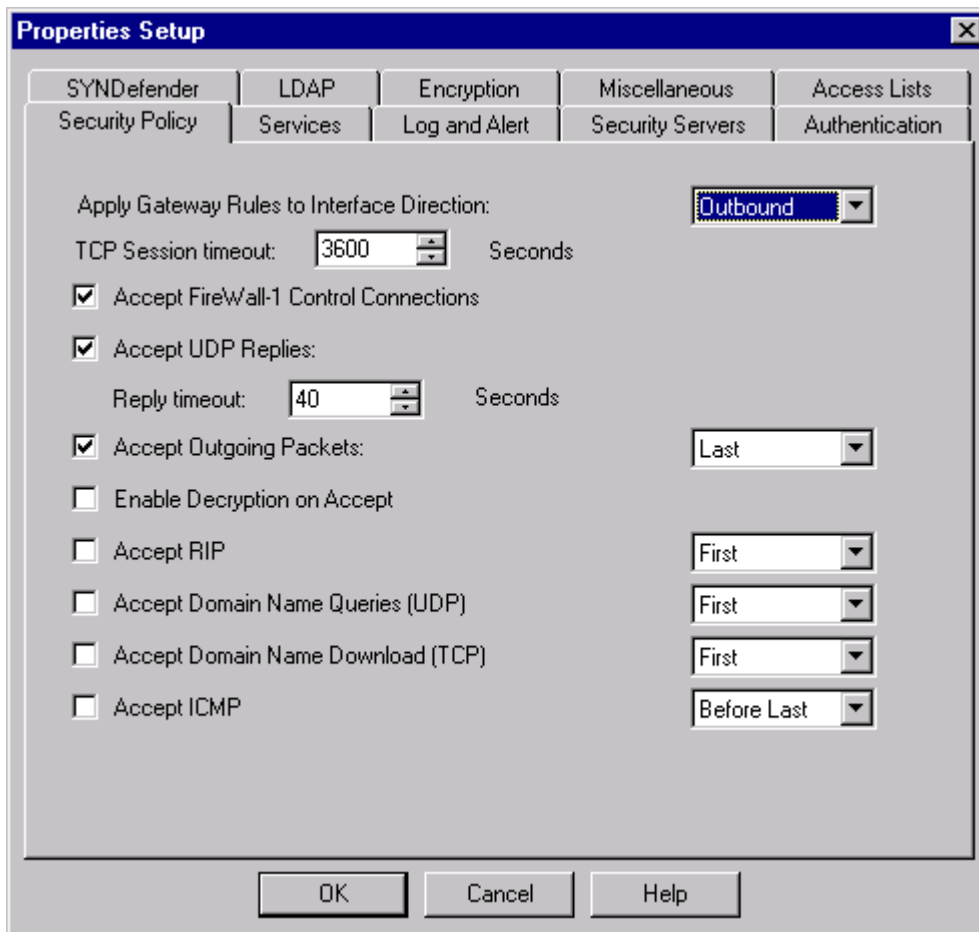
- The remaining traffic is everything left. We drop this and log it using Firewall-1's "long" tracking feature. Although this traffic is usually harmless, it is useful to analyze because it usually points out misconfigured or improperly hardened systems that are generating noise. Once in a while this traffic may show you an attack against your network.

## Firewall Rules

The complete firewall policy is in an Appendix. This section will describe them on a rule-by-rule basis.

## General Firewall-1 Setup

Default Firewall-1 comes with some services enabled by default. In order to avoid administrative confusion, we want to see as much as possible in explicit rules, not in Firewall-1 implied rules. In order to do this we need to ensure that our configuration appears as below:



The following options should be enabled:

- TCP Session Timeout - the time period after which a idle TCP session will be timed out, and terminated, by the firewall software. 3600 seconds is the default value.

- Accept FireWall-1 Control Connections – permits the Firewall-1 management station to send and install firewall policies. You can disable this option if the firewall and management station are on the same server. In this case they are separate.
- Accept UDP Replies – permits two-way UDP communication. This permits UDP traffic, to be somewhat stateful. In a nutshell, if a UDP packet goes out the firewall will wait for a response to that request. This is required to permit DNS queries to succeed.
- Reply Timeout -the amount of time the firewall will wait for a UDP reply before timing out, and discontinuing the communication. Forty seconds is the default value.
- Accept Outgoing Packets - accept all outgoing packets from the firewall itself. This permits connection debugging to be done directly from the firewall.

All other options should be disabled. We will enable them through explicit firewall rules, if required.

## Customer Access Rules

Number	Source	Destination	Service	Action	Track
1	Internet	Customer Web Server	HTTP HTTPS	Accept	short

This rule is to permit Customers to access the customer web server on the external service network. All access is via http and https. The breakdown of the rules is as follows:

- **Number:** This is rule 1 in the policy set. There is no security reason this rule must be first. This rule is first because due to latency and efficiency considerations we are giving precedence to customer rules.
- **Source:** The source is Internet because we don't know where our customer will come from on the Internet. If we had a limited set of customers we could restrict this to a set of customer subnets or servers.
- **Destination:** This is the customer web server. We restrict the access to only the machine that is required to reduce the possibility of attacks on other machines which may be running http or https services for maintenance, may not have been properly hardened, or may have had service opened up by accident through maintenance actions such as patching.
- **Service:** All customer applications are web based. Therefore HTTP and HTTPS are the only services required. Web servers are popular targets of crackers, so we will pay careful attention to any alerts from the IDS for this machine.
- **Action:** We are accepting this traffic.
- **Track:** We log this short because it is considered normal traffic. This provides enough information to do some analysis on the traffic if required.

Number	Source	Destination	Service	Action	Track
2	Customer Web Server	Customer Database Server Fortune Database Server	SQL	Accept	short

This rule is to permit the customer web server on the external service network to access the customer database server on the application service network. Our database servers are all Oracle web servers running SQL. The breakdown of the rules is as follows:

- **Number:** This is rule 2 in the policy set. There is no security reason this rule must be near the top of the rule set. This rule is rule 2 because due to latency and efficiency considerations we are giving precedence to customer rules.
- **Source:** The source is Customer Web Server because we want to restrict access only to the machine that requires it.
- **Destination:** This is limited only to the two database servers that are required by the customer web server.
- **Service:** Our database servers are all Oracle databases running SQL. The Firewall-1 SQL service opens port 1521/tcp, 1525/tcp, and 1526/tcp. Not all of these ports may be required by our servers, so we may need to create our own service definition to restrict the ports.

- **Action:** We are accepting this traffic.
- **Track:** We log this short because it is considered normal traffic. This provides enough information to do some analysis on the traffic if required.

