Global Information Assurance Certification Paper

GIAC Level Two Firewalls, Perimeter Protection, and VPNs

Practical Assignment for Capitol SANS

December 10-15, 2000

Version 1.4

Gavin Vallance

Assignment 1: Security Architecture

The purpose of this section is to give an overview of the network we will be implementing and what technologies we will be using to secure it.

The primary requirements for the network are as follows:

- Customers need access to purchase and retrieve bulk online fortunes
- Suppliers need access to supply fortunes
- Partners need access to translate and resell fortunes

The secondary requirements for the network are as follows:

- Internet access from GIAC corporate workstations (web browsing, e-mail, etc.)
- Secure network access for corporate employees away from the office

The following diagram shows an overview of the proposed network:

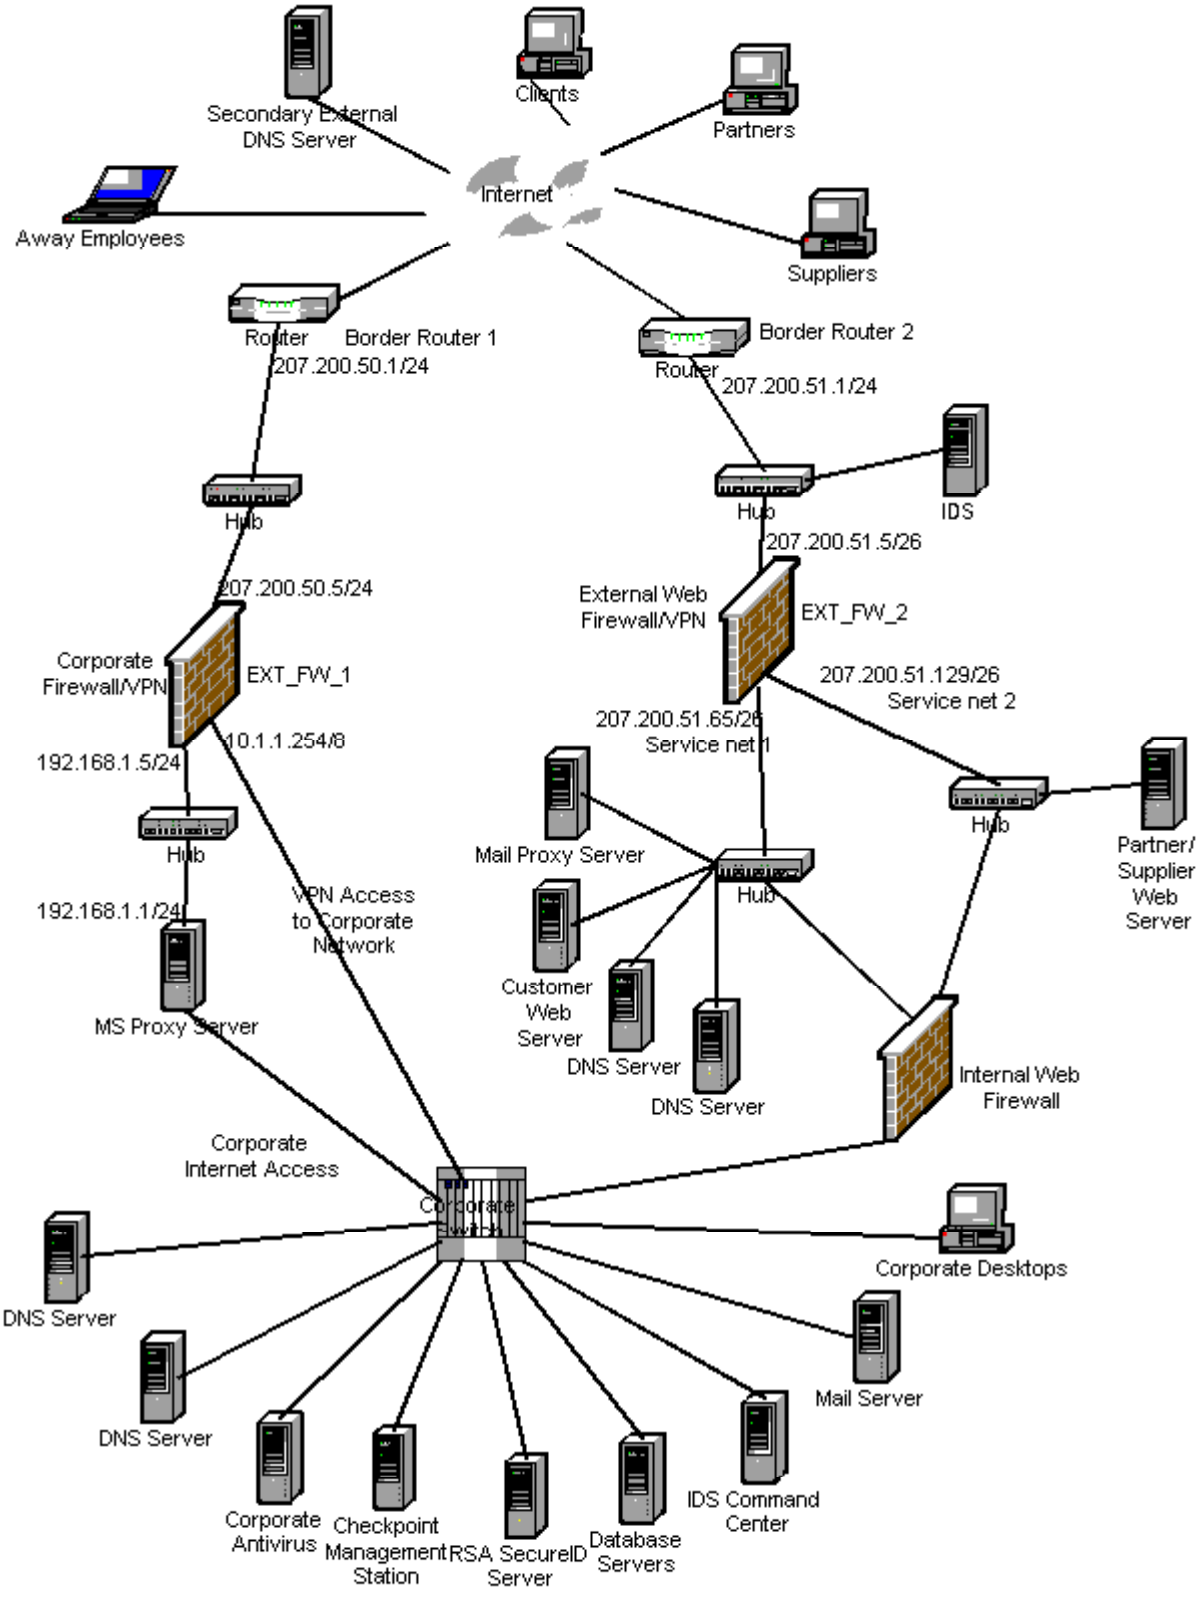

This diagram has been broken into four sections.

Section 1: The Internet

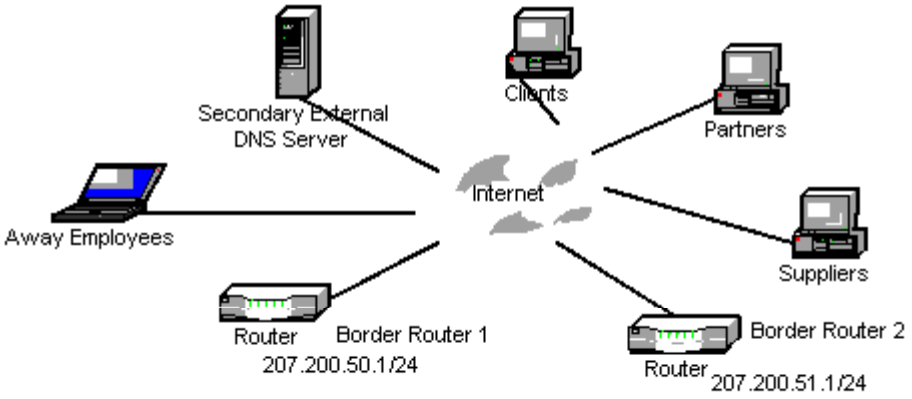

This section shows the Internet cloud with various types of computers that we want to connect to our network. The diagram shows Away Employees, Clients, Partners, Suppliers, and an External DNS server.

Away employees can be described as traveling salespeople or employees working from home that have remote access to the internal corporate network. This remote access is provided though VPN that is described in section 2.

Clients are companies or individuals that are accessing the network to browse the web site or engage in e-commerce. Clients utilize standard HTTP to browse the web site and SSL when engaging in e-commerce.

Partners are those companies or individuals who do special work for GIAC Enterprises and therefore need access to parts of the network. Suppliers are those companies or individuals who provide the fortunes and need access to parts of the network to do so.

Finally, the external DNS server is a secondary DNS server that is hosted offsite and acts as a backup for the two locally hosted DNS servers. This will allow users on the Internet to resolve the sites addresses even if the two local DNS servers fail.