## Supplier and Partner Access Rules

Number	Source	Destination	Service	Action	Track
3	Partner_Nets Supplier_Nets	Partner/Supplier VPN Server	IPSec	Accept	short

This rule is to permit partners and suppliers to access the partner/supplier VPN Server on the external service network. All access is via http and https. The breakdown of the rules is as follows:

- **Number:** This is rule 3 in the policy set. There is no security reason this rule must be near the beginning of the rule set. This rule is located where it is because due to latency and efficiency considerations we are making supplier and partner rules second priority after customer rules.
- **Source:** The source is limited to just the set of networks which belong to our partners and suppliers. Partner\_Nets is the list of partner networks. Supplier\_Nets is the list of supplier networks. If a new partner or supplier is added their networks will need to be added into these groups.
- **Destination:** This is the partner/supplier VPN server. We restrict the access to only the machine that is required
- **Service:** The service allowed is IPSec. The Firewall-1 service object for IPSEC contains AH (authentication header), ESP (Encapsulating Security Payload), ISAKMP, and SKIP. Although the possibility of a compromise is small, we may wish to consider AH and SKIP from our object definition since we don't use them for our implementation.
- **Action:** We are accepting this traffic.
- **Track:** We log this short because it is considered normal traffic. This provides enough information to do some analysis on the traffic if required.

Number	Source	Destination	Service	Action	Track
4	Partner/Supplier VPN Server	Partner_Nets Supplier_Nets	ESP	Accept	short

Because we are using ESP for our VPN authentication, we need to permit ESP back out from our VPN server to the partner and supplier networks. The breakdown of the rules is as follows:

- **Number:** This is rule 4 in the policy set. This rule is located where it is because due to latency and efficiency considerations we are making supplier and partner rules second priority after customer rules.
- **Source:** This is the partner/supplier VPN server. We restrict the access to only the machine that is required.
- **Destination:** The destination is limited to just the set of networks which belong to our partners and suppliers. Partner\_Nets is the list of partner networks. Supplier\_Nets is the list of supplier networks. If a new partner or supplier is added their networks will need to be added into these groups.
- **Service:** The service allowed is ESP. The IPSec Encapsulating Security Payload (ESP) attempts to guarantee the integrity of the data.
- **Action:** We are accepting this traffic.
- **Track:** We log this short because it is considered normal traffic. This provides enough information to do some analysis on the traffic if required.

Number	Source	Destination	Service	Action	Track
5	Partner VPN IP Pool	Partner Web Server	HTTP	Accept	short

			HTTPS		
--	--	--	-------	--	--

This rule is to permit partners to access the partner application web server on the application service network. All access is via http and https. The breakdown of the rule is as follows:

- **Number:** This is rule 5 in the policy set. This rule is located where it is because due to latency and efficiency considerations we are making supplier and partner rules second priority after customer rules.
- **Source:** The source is Partner VPN IP Pool because as each partner connects to the VPN server they are allocated an IP out of a pool reserved for their company. All of the partner IP pools are in a contiguous range. This range is the Partner VPN IP Pool. This methodology permits us to implement individual partner rules if we wish to at some point in the future.
- **Destination:** This is the partner web application web server on the application service network. We restrict the access to only the machine that is required to reduce the possibility of attacks from partners on other machines which may be running http or https services for maintenance, may not have been properly hardened, or may have had service opened up by accident through maintenance actions such as patching.
- **Service:** All the partner applications are web based. Therefore HTTP and HTTPS are the only services required.
- **Action:** We are accepting this traffic.
- **Track:** We log this short because it is considered normal traffic. This provides enough information to do some analysis on the traffic if required.

Number	Source	Destination	Service	Action	Track
6	Supplier VPN IP Pool	Supplier Web Server	HTTP HTTPS	Accept	short

This rule is to permit suppliers to access the supplier application web server on the application service network. All access is via http and https. The breakdown of the rule is as follows:

- **Number:** This is rule 6 in the policy set. This rule is located where it is because due to latency and efficiency considerations we are making supplier and partner rules second priority after customer rules.
- **Source:** The source is Supplier VPN IP Pool because as each supplier connects to the VPN server they are allocated an IP out of a pool reserved for their company. All of the supplier IP pools are in a contiguous range. This range is the Supplier VPN IP Pool. This methodology permits us to implement individual supplier rules if we wish to at some point in the future.
- **Destination:** This is the supplier web application web server on the application service network. We restrict the access to only the machine that is required to reduce the possibility of attacks from suppliers on other machines which may be running http or https services for maintenance, may not have been properly hardened, or may have had service opened up by accident through maintenance actions such as patching.
- **Service:** All the supplier applications are web based. Therefore HTTP and HTTPS are the only services required.
- **Action:** We are accepting this traffic.
- **Track:** We log this short because it is considered normal traffic. This provides enough information to do some analysis on the traffic if required.

## Maintenance and Support Rules

Number	Source	Destination	Service	Action	Track
7	Corporate Network	DMZ External Service Network Application Service Network	HTTP HTTPS IPSEC SSH Telnet	Accept	short

This rule is to permit support people on the corporate network to access all of the e-business network components for troubleshooting, and ongoing maintenance and support. The breakdown of the rule is as follows:

- **Number:** This is rule 7 in the policy set. This rule is located where it is because due to latency and efficiency considerations we gave higher priority to customer, supplier and partner rules.
- **Source:** The source is the entire corporate network IP range. This is probably not the most secure arrangement. I would have rather had this reduced to only designated support nodes, however numerous groups are involved in the support of this infrastructure, and they did not want to be restricted.
- **Destination:** This is all of the external networks in our e-business infrastructure. This permits support and maintenance of all devices involved. Even though IPSec is only supported on the supplier/partner VPN server, IPSec is permitted to the entire network to avoid adding another rule, and marginally decreasing the efficiency of the firewall.
- **Service:** HTTP and HTTPS are required for troubleshooting and support of the customer, partner, and supplier web servers. IPSec is required to troubleshoot and support IPSEC connections to the partner/supplier VPN Server. SSH and telnet are required for support and maintenance of the individual components in the e-business networks. Most of the devices can handle SSH, but telnet is required for the couple that don't (routers and switches).
- **Action:** We are accepting this traffic.
- **Track:** We log this short because it is considered normal traffic. This provides enough information to do some analysis on the traffic if required.