Section 2: Corporate Access

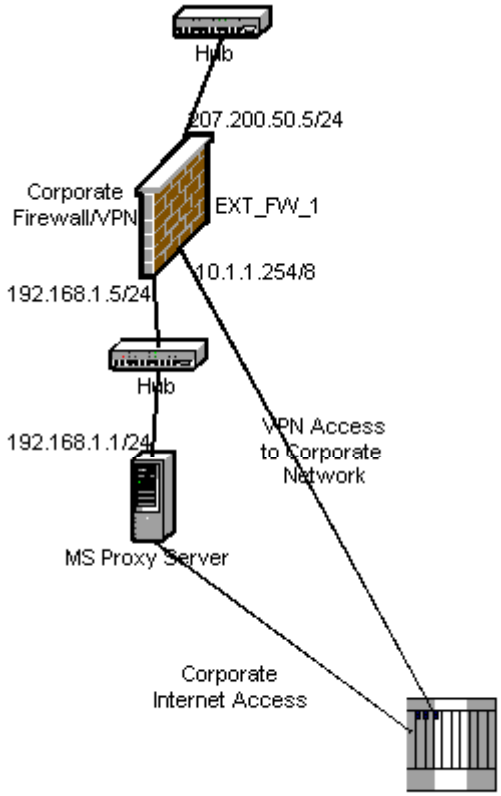

This section describes how Internet access will be given to corporate users behind the firewall and how will access the corporate network securely while at other locations.

This leg of the network consists of a dedicated fractional T3 line from the ISP to our border router, a Cisco 3600. This router acts as a basic packet filter that will eliminate some basic attacks such as an external host spoofing packets with internal addresses. All connections from here on are attached to a hub or a switch with a promiscuous monitor port. This is to easily allow the hookup of an IDS or packet sniffer as needed.

The next connection is a Nokia 440 appliance firewall running Checkpoint VPN-1. This device acts as a high performance firewall, NAT, and VPN device. The benefits of having the firewall and VPN in the same box are simplified rule base, increased security due to the VPN portion being protected by the firewall, and consolidated logging. Since the Nokia 440 is a high performance box, it should easily be able to handle the task of encrypting VPN traffic and protecting the internal network.

The firewall has two internal interfaces. One is for VPN access connected directly to the corporate network. The other interface connects with a dual-homed server running Microsoft Proxy Server 2.0. This server's second interface is connected to the corporate network. This is the path that internal corporate users will access the Internet via HTTP, FTP, or telnet. The purpose of Microsoft Proxy server is to restrict who internally can access the Internet using what protocols. Being a proxy server, it will help ensure the traffic going out on the open ports is actually what it should be. For example, this would

stop someone on the inside from connecting to an IRC server listening on port 21 on the outside via the FTP port. The log files from MS Proxy Server can be turned into some fantastic reports using a reporting product such as WebTrends Log Analyzer.

Section 3: Suppliers, Partners, and Clients

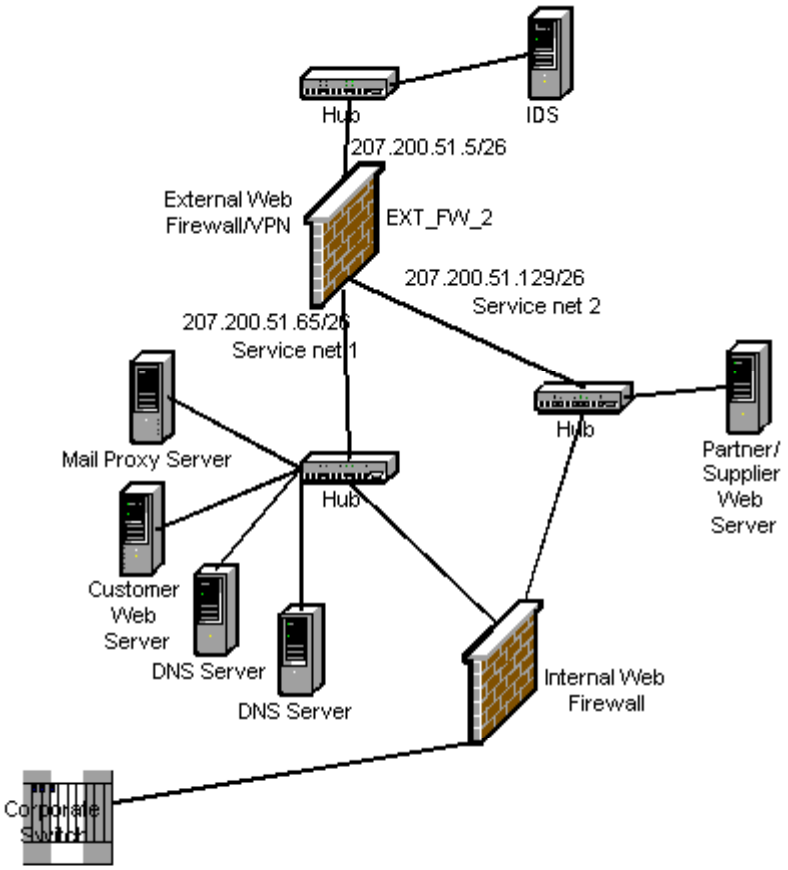

Section three is the public part of GIAC Enterprises' network. This section is fed by its own T3. The purpose of having separate Internet connections for sections 2 & 3 is to prevent them from interfering with each other and to separate business functions. As in section 2, the T3 feeds into a separate Cisco 3600, which is also acting as a basic packet filter. The router feeds into a hub that connects to an IDS and the external firewall. The IDS will monitor all traffic between the router and firewall for suspicious activity. The IDS used in this case is SNORT 1.7 running on Red Hat 7.

The external firewall is again a Nokia 440 running Checkpoint VPN-1. The purpose of the external firewall is to protect the servers in the service network from unnecessary or malicious traffic. In this case, the firewall will accept VPN connections from suppliers and partners, HTTP and SSL traffic from clients, SMTP mail in and out of the corporation, and primary DNS requests for the domain.

The firewall has two internal interfaces to the service network. The service network on the left is where the mail proxy, external DNS servers, and the public web server reside. SMTP mail coming into GIAC Enterprises does not get directly sent to the corporate mail server. Instead it is delivered first to a mail proxy server in this service network. The purpose of this is to protect the internal mail server from direct attacks and to filter and scan incoming and outgoing e-mail for viruses and/or other types of dangerous content/attachments. The two DNS servers are the authoritative servers for GIAC Enterprises domain. They allow zone transfers only to the external DNS server that acts as a secondary if these two should fail. The DNS servers also act as primary timeservers. They synchronize with an external NTP server such as one hosted by the US Navy. All other servers in the service networks and corporate network synchronize off these two servers to provide a global time consistency. This is important when tracing events across multiple logs or servers. Finally we have the web server that acts as the main web site for customers that provides company and product information as well as a place to purchase and receive fortunes via e-commerce. Regular browsing traffic is done via HTTP and e-commerce is secured with 128-bit SSL encryption.

The service network on the right will only accept VPN traffic authenticated by the firewall for suppliers and clients. This leads to the Partner/Supplier Web Server. This server is where partners and suppliers will connect to conduct business via HTTP and FTP. We do not have concerns about the clear-text HTTP and FTP sessions because they are encrypted via the VPN connection.

All servers in these two zones are kept to current patch and hot fix levels for the OS and applications. The operating systems are hardened and have unnecessary services turned off.

Both service networks connect to an internal firewall. This firewall is Symantec's Raptor 6.5 running on Windows NT 4.0 SP6a+current hot fixes. The job of this firewall is to protect the corporate network form any server in the service network that might be compromised and to protect the service network from unauthorized corporate users. This firewall is purposely different from the external firewall two aspects: platform (Nokia vs. NT) and type (state full inspection vs. application proxy). This difference increases our security by preventing a hacker who has used an exploit on the external firewall from using the same exploit on the internal firewall.