Number	Source	Destination	Service	Action	Track
8	External Service Network	Corporate Network	ESP	Accept	short

Because we are using ESP for our VPN authentication, we need to permit ESP back to our corporate network. Since there are no devices that support IPSec on the other networks, this traffic is only permitted from the external service network. This rule could have been combined with rule number four. It was not due to security reasons. If this rule were combined with rule four and somebody were to add another service to the rule, that service would also have access to the corporate network. The breakdown of the rules is as follows:

- **Number:** This is rule 8 in the policy set. This rule is located where it is because due to latency and efficiency considerations we gave higher priority to customer, supplier and partner rules.
- **Source:** This is the external service network. This could be restricted to just the supplier/partner VPN server for added security.
- **Destination:** The destination is limited to the corporate network.
- **Service:** The service allowed is ESP. The IPSec Encapsulating Security Payload attempts to guarantee the integrity of the data in the VPN.
- **Action:** We are accepting this traffic.
- **Track:** We log this short because it is considered normal traffic. This provides enough information to do some analysis on the traffic if required.

## Infrastructure Rules

Number	Source	Destination	Service	Action	Track
9	Any	External DNS Server	DNS (53/UDP)	Accept	short

This rule is to permit DNS queries from the Internet and forwarders on our corporate network. The breakdown of the rule is as follows:

- **Number:** This is rule 9 in the policy set. This rule is located where it is because due to latency and efficiency considerations we gave higher priority to customer, supplier, partner and support and maintenance rules.

- **Source:** The source is Any. This permits all IPs whether on the Internet, the corporate network, or in this e-business infrastructure to access the DNS server. The corporate network requires access for the forwarders on the corporate network to access the external DNS addresses.
- **Destination:** This is our external DNS server. Although other machines in this infrastructure should not be running DNS, we restrict the access to only the DNS machine to reduce the possibility of attacks on other machines may not have been properly hardened, or may have had DNS turned on by accident through maintenance actions such as patching.
- **Service:** The enabled service is DNS (domain-udp, port 53/udp). Firewall-1's DNS service by default contains both domain-udp (53/udp) and domain-tcp (53/tcp). We have removed domain-tcp from the object definition, on the grounds that we will not be permitting zone transfers. It will be necessary to watch carefully since removing domain-tcp also means that long dns-queries will not be supported. It is important to note that this will not work unless "Accept UDP replies" is enabled on the Firewall-1 Security Properties screen. Without "Accept UDP replies" enabled, the queries will still be allowed through the firewall, but the replies will be dropped on the firewall.
- **Action:** We are accepting this traffic.
- **Track:** We log this short because it is considered normal traffic. This provides enough information to do some analysis on the traffic if required. It may be necessary, if the volume of these queries becomes high enough, to stop logging this traffic to stop log flooding, and to reduce the load on the firewall. For now we are logging it to see the volume generated.

Number	Source	Destination	Service	Action	Track
10	Management Stations	DMZ External Service Network Application Service Network	SNMP ICMP_echo _request	Accept	short

This rule is to permit SNMP monitoring of the e-business infrastructure from designated HP OpenView management stations on the corporate network. Since SNMP can be used as a data-gathering tool by crackers, access is only permitted from designated management stations. The breakdown of the rule is as follows:

- **Number:** This is rule 10 in the policy set. This rule is located where it is because due to latency and efficiency considerations we gave higher priority to customer, supplier, partner and maintenance and support rules.
- **Source:** The source is designated HP OpenView management stations.
- **Destination:** Destination is all of the e-business networks. This will permit monitoring of all devices on these networks.
- **Service:** Two services are required for HP OpenView, SNMP (udp 162), and ICMP echo request. According to the documentation HP OpenView does not require ICMP echo request if SNMP is available, but it is our experience that it works much better with it. It is important to note that SNMP will not work unless "Accept UDP replies" is enabled on the Firewall-1 Security Properties screen. Without "Accept UDP replies" enabled, the SNMP queries will still be allowed through the firewall, but the replies will be dropped on the firewall.
- **Action:** We are accepting this traffic.
- **Track:** We log this short because it is considered normal traffic. This provides enough information to do some analysis on the traffic if required.

Number	Source	Destination	Service	Action	Track
11	DMZ External Service Network Application Service Network	Management Stations	SNMP_Trap ICMP_echo _reply	Accept	short

This rule is to aid SNMP monitoring of the e-business infrastructure by permitting SNMP traps to designated HP OpenView management stations on the corporate network. The breakdown of the rule is as follows:

- **Number:** This is rule 11 in the policy set. This rule is located where it is because due to latency and efficiency considerations we gave higher priority to customer, supplier, partner and maintenance and support rules.
- **Source:** The source is all of the e-business networks. This will permit all devices on these networks to send SNMP traps back to the management stations.
- **Destination:** The destination is designated HP OpenView management stations.
- **Service:** Two services are on this rule. The first is SNMP trap (udp 162), this is permitted so the servers and network devices that support SNMP trap can trap back to the HP OpenView management servers. The second is ICMP echo reply. This is permitted to allow the ICMP echo requests in rule 10 to complete successfully.
- **Action:** We are accepting this traffic.
- **Track:** We log this short because it is considered normal traffic. This provides enough information to do some analysis on the traffic if required.

Number	Source	Destination	Service	Action	Track
12	DMZ External Service Network Application Service Network	Internet	NTP	Accept	short

This rule is to permit Network Time Protocol to access the Internet. This will ensure that all devices in the infrastructure are time synchronized to permit correlation of logs for investigative purposes. The breakdown of the rule is as follows:

- **Number:** This is rule 12 in the policy set. This rule is located where it is because due to latency and efficiency considerations we gave higher priority to customer, supplier, partner and maintenance and support rules.
- **Source:** The source is any of the e-business networks.
- **Destination:** Destination is any NTP server on the Internet. It would probably be better to designate only a couple of servers, but the NTP client that most of the servers are using has a large set of servers embedded, so it was easier to open it up, than to modify the NTP client.
- **Service:** The service is Network Time Protocol (NTP), 123/tcp. Note that the Firewall-1 NTP service object also contains ntp-udp (123/udp). We removed it from this object.
- **Action:** We are accepting this traffic.
- **Track:** We log this short because it is considered normal traffic. This provides enough information to do some analysis on the traffic if required.

## Explicit Drop Rule

Number	Source	Destination	Service	Action	Track
13	Any	Any	Any	Drop	Long

This rule is the explicit drop rule. There is an implied drop rule in Firewall-1 that is executed last which drops all remaining traffic. However that rule does not log. We want to log the dropped traffic for later analysis. We want to analyze this traffic for a number of reasons. First, to detect misconfigured servers that are generating noise, so they can be properly configured and reduce the load on the firewall. Second, to detect possible malicious traffic that your firewall is blocking. Third to reinforce the deny everything except what is permitted policy. The breakdown of the rule is as follows:

- **Number:** This is rule 13 in the policy set. This rule needs to be last because any rules placed after it will never execute.
- **Source:** The source is any. We want to drop all remaining traffic regardless of origin.
- **Destination:** The destination is any. We want to drop all remaining traffic regardless of destination.
- **Service:** The service is any. We want to drop all remaining traffic regardless of service.
- **Action:** We are dropping this traffic. It has not been explicitly permitted so it must be dropped.