Section 4: The Corporate Network

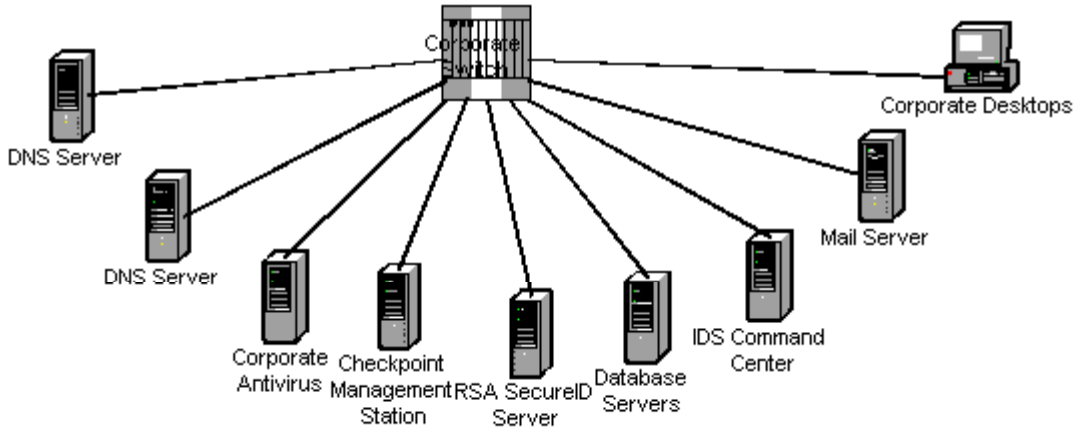

The corporate network consists of many servers and clients. I have shown only the ones that have relevance to the overall security plan. These include internal DNS servers, a corporate anti-virus platform, the backend database servers, and IDS command center, the mail server, the Checkpoint VPN-1 Management Console, and the RSA SecureID Authentication Server. The DNS servers will resolve internal and external addresses for internal hosts through the corporate firewall. The corporate anti-virus platform will get external anti-virus updates through the corporate firewall and scan internal servers and update the mail proxy definitions. The backend database servers contain the data for the two web servers in the service networks. Their communication is controlled through a SQL proxy in the Raptor firewall. The IDS command center is the central console for all IDS modules that are installed on the service network servers. This console receives alerts and data from these modules and notifies the appropriate staff when needed. The Checkpoint VPN-1 Management Console is where the rule base for the two Checkpoint firewalls are written and saved and also the central repository for their log files. The RSA SecureID Authentication Server (RSA ACE server 4.0) provides the authentication for the two-factor authentication keys that are used for VPN access. Corporate users, partners, and suppliers are given a hardware or software token that generates a number that when combined with the users PIN allows strong authentication when establishing a VPN session. The VPN/RSA solution will be discussed more in assignment 2.

Assignment 2: Security Policy

Now that the network has been diagramed a security policy must be designed for the border routers and firewalls. In this section we will detail the security policy of the two border routers and the two external firewalls, including VPN access.

Border Router 1:

This router provides access to the corporate browsing and VPN leg of the network.

Cisco 3600

```
1 service password encryption

2 no service tcp-small-servers
3 no service udp-small-servers
4 no service finger

5 no ip bootp server
6 no ip http server

7 ip access-group 110 in
8 ip access-group 120 out

9 access-list 110 deny ip 10.0.0.0 0.255.255.255 any log
10 access-list 110 deny ip 172.16.0.0 0.15.255.255 any log
11 access-list 110 deny ip 192.168.0.0 0.0.255.255 any log
12 access-list 110 deny ip 127.0.0.0 0.255.255.255 any log
13 access-list 110 deny ip 207.200.50.0 0.0.0.255 any log
14 access-list 110 permit any

15 access-list 120 permit ip 202.200.50.0 0.0.0.255 any
16 access-list 120 deny ip any any log

17 no ip directed-broadcast
18 no ip source route
19 no ip unreachables
20 no snmp

21 logging 10.1.1.10

22 banner / Warning! Authorized Users Only! /
```


This table shows a basic design for the border router. Some modifications may have to be made such as *no ip direct-broadcast* must be assigned to the external interface. The line numbers would not be included on an actual router but have been added for reference.

1. This line encrypts the router password instead of displaying it clear-text in the configuration file. This encryption is reversible and is only meant to make it difficult to discover the password by inspection
2. This disables access to small port services such as echo, discard, chargen, and

daytime. These services are rarely used anymore and are unnecessary in most cases.
3. This disables access to small port services such as echo, discard, chargen, and daytime.
4. This disables access to the finger utility, which can give away information you would rather keep unknown!
5. Disables the bootp server
6. Disables the HTTP server
7. Defines the access list number for incoming traffic to be 110
8. Defines the access list number for outgoing traffic to be 120
9. Prevents RFC-1918 source addressed packets from entering from the Internet. Packets coming from these addresses could be from a misconfigured host or somebody trying to do something naughty!
10. Prevents RFC-1918 source addressed packets from entering from the Internet.
11. Prevents RFC-1918 source addressed packets from entering from the Internet.
12. Prevents the loop-back address from being used as a source address.
13. Prevents external packets from being spoofed with our registered internal addresses
14. Allows anything else in
15. Allow packets with the correct registered internal addresses go out.
16. Deny everything else. This prevents internal hosts from sending out spoofed packets. This makes us a good net neighbor.
17. Assign to external interface and prevents against smurf attacks
18. Prevents packets from using loose source routing, which can be used to attempt to slip packets past our defenses.
19. Prevents the router from sending ICMP error messages that can give out network information
20. Disables SNMP
21. Sets the logging server
22. Gives a warning to all users attempting to log into the router. This is important to show that you don't want just anyone to have access to the router and is necessary if you're going to prosecute!

Border Router 2:

This router provides access to the service network and partner/supplier VPN leg of the network.

Cisco 3600

```
13 access-list 110 deny ip 207.200.51.0 0.0.0.255 any log
15 access-list 120 permit ip 202.200.51.0 0.0.0.255 any
```

This router will have the same configuration as border router 1 with the exception of rules 13 and 15. Rule 13 prevents outside users from spoofing packets with valid internal

addresses. Rule 15 prevents internal users from sending out spoofed packets. Therefore these rules are different on this router because the internal addresses are different.

External Firewall 1:

This firewall protects the corporate browsing and VPN leg of the network.

Checkpoint VPN-1

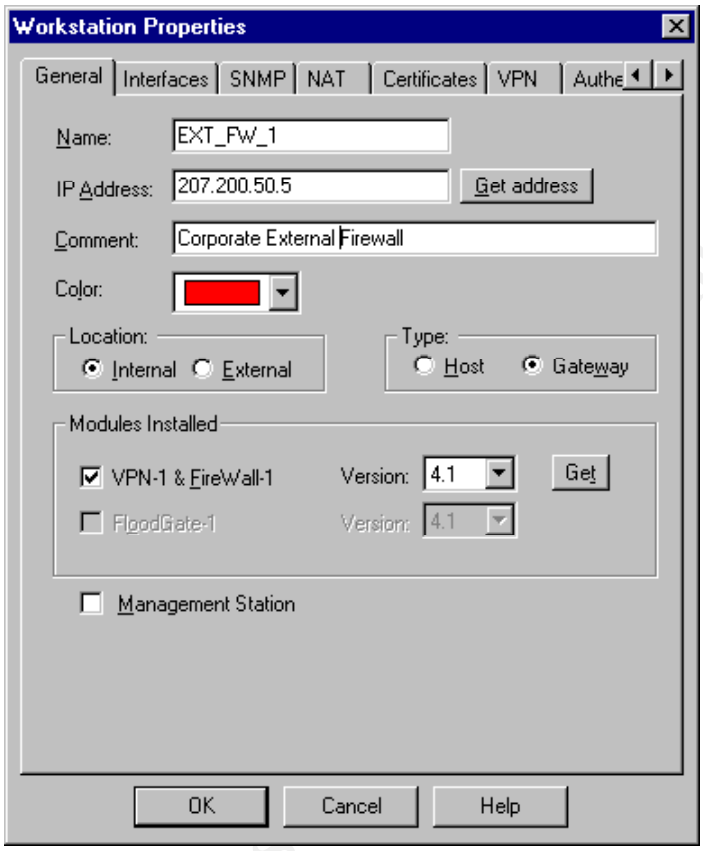

No.	Source	Destination	Service	Action	Track	Install On	Time	Comment
1	ANY	EXT_FW_1	ANY	DROP	LONG	EXT_FW_1	ANY	
2	PROXY_NET	ANY	HTTP HTTPS TELNET FTP	ALLOW	LONG	EXT_FW_1	ANY	
3	DNS_INTERNAL	ANY	DNS	ALLOW	LONG	EXT_FW_1	ANY	

4	CORP_USERS@ANY	CORP	ANY	CLIENT ENCRYPT	LONG	EXT_FW_1	ANY	
5	CORP PROXY_NET	ANY	NBT	DROP		EXT_FW_1	ANY	
6	ANY	ANY	ANY	DROP	LONG	EXT_FW_1	ANY	

This is the rule set for the first VPN-1 firewall. The explanations for the rules are as follows:

1. This rule states that any packet directly addressed to the firewall should be dropped and logged. Packets addressed directly to the firewall could be someone trying to attack the firewall or scan for vulnerabilities. We drop instead of reject because it makes the firewall more "silent" but don't be fooled into thinking this makes it invisible. They will get less response this way.
2. This rule allows traffic from PROXY_NET, which is the private network that the proxy server resides on, to go out on allowed protocols. Traffic from corporate users go to the proxy server which permits or denies it. If permitted, the traffic is forwarded to the firewall, but as a translated address of 192.168.1.1. The firewall then translates the address to a proper registered IP address and sends it on its way.
3. This rule allows the internal DNS servers to make external DNS queries to resolve address requests from internal hosts.
4. This is the VPN rule. Since VPN is integrated into the firewall, all we need to do is add a rule specifying what users have access to what services. Here we say that CORP_USERS@ANY can access anything on the internal corporate network. CORP_USERS is a defined group of users who are allowed access. @ANY means allowed users from any host. This traffic is sent to a separate interface that is connected directly to the corporate network. The CLIENT ENCRYPT states that the user must be using Checkpoint SecuRemote or SecureClient to establish an encrypted VPN tunnel. This will be described after the rule explanations.
5. This rule drops and does NOT log any NetBIOS packets that either internal interface sees. This stops the logs from filling up with the constant NetBIOS broadcasts that Microsoft machines frequently do. These are generally benign and not worth logging. This does not disable logging of NetBIOS traffic seen from the external interface. This is because such traffic could be someone running a NetBIOS scan or vulnerability scan and should be logged.
6. The last rule is the generic deny all other traffic and log it.

The order of these rules is important since they are checked in descending order until a rule matches. The first rule should always be deny packets addressed directly to the firewall. The only time you may want to put a rule before it is if you wish to enable ICMP echo requests and replies (ping) to one or more of the firewalls interfaces. VPN may

access the firewall's external interface directly but is not denied by rule one because there are implied rules that allow clients to authenticate with the firewall. These are considered rule 0. In this limited rule set the order of rules 2-5 are not that important since none of them supersede another. I generally group the rules by whether they are coming in or going out. The last rules I have are usually the ones that do not log constant broadcast traffic.