- **Track:** We log this long, because it is not considered normal traffic. This permits us to have lots of data for later analysis.

An excellent treatise on screening of inbound traffic, and specifically implementation of firewall filters is available on the SANS website at [http://www.sans.org/infosecFAQ/firewall/fw\\_filters.htm](http://www.sans.org/infosecFAQ/firewall/fw_filters.htm)<sup>6</sup>

## **VPN**

The partner and supplier VPN server is a Nortel Networks Contivity 2600 VPN Server. All VPN tunnels are initiated using the Nortel Networks Contivity Extranet Client version 2.62 from the supplier/partners Microsoft Windows desktop on their corporate network.

Partner and supplier access is provided via IPSec VPN to the partner/supplier VPN server on the external service network. Authentication is performed with a protected User ID and Password through the ISAKMP key management protocol. As well as individual logins group userids and passwords are also used. The group userids permit the VPN server to assign each partner to a unique group and to assign them an IP address out of a pool explicitly defined for their company. This IP is used to limit access through the firewall to only appropriate web server on the application service network. A client policy is also applied on the VPN server to each group that limits which protocols and destinations are supported.

## **General Configuration of the VPN Server**

Two methods of encryption are supported:

- ESP - Triple DES with MD5 Integrity
  - Encapsulated Security Payload Triple DES 168-bit key. This is the highest encryption supported. It is important to note that if the client (or server) does not support triple DES, the client and server will negotiate downward until they reach a compatible encryption setting.
- AH - Authentication Only (HMAC-MD5)
  - The Authentication Header Message Authentication Code Message Digest 5 It uses a 128-bit hash for authentication of source, destination and payload. It does not encrypt data. In this case the only encryption provided is through compression. Authentication does not work through Network Address Translation (NAT).

Perfect Forward Secrecy (PFS) is enabled. This ensures that subsequent encryption keys are not derived from previous encryption keys. This reduces the possibility that if a key has been compromised that subsequent keys can be derived from that key.

All tunnels are terminated after 4 hours. This reduces the possibility of a partner/supplier leaving a tunnel up and their machine being compromised.

Client Screen Saver Password Required is enabled. This forces the partner/supplier to use a password in association with their machines screen saver. If the user does not use a screensaver on his PC and the user is connected via VPN, and the screen saver starts up, the tunnel is terminated. This reduces the exposure when a user leaves his PC unattended without the screen saver locked.

Allow Password Storage on Client is disabled. The Extranet Access Client supports the storage of passwords with the tunnel configuration. By disabling this feature we can require that the partner/supplier enter his password each time he requests an IPSec tunnel to the partner/supplier VPN server. This reduces the possibility of someone gaining access to the partner/suppliers PC and being able to initiate a tunnel.

The Rekey Timeout is set to 15 minutes. This means that the encryption key used by the VPN tunnel will be renegotiated every 15 minutes. This reduces the impact of someone compromising the key.

VPNs not initiated by the Extranet Access Client are not supported. Non-IPSec VPNs, and IPSec VPNs from other VPN clients are disallowed.

Compression is enabled using LZS compression for all connections.

Split tunneling (split horizon) is permitted. Initially this was disabled, but suppliers and partners complained because they were not able to print their order information to local printers. Several precautions were taken that go a long way toward mitigating the security impact of this decision.

- The Contivity Extranet Switch has a security feature which drops packets that do not have the IP address assigned to the VPN as its source address. This ensures that packets originate from the appropriate host.
- Each user group has a client policy defined which permits only HTTP/HTTPS access to the appropriate server. All other traffic is disallowed. This means that when a partner or supplier connects the only successful traffic will be to the appropriate web server.

## **Internal Filter Router**

The basic security policy for the internal filter router is to deny all traffic except what is explicitly required.

The filters for the internal router were not required for this practical. Because of this, and for brevity, they will not be included.

## **Testing Methodology**

Testing was performed to verify the validity of this infrastructure using ISS Internet Scanner. After the networks were put in place, and before any servers were put on the network, the router, switches and firewalls were explicitly scanned as well as the entire e-business address space. A detailed analysis was performed on the results and appropriate action taken.

After the servers were all put in place, and all the services tested. The scans were repeated and again the results were analyzed to verify the router filters and firewall rules. Action was taken to understand and repair any abnormalities.

This work was performed on top of the individual server scans that are performed before a server may be placed in service, and the regular monthly scans.

For more information on the scans, look at the methodology in part 3 of this practical.

### Assignment 3 - Audit Your Security Architecture (25 Points)

You have been assigned to provide technical support for a comprehensive information systems audit for GIAC Enterprises. You are required to audit the Primary Firewall described in Assignments 1 and 2. Your assignment is to:

1. *Plan the assessment. Describe the technical approach you recommend to assess your perimeter. Be certain to include considerations such as what shift or day you would do the assessment. Estimate costs and level of effort. Identify risks and considerations.*
2. *Implement the assessment. Validate that the Primary Firewall is actually implementing the security policy. Be certain to state exactly how you do this, including the tools and commands used. Include screen shots in your report if possible.*
3. *Conduct a perimeter analysis. Based on your assessment (and referring to data from your assessment), analyze the perimeter defense and make recommendations for improvements or alternate architectures. Diagrams are strongly recommended for this part of the assignment.*

*Note: DO NOT simply submit the output of nmap or a similar tool here. It is fine to use any assessment tool you choose, but annotate the output.*

## Planning The Assessment

Since the networks and servers in this practical don't really exist, it is impossible for me to perform an assessment on them. But here is how I would assess them.

The assessment will focus on multiple areas. First of all, an assessment of just the firewall, without considering the other e-business network components, and the policies, procedures and best practices involved in operating the components would not be a thorough assessment. For the sake of this practical I will describe how I would plan the assessment for all components, and only describe the assessment of the firewall.

I would plan my assessment to include the following components.

**Network Assessment** – this would include comprehensive vulnerability scans from all angles. The network should be scanned from the Internet, from all of the e-business networks, and from our corporate network. It is important that the Firewall-1 management server be considered in this evaluation. Even though it is on the corporate network, the security of the management station is important.

Scans are carried out off hours to minimize the possibility of impacting production machines.

After each scan the scan output, the firewall logs, and the IDS logs will be analyzed to ensure that each component is as it should be.