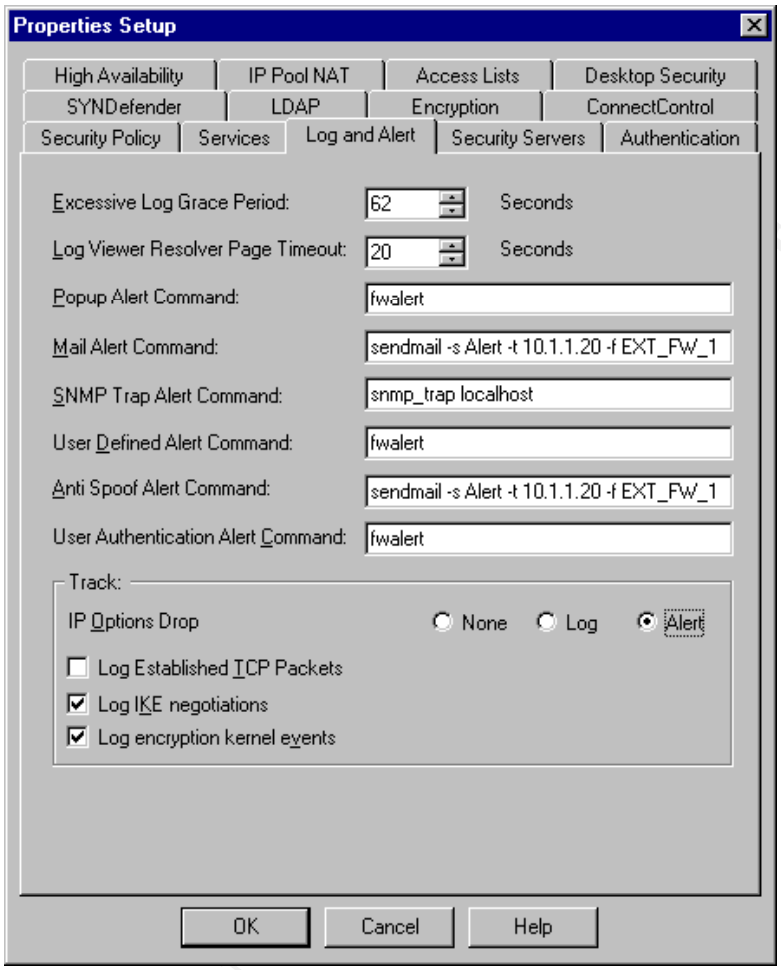

All rules (except rule 5) are logged in long form, which records all information that Checkpoint can about the packet. The tracking function also supports different alert methods such as SMTP mail, SNMP trap, or run a program. This is useful if you want to be made aware by pager or email if a certain type of traffic occurs. You can also restrict rules by time by creating a time object. This is useful if you need to open a window for backups using SSH for only 2 hours from one network to another. This would reduce your vulnerability time. The picture above shows the alert settings for this firewall.

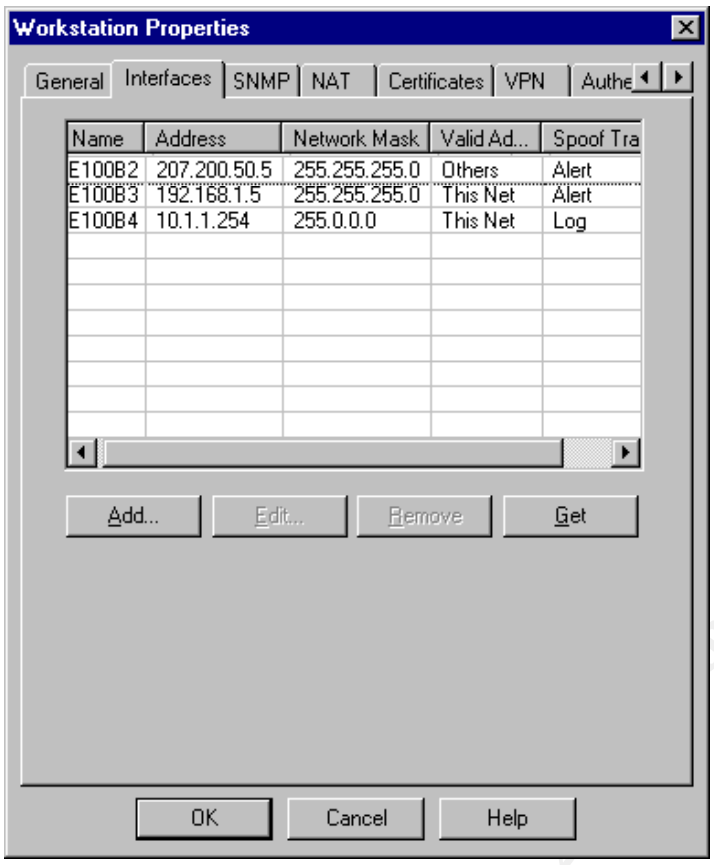

Other settings not shown are anti-spoof rules. These are done though the configuration of the Firewall object. Each interface is assigned a network of allowed addresses. The internal interfaces are set to "This Net" meaning allow traffic from this network that have source address of that subnet. The external interface is listed as "Others." This means allow packets with source address that are not members of any of the internal interfaces subnets. Each interface also lets you set whether you want to be alerted when a spoofed packet is detected or just log it. I generally set the external interface to alert and internal to log. This is because on a large network you can have some hosts sending out packets that are not addressed or have broadcast addresses. These are detected as spoofed addresses. Anti-spoofing on the firewalls is almost redundant since the border router should block spoofed addresses. It never hurts to have an extra layer of protection! The picture above shows the network interfaces for this firewall.

The VPN component of the Checkpoint firewall is an added option. This option comes with a choice of encryption and authentication. For this network I have chosen to use IKE key exchange with 3DES encryption and RSA Secure Authentication. IKE is a standard for managing keys between two hosts using IPSec. The RSA Secure Authentication scheme requires the RSA ACE server. This server is the authentication server that will allow users in based on the numbers they enter. The external user has either a fob or a software key, known as SecureID, that can run on a laptop or palm that generates new numbers on a regular interval. When a user attempts a VPN connection,

they have to submit the number seen on the generator and a pin. This will authenticate them with something they know and something they have. This is known as strong authentication. From there the session will be encrypted.

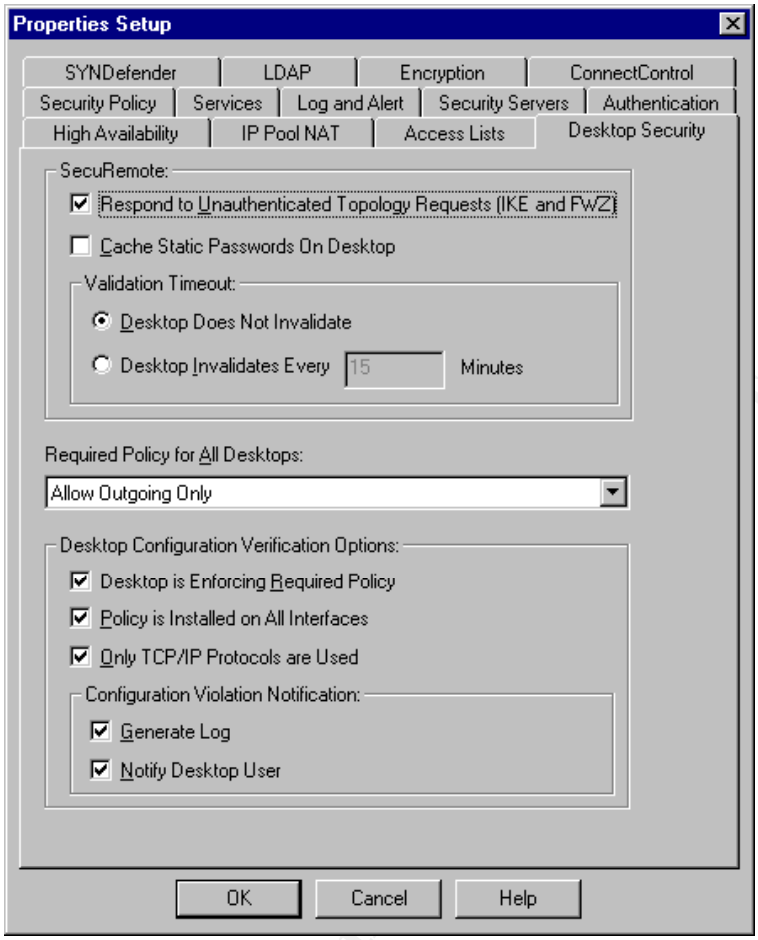

The two VPN clients offered by Checkpoint are SecuRemote and SecureClient. They both serve the same purpose of establishing an encrypted tunnel over the Internet or though a network but they have two main differences. The first difference is that SecuRemote is free and SecureClient is not. The second difference is that Secure Client also acts as a personal firewall when the VPN is in session. With SecureClient, the administrator can restrict what traffic can access the tunnel. For example, a workstation at a partner is on her local area network. She dials up to GIAC Enterprises and establishes a VPN tunnel. Now she is on two networks. If someone gains access to her machine though the LAN, he would be able to access the VPN tunnel, compromising the GIAC internal network. In this case it makes sense to make sure all VPN users with unlimited or sensitive network access use SecureClient to maintain better security. The picture above shows that the security policy can be pushed to the remote desktop.

External Firewall 2:

This firewall protects the service network leg of the network.

Checkpoint VPN-1

No.	Source	Destination	Service	Action	Track	Install On	Time	Comment
1	ANY	EXT_FW_2	ANY	DROP	LONG	EXT_FW_2	ANY	
2	ANY	CUST_WEB_SERV	HTTP HTTPS	ALLOW	LONG	EXT_FW_2	ANY	
3	ANY	MAIL_PROXY	SMTP	ALLOW	LONG	EXT_FW_2	ANY	
4	MAIL_PROXY	ANY	SMTP	ALLOW	LONG	EXT_FW_2	ANY	
5	ANY	DNS1 DNS2	DNS	ALLOW	LONG	EXT_FW_2	ANY	
6	DNS1 DNS2	ANY	DNS NTP-UDP	ALLOW	LONG	EXT_FW_2	ANY	
7	PART_SUPP@ANY	PART_SUPP_WEB	HTTP FTP	CLIENT ENCRYPT	LONG	EXT_FW_2	ANY	
8	PART_SUPP_WEB	DNS1 DNS2	NTP-UDP	ALLOW	LONG	EXT_FW_2	ANY	
9	SVC_NET1 SVC_NET2	ANY	NBT	DROP		EXT_FW_2	ANY	
10	ANY	ANY	ANY	DROP	LONG	EXT_FW_2	ANY	

This is the rule set for the second VPN-1 firewall. The explanations for the rules are as follows:

1. This rule states that any packet directly addressed to the firewall should be dropped and logged just as it was in the first firewall.
2. This rule allows HTTP and HTTPS to come from anywhere to the client web server. This provides generic browsing and encrypted SSL sessions.
3. This allows anyone to start an SMTP session with the mail server. The primary use of this will be when outside servers are sending SMTP mail to the GIAC domain. The MX record of the DNS will reference the mail proxy server as the

primary mail server. The mail proxy will not allow mail relays, VRFYs, or EXPNs. These functions are used by spammers to relay mail to other domains or enumerate your own domain for abuse. If the mail proxy decides the message is valid to forward to the internal mail server it will do so through another mail proxy in the Raptor firewall.