The purpose of the network assessment is three-fold. First, to validate the rules and filters implemented on the routers, and firewall. Second, to verify the hardening of all servers and network components. Third, to assess all components for vulnerabilities.

**Organizational Assessment** – Review and assess the security policy and procedures used in the installation, support, and ongoing maintenance of the e-business environment.

The purpose of this review is to ensure that the security policy is up to date, and that sound management procedures are being used in the administration and maintenance of this environment.

**Physical Security** – Review and assess the physical security of all components in the network. Once again it is important that the Firewall-1 management server be considered.

The purpose of this assessment is to ensure the physical security of all components, and to ensure that sound procedures are being followed for the physical maintenance of all components.

The output of the assessments is a comprehensive report assessing the current state of the network components and recommendations on how to improve the security.

## **Cost estimate**

Assuming a loaded labor rate of \$400 per day per person

Assessment planning – 2 people, 5 days  
Network assessment – 2 people, 6 days  
Organizational Assessment – 2 people, 5 days  
Physical Security Assessment – 2 people, 1 day  
Detailed Report – 2 people, 5 days

Total – 44 man days = \$17,600

## **Pre-Planning Data Gathering**

Before you can adequately plan the assessment a certain amount of data gathering is required. You will need to collect at least the following information.

- Network diagrams detailing the topology being tested.
- Corporate policies, procedures, and best practices involved in the implementation, support and maintenance of this networks.
- A list of organizations involved in the implementation, support and maintenance of the e-business components, and a list of contacts in those organizations. These contacts will be your sources for data gathering, as well as the contacts in case any difficulties are experience during the assessments.

## **Planning the Assessment**

Once you have all collected enough information it is time to plan the assessment. This will include a detailed project plan, including timing of tests and the methods to be used for testing.

## **Pre-Assessment Executive Buy-In**

Before beginning the assessments, you should have a meeting with GIAC enterprises management. The intent of this meeting is to ensure a clear understanding of the scope of the assessment is, and when it will occur. One of the outputs from this meeting is sign off on the assessment plan.

## **Pre-Assessment Meeting**

Before implementing the assessments a meeting needs to be held with all parties involved in the support and maintenance of the e-business components. The meeting is to assure that all parties clearly understand what is to be undertaken, when it is to occur, and who they are to call if any difficulties should occur.

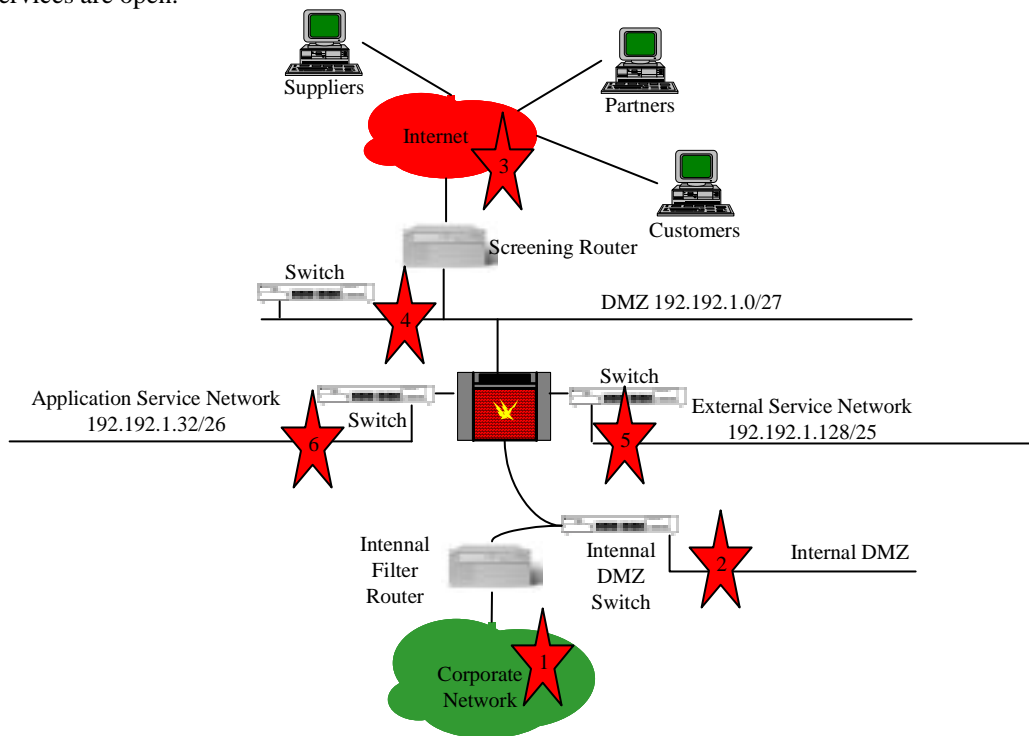
## **Implementing the Assessment**

### ***Network Assessment***

The assessment tool that we are going to use for the network assessment is ISS Internet Scanner. Although several good freeware and shareware assessment tools exist, such as SAINT, SARA, and nmap, ISS Internet Scanner is a product that we have easily available to us, and which we have a good working knowledge of.

We will be scanning our e-business network from the six points on the network diagram below. We will be scanning all of the e-business address range (192.192.1.0/24) on every scan. For the scans originating on the e-business networks, we will also be attempting to scan our corporate network. No traffic should get through to the corporate network, but since the routing exists for return traffic to get back to the corporate

network we are going to attempt to scan it anyway. After each scan we will stop and analyze the firewall logs, the IDS logs when applicable, and the scan output. We will be paying special attention to components on the network segment we are scanning from, to ensure that they are properly hardened and no unessential services are open.



## Scan 1: Corporate Network to e-business networks

**Methodology:** Initiate a scan from the corporate network to all of the e-business address range. Also scan the Internal DMZ network range, and the Firewall-1 management server

### Addresses to be scanned:

- 192.192.1.0/24
- Internal DMZ Network
- Firewall-1 Management Station
- Corporate Network side of the internal filter router

### What to watch for:

- Non-essential services disabled on the internal filter router.
- No machines exist on the Internal DMZ except the filter router, the internal DMZ switch and the firewall.
- No unexpected hosts on the e-business network range.
- Internal filter router passes only expected traffic.
- Firewall passes only expected traffic.
- Firewall-1 management server is appropriately hardened.

## Scan 2: Internal DMZ Network to e-business networks

**Methodology:** Initiate a scan from the internal DMZ network to all of the e-business address range and the internal DMZ network.

### Addresses to be scanned:

- 192.192.1.0/24
- IMZ Network

### What to watch for:

- No non-essential services present on the internal DMZ switch.

- No machines exist on the Internal DMZ except the filter router, the internal DMZ switch and the firewall.
- Firewall passes only expected traffic.

### **Scan 3: Internet to e-business networks**

**Methodology:** Initiate a scan from the Internet network to all of the e-business address range and the ISP side of the external screening router.

**Addresses to be scanned:**

- 192.192.1.0/24
- ISP side of the external screening router.

**What to watch for:**

- No non-essential services visible on the external screening router.
- Telnet cannot connect to the external screening router from the Internet
- Cannot access any unexpected hosts or unexpected services on existing hosts.
- Firewall passes only expected traffic.

### **Scan 4: E-business DMZ to e-business networks, internal DMZ network and corporate network**

**Methodology:** Initiate a scan from the internal DMZ network to all of the e-business address range, and the corporate network including the internal DMZ network.

**Addresses to be scanned:**

- 192.192.1.0/24
- Corporate network address range, including the internal DMZ network.

**What to watch for:**

- No non-essential services visible on the external screening router.
- No hosts appear on the DMZ except the firewall, the external screening router, and the switch. The IDS machine should not show up because it is supposed to be stealth.
- Firewall passes only expected traffic.
- No non-essential services present on the DMZ switch.
- No traffic passes to corporate network address range.

### **Scan 5: External Service Network to e-business networks, internal DMZ network and corporate network**

**Methodology:** Initiate a scan from the external service network to all of the e-business address range, the internal DMZ network.

**Addresses to be scanned:**

- 192.192.1.0/24
- Corporate network address range, including the internal DMZ network.

**What to watch for:**

- No unexpected hosts are present.
- No non-essential services are visible on the web server, dns server or VPN server.
- Firewall passes only expected traffic.
- No non-essential services present on the external service network switch.
- No traffic passes to corporate network address range.

### **Scan 6: External Service Network to e-business networks, internal DMZ network and corporate network**

**Methodology:** Initiate a scan from the application service network to all of the e-business address range, the internal DMZ network.

**Addresses to be scanned:**

- 192.192.1.0/24

- Corporate network address range, including the internal DMZ network.

**What to watch for:**

- No unexpected hosts are present.
- No non-essential services are visible on the existing servers.
- Firewall passes only expected traffic.
- No non-essential services present on the application service network switch.
- No traffic passes to corporate network address range.

#### Assignment 4 - Design Under Fire (25 Points)

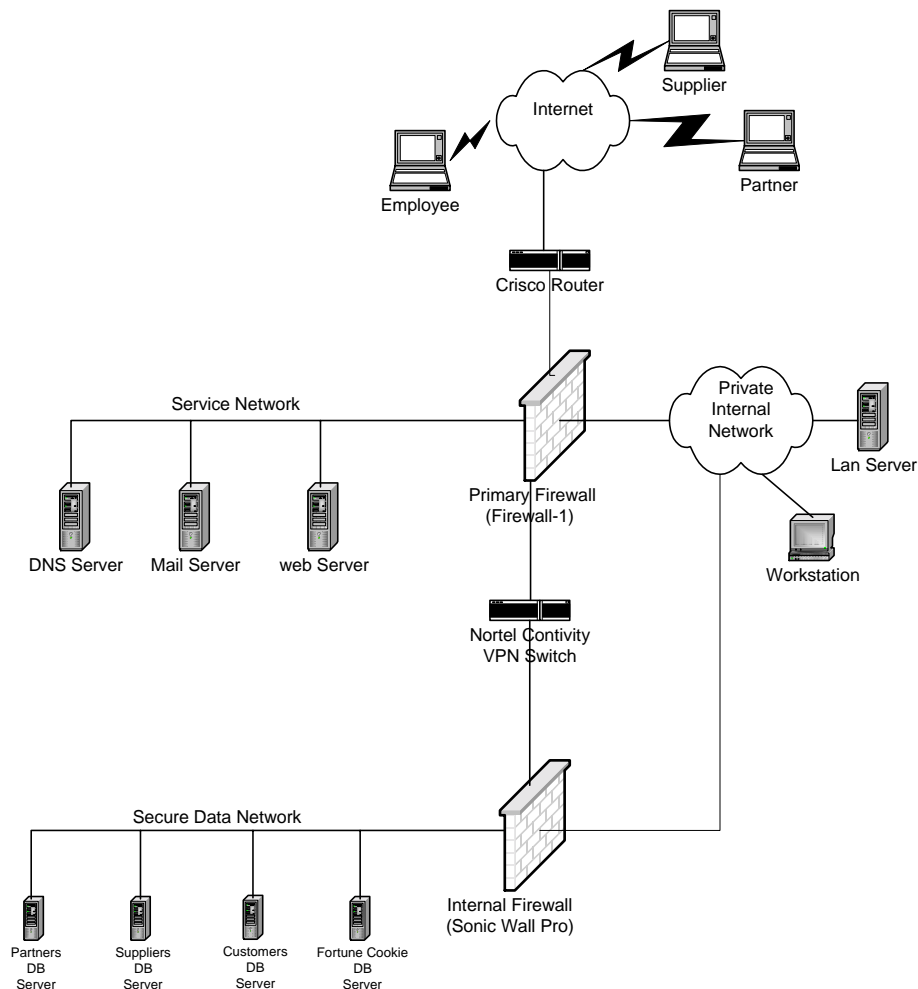
The purpose of this exercise is to help you think about threats to your network and therefore develop a more robust design. Keep in mind that the next certification group will be attacking your architecture!

Select a network design from any previously posted GCFW practical (<http://www.sans.org/gjactc/gcfw.htm>) and paste the graphic into your submission. Be certain to list the URL of the practical you are using. Design the following three attacks against the architecture:

1. An attack against the firewall itself. Research vulnerabilities that have been found for the type of firewall chosen for the design. Choose an attack and explain the results of running that attack against the firewall.
2. A denial of service attack. Subject the design to a theoretical attack from 50 compromised cable modem/DSL systems using TCP SYN, UDP, or ICMP floods. Describe the countermeasures that can be put into place to mitigate the attack that you chose.
3. An attack plan to compromise an internal system through the perimeter system. Select a target, explain your reasons for choosing that target, and describe the process to compromise the target.

**Note:** this is the second time this assignment has been used. The first time, a number of students came up with magical "hand-waving" attacks. You must supply documentation (preferably a URL) for any vulnerability you use in your attack, and the exploit code that you use to accomplish the attack. The purpose of this exercise is for the student to clearly demonstrate they understand that firewall and perimeter systems are not magic "silver bullets" immune to all attacks.

The network design I chose for attack is Said Nurhussein's at [http://www.sans.org/y2k/practical/said\\_nurhussein\\_gcfw.doc](http://www.sans.org/y2k/practical/said_nurhussein_gcfw.doc)



# Attack Against the Firewall

The firewall is a Checkpoint Firewall-1 version 4.1 running on Windows NT. Using Security Focus's Bugtraq vulnerabilities depository at <http://www.securityfocus.com/vdb/top.html> we find the following vulnerabilities published in the last year against Firewall-1:

- [2001-01-17: Check Point Firewall-1 4.1 Denial of Service Vulnerability](#)
- [2000-12-14: Check Point Firewall-1 Fast Mode TCP Fragment Vulnerability](#)
- [2000-11-01: Checkpoint Firewall-1 Valid Username Vulnerability](#)
- [2000-08-15: Check Point Firewall-1 Session Agent Dictionary Attack Vulnerability](#)
- [2000-08-02: Check Point Firewall-1 Unauthorized RSH/REXEC Connection Vulnerability](#)
- [2000-07-05: Check Point Firewall-1 Spoofed Source Denial of Service Vulnerability](#)
- [2000-06-30: Check Point Firewall-1 SMTP Resource Exhaustion Vulnerability](#)
- [2000-06-06: Check Point Firewall-1 Fragmented Packets DoS Vulnerability](#)
- [2000-03-11: Check Point Firewall-1 Internal Address Leakage Vulnerability](#)
- [2000-03-10: Multiple Firewall Vendor FTP "ALG" Client Vulnerability](#)

The ICAT vulnerability database (<http://icat.nist.gov/icat.taf>) shows the same list of vulnerabilities, but in some cases shows better details.

The fourth one in the list sounds particularly interesting, in that it may get us past the firewall, is there is not a patch yet for the vulnerability and there is a published exploit.

## [2000-08-15: Check Point Firewall-1 Session Agent Dictionary Attack Vulnerability](#)

*A vulnerability exists in all versions of the Check Point Session Agent, part of Firewall-1. Session Agent works in such a way that the firewall will establish a connection back to the client machine. Upon doing so, it will prompt for a username, and if the username exists, a password. Upon failure, it will reprompt indefinitely. This allows for a simple brute force attack against the username and password.*

*Currently the SecurityFocus staff are not aware of any vendor supplied patches for this issue. If you feel we are in error or are aware of more recent information, please mail us at: [vuldb@securityfocus.com](mailto:vuldb@securityfocus.com).*

*Nelson Brito <[nelson@secunet.com.br](mailto:nelson@secunet.com.br) || [nelson@sekure.org](mailto:nelson@sekure.org)> provided brute-fw1-agent.pl exploit.*

*Gregory Duchemin <[c3rb3r@hotmail.com](mailto:c3rb3r@hotmail.com)> submitted fwsa.sh exploit.*

- [/data/vulnerabilities/exploits/fwsa.sh](#)
- [/data/vulnerabilities/exploits/brute-fw1-agent.pl](#)

More information is available on this exploit at <http://www.securityfocus.com/archive/1/76389>.

To exploit this vulnerability you have to count on an idiosyncrasy of the Firewall-1 session agent. The session agent permits you to authenticate your session using a userid and password. The userid and password are 8 characters in length and have to be ASCII printable characters.

When you authenticate to the session agent it prompts with

331 User:

If you enter an invalid userid you will get the response

```
220 User .... not found
530 NOTOK
```

If the userid was valid the firewall will respond with

```
331 *FireWall-1 password:
```

This means that we can tell which userids are valid based on the response from the firewall.

The second part of this vulnerability is that if an invalid password is entered session agent does not close, but rather waits until it receives a valid password.

The first exploit (fwsa.sh) is a bash shell program takes an IP address list. The documentation suggests you use nmap to probe an IP range looking for session agent listening on port 261.

The exploit can operate in 4 different modes:

- Attack 1 is called password recovery. Given a userid it attempts to brute force the password.
- Attack 2 is a simple denial of service attack that opens session agent connections and enters nothing. It will tie up all the connections so nobody can log in.
- Attack 3 is a malicious denial of service attack. It repeats session agent connections in a infinite loop, sending garbage every time. This will apparently crash some systems.
- Attack 4 is a brute force dictionary password attack. It will attempt to guess userids and then attempt a brute force dictionary attack against that id.

The second exploit (brute-fw1-agent.pl) is a perl script which given a text file runs each of the text file entries against a Firewall-1 session agent to see if any of them are valid userids. It outputs the valid userids.

## Denial of Service Attack

The denial of service attack I chose to launch is trinoo<sup>16</sup>. Trinoo is a pretty standard distributed UDP flood style denial of service attack tool. It is composed of a master portion, and a daemon portion. The master runs on a Red Hat Linux 6.0 system, and the daemons can be installed on either Linux 6.0 or Solaris 2.5.1.

Typically the attacker will use nmap or some other such tool to compile an IP list of Linux or Solaris 2.X machines that have remotely exploitable buffer overflow vulnerabilities. Once the list is compiled an exploit is run that exploits the vulnerability and installs a root command shell that listens on port 1524. This is not the trinoo daemon itself, but rather a shell that can be used to verify that the exploit succeeded and that can be used later to install whatever the attacker wishes.

From this list of compromised systems a subset of machines with the appropriate architecture are chosen for the trinoo daemon installation. Pre-compiled binaries are created on an account somewhere on the Internet. An install script is downloaded to the “owned” machines through the root shell on port 1524, and either the trinoo daemon or the trinoo master is installed on the compromised hosts. The attacker has setup a network of masters and daemons, thus giving the attacker a large number of geographically distributed systems from which to launch a distributed denial of service attack.

By sending the command “aaa l44adsl IP” to the master, where IP is the IP address of the host which is to be DOSed, the master will send out a command of “aaa pass IP” to all of the daemons. This will result in all of the daemons sending UDP packets to random ports on the attacked machine for 120 seconds. Master commands exist which allow you to change the length of an attack (up to 2000 seconds), and the packet size to be sent, as well as a command to permit attacks on multiple hosts at the same time.

A detailed analysis of trinoo is available at <http://staff.washington.edu/dittrich/misc/trinoo.analysis.txt>.

Defending against Trinoo is not an easy task. The easiest way, but one that may not be popular with users, is to block all inbound UDP unless it is explicitly required.

## Compromising an Internal Server

There are three tempting targets on the network; the web server, dns server, and mail server. The typical software running on these is often to buffer overflow type problems, which with luck can result in achieving a root shell on the server.

Let's take a look at the web server. If I probe it with nmap, and find out it is a Microsoft IIS version 4.0 web server, I can go to [www.securityfocus.com](http://www.securityfocus.com) and go through the bugtraq vulnerability archives and discover there are two exploits which if executed could give me an Administrator login on the web server.

### [2000-10-17: Microsoft IIS and PWS Extended Unicode Directory Traversal Vulnerability](#)

*Microsoft IIS 4.0 and 5.0 are both vulnerable to double dot "." directory traversal exploitation if extended UNICODE character representations are used in substitution for "/" and "\".*

*Unauthenticated users may access any known file in the context of the IUSR\_machinename account. The IUSR\_machinename account is a member of the Everyone and Users groups by default, therefore, any file on the same logical drive as any web-accessible file that is accessible to these groups can be deleted, modified, or executed. Successful exploitation would yield the same privileges as a user who could successfully log onto the system to a remote user possessing no credentials whatsoever.*

There are a large number of exploits available for this vulnerability, a number of which will permit the attacker to run commands as Administrator. A partial list is available at <http://www.securityfocus.com/vdb/bottom.html?section=exploit&vid=1806>.

The second exploit that can be used after the first is successful is:

### [2000-11-06: Microsoft IIS 4.0 ISAPI Buffer Overflow Vulnerability](#)

*The ASP ISAPI file parser does not properly execute certain malformed ASP files that contain scripts with the LANGUAGE parameter containing a buffer of over 2200 characters and have the RUNAT value set as 'server'. Depending on the data entered into the buffer, a denial of service attack could be launched or arbitrary code could be executed under the SYSTEM privilege level in the event that a malicious ASP file were locally executed on IIS.*

*This issue has been resolved by a number of Microsoft IIS patches. The patch below will eliminate this vulnerability:*

*Microsoft IIS 4.0:*

*Microsoft patch secsesi*

*<http://download.microsoft.com/download/winntsp/Patch/Q274149/NT4/EN-US/secsesi.exe>*

*eEye Digital Security <[info@eEye.com](mailto:info@eEye.com)> has released the following exploit:  
[/data/vulnerabilities/exploits/IISHack1.5.zip](#)*

## References

1. Irby, David, "Firewalk: Can Attackers See Through Your Firewall", December 10, 2000.  
<http://www.sans.org/infosecFAQ/firewall/firewalk.htm>
2. SANS Institute, "How To Eliminate The Ten Most Critical Internet Security Threats, The Experts' Consensus", Version 1.3 2 January 18, 2001. <http://www.sans.org/10threats.doc>
3. Richard Langley, "Securing Your Internet Access Router", January 23,2001.  
<http://www.sans.org/infosecFAQ/firewall/router.htm>
4. Winters, Scott, "Top Ten Blocking Recommendations Using Cisco ACLs", August 15, 2000.  
[http://www.sans.org/infosecFAQ/firewall/blocking\\_cisco.htm](http://www.sans.org/infosecFAQ/firewall/blocking_cisco.htm)
5. Brockaway, Rockie, "Layered Security: An ISP Case Study with Cisco and Solaris".  
[http://www.sans.org/infosecFAQ/firewall/layered\\_sec.htm](http://www.sans.org/infosecFAQ/firewall/layered_sec.htm)
6. Thompson, Rick, "GIAC Practical: Implementation of Firewall Filters", August 14, 2000.  
[http://www.sans.org/infosecFAQ/firewall/fw\\_filters.htm](http://www.sans.org/infosecFAQ/firewall/fw_filters.htm)
7. Nortel Networks, "BCC Quick Reference", date unknown.  
[http://www25.nortelnetworks.com/library/tpubs/pdf/router/soft1420/308602\\_1420\\_00.PDF](http://www25.nortelnetworks.com/library/tpubs/pdf/router/soft1420/308602_1420_00.PDF)
8. Nortel Networks, "Using the Bay Command Console", August 2000.  
[http://www25.nortelnetworks.com/library/tpubs/pdf/router/soft1420/308602\\_1420\\_00.PDF](http://www25.nortelnetworks.com/library/tpubs/pdf/router/soft1420/308602_1420_00.PDF)
9. Internet Assigned Numbers Authority, "Internet Protocol Address Space", date unknown.  
<http://www.isi.edu/in-notes/iana/assignments/ipv4-address-space>
10. Hewes, Douglas, "I Can See you Behind Layer 2...Overcoming the difficulties of Packet Capturing on a Switched Network", September 14, 2000. <http://www.sans.org/infosecFAQ/switchednet/layer2.htm>
11. Spitzner, Lance, "Auditing Your Firewall Setup", December 12, 2000.  
<http://www.enteract.com/~lspitz/audit.html>
12. Spitzner, Lance, "Auditing Your Firewall Setup", December 12, 2000.  
<http://www.enteract.com/~lspitz/audit.html>
13. Spitzner, Lance, "Armoring Solaris", October, 22, 2000.  
<http://www.enteract.com/~lspitz/armoring.html>
14. Spitzner, Lance, "Armoring NT", April 16, 2000. <http://www.enteract.com/~lspitz/nt.html>
15. IETF Network Working Group, "RFC 1918, Address Allocation for Private Internets", February 1996.  
<http://www.cis.ohio-state.edu/cgi-bin/rfc/rfc1918.html>
16. Dittrich, David, "The DoS Project's "trinoo" distributed denial of service attack tool", October 21, 1999. <http://staff.washington.edu/dittrich/misc/trinoo.analysis.txt>
17. Stevens, W. Richard, TCP/IP Illustrated, Volume 1, The Protocols, 1994, Addison-Wesley

## Appendix: Aggregate Firewall Rules

Number	Source	Destination	Service	Action	Track
1	Internet	Customer Web Server	HTTP HTTPS	Accept	short
2	Customer Web Server	Customer Database Server Fortune Database Server	SQL	Accept	short
3	Partner_Nets Supplier_Nets	Partner/Supplier VPN Server	IPSec	Accept	short
4	Partner/Supplier VPN Server	Partner_Nets Supplier_Nets	ESP	Accept	short
5	Partner VPN IP Pool	Partner Web Server	HTTP HTTPS	Accept	short
6	Supplier VPN IP Pool	Supplier Web Server	HTTP HTTPS	Accept	short
7	Corporate Network	DMZ External Service Network Application Service Network	HTTP HTTPS IPSEC SSH Telnet	Accept	short
8	External Service Network	Corporate Network	ESP	Accept	short
9	Any	External DNS Server	DNS (53/UDP)	Accept	short
10	Management Stations	DMZ External Service Network Application Service Network	SNMP ICMP_echo _request	Accept	short
11	DMZ External Service Network Application Service Network	Management Stations	SNMP_Trap ICMP_echo _reply	Accept	short
12	DMZ External Service Network Application Service Network	Internet	NTP	Accept	short
13	Any	Any	Any	Drop	Long