4. This rule allows mail to be sent from the mail proxy to the outside. The mail proxy will receive messages from the internal mail server to be routed to the outside via the Raptor firewall. It will then resolve the Internet mail server it must connect to and send it.
5. This rule allows anyone to make a DNS query to either DNS server in service net 1. This includes any machine on the Internet or in service net 2. The DNS servers will be set to only allow zone transfers to one trusted external host that acts as a secondary DNS server. Allowing zone transfers to anyone allows them to see the entire DNS record, exposing all of your hosts. In this case, the only hosts in the DNS server are publicly available so it would not be as tragic. It is good practice anyway to restrict them. The DNS servers will also be set up to only do recursive queries for hosts in either service network. This prevents outside users from using the DNS servers as resolvers for any host on the Internet.
6. This rule allows the DNS servers to make DNS queries to other DNS servers. This is necessary because they will resolve addresses for the mail proxy and any other server that resides in either service network. Only hosts on the service network are allowed to make recursive queries. This rule also allows the DNS servers to make NTP-UDP queries to an external NTP server, such as one hosted by the US Navy. All other servers in the service networks will synchronize time off these machines to maintain a consistent time. This makes log tracing across multiple machines easier.
7. This rule sets up VPN access for suppliers and partners. Following the VPN example from the first firewall we see that users in PART_SUPP group have VPN access from any machine. This VPN rule allows only HTTP and FTP access to the partner/supplier web server, which resides on the second internal interface of the firewall. Traffic over this connection will be encrypted and authenticated via the RSA ACE server. Because of this, clear-text FTP and HTTP are acceptable. All partners and suppliers will be provided with a SecureID fob or software and SecuRemote. The reason for not using SecureClient is that partners and suppliers have only limited access to one server, HTTP and FTP. Even if the remote client is compromised, it is not a great threat. Also regulating how a partner or suppliers computer participates on their own network for this simple access seems like overkill.
8. This rule allows the partner/supplier web server in the second service network synchronize time with the two DNS servers in the first service network.
9. This rule silently drops all NetBIOS packets received on the two internal interfaces of the firewall. This rule follows the same reasoning as rule 5 on External firewall 1.
10. The last rule is the generic deny all other traffic and log it

Again as in External Firewall 1 all rules are set to long logging (except 9) and anti-spoofing is set.

Assignment 3: Audit you Security Architecture

Once the network has been implemented and functioning the next step is to audit your environment. There are several ways to audit you security setup, you can do it yourself or you can hire a company to do it for you. I believe in a balanced approach. Since knowing about vulnerabilities can help you defend against them, you should remain current on basic penetration techniques. There are many free utilities such as nmap, saint, SARA, and others. On the other hand, many penetrations take some level of skill. If you are not an experienced hacker then bringing in a consultant to give a full security audit is probably a good idea.

To start we can try to ping each of the firewall's interfaces and the hosts in the service network. The ping results should be something like this:

```
C:\>ping 207.200.50.1

Request timed out.
Request timed out.
Request timed out.
Request timed out.
```

This shows that the interface to properly not responding to ICMP echo requests. Of course this could also mean that the firewall is not functioning but you can test that by making sure you can access allowed services such as HTTP to the web server.

A better tool than ping is nmap, by insecure.org. This tool scans the ports on a machine for responsiveness. You should run this scan against all of your machines on their local subnet (not firewalled) to see what ports are open on the machine and decide if it can be hardened more. Once this is done, try the scan again from outside the firewall. The results should be different. The only ports that should be open on the server from the outside are those that are explicitly allowed by the firewall.

Here is an example of an nmap scan:

```
Nmap -sS 207.200.51.70

Starting nmap V. 2.54BETA7 ( www.insecure.org/nmap/ )
Interesting ports on  (207.200.51.70):
(The 1533 ports scanned but not shown below are in state:
closed)
Port       State          Service
25/tcp     open           smtp
```

Nmap run completed − 1 IP address (1 host up) scanned in 11 seconds

After running the scans from a host outside of the firewall, check your firewall logs to see what it looks like from the firewall side. This will give you an indication on what a basic scan would look like. You can test IP spoofing easily by plugging a laptop with an internal address outside the border router. Then try and ping or browse a host behind the firewall. Packets coming from your computer should have an internal source address. The external interface of the router should see that internally addressed packets are attempting to come in the external interface, this should set off a spoof alert and deny the packets. It is also important to attack from inside your network. Many hacking attempts come from the inside!

The ping and nmap scans will test your firewall's rule set but does not check for application errors such as allowing SMTP relaying and weak or default passwords. You can test your SMTP relay rules with a Windows laptop outside the firewall. Use a simple e-mail program such as Outlook Express and send a e-mail to an external account, such as hotmail) using the mail proxy in the service network as your SMTP server. The e-mail client should connect via SMTP to the mail proxy but when the proxy realizes that you are trying to send mail to a host outside of its domain, it will reject the message or just not route it. If it accepts the message and sends it to the external account then you are vulnerable to spammers! To test for default passwords you can plug your laptop into the various hubs/switches and try and log into various devices with their default passwords. Many routers and manageable hubs/switches have default passwords.

Depending on your skill level, you can do more or less, but in the end you should always have someone else check your work. A good way to do this is to hire an outside consulting company to review your setup and attempt to penetrate your defenses. Many consulting companies will come in and look over what you have done and make recommendations on how to improve your setup. They also may offer to do a blind penetration test. This means that they will try and breach your security with little or no knowledge about your network. I believe both can be helpful, but if the same company does both, make sure the group that evaluates your setup does not also do the penetration test (otherwise it wouldn't be blind!).

Companies that do this work can be found on the web or you can get recommendations from people you know. Once you have some companies selected make sure you get references that you can call to confirm the quality of their work. This can be difficult because many companies do not want to disclose what the do regarding security.

A good penetration test will test all that you have already done and much more. They will try and exploit vulnerabilities specific to the applications you are using. For example, they might try and crash your Microsoft IIS web server using the vulnerability of the week. If you have not applied the most recent service pack then you have just lost a server. Sometimes the penetrators will specify a goal of placing a file in a certain location

on one of your servers. Another method of testing your security is a social method. They may call up your users and pretend to be IS personnel and request their password. This makes cracking that much easier.

In the end you will have a detailed report of your network and its security status. Hopefully there will be nothing wrong but likely they will make some helpful recommendations that will secure your network further.

Assignment 4: Design Under Fire

The network that I am going to design attacks for was submitted by Alexander Usenko. The URL for his submission is:
ww.sans.org/y2k/practical/Alexander_Usenko_GCFW.doc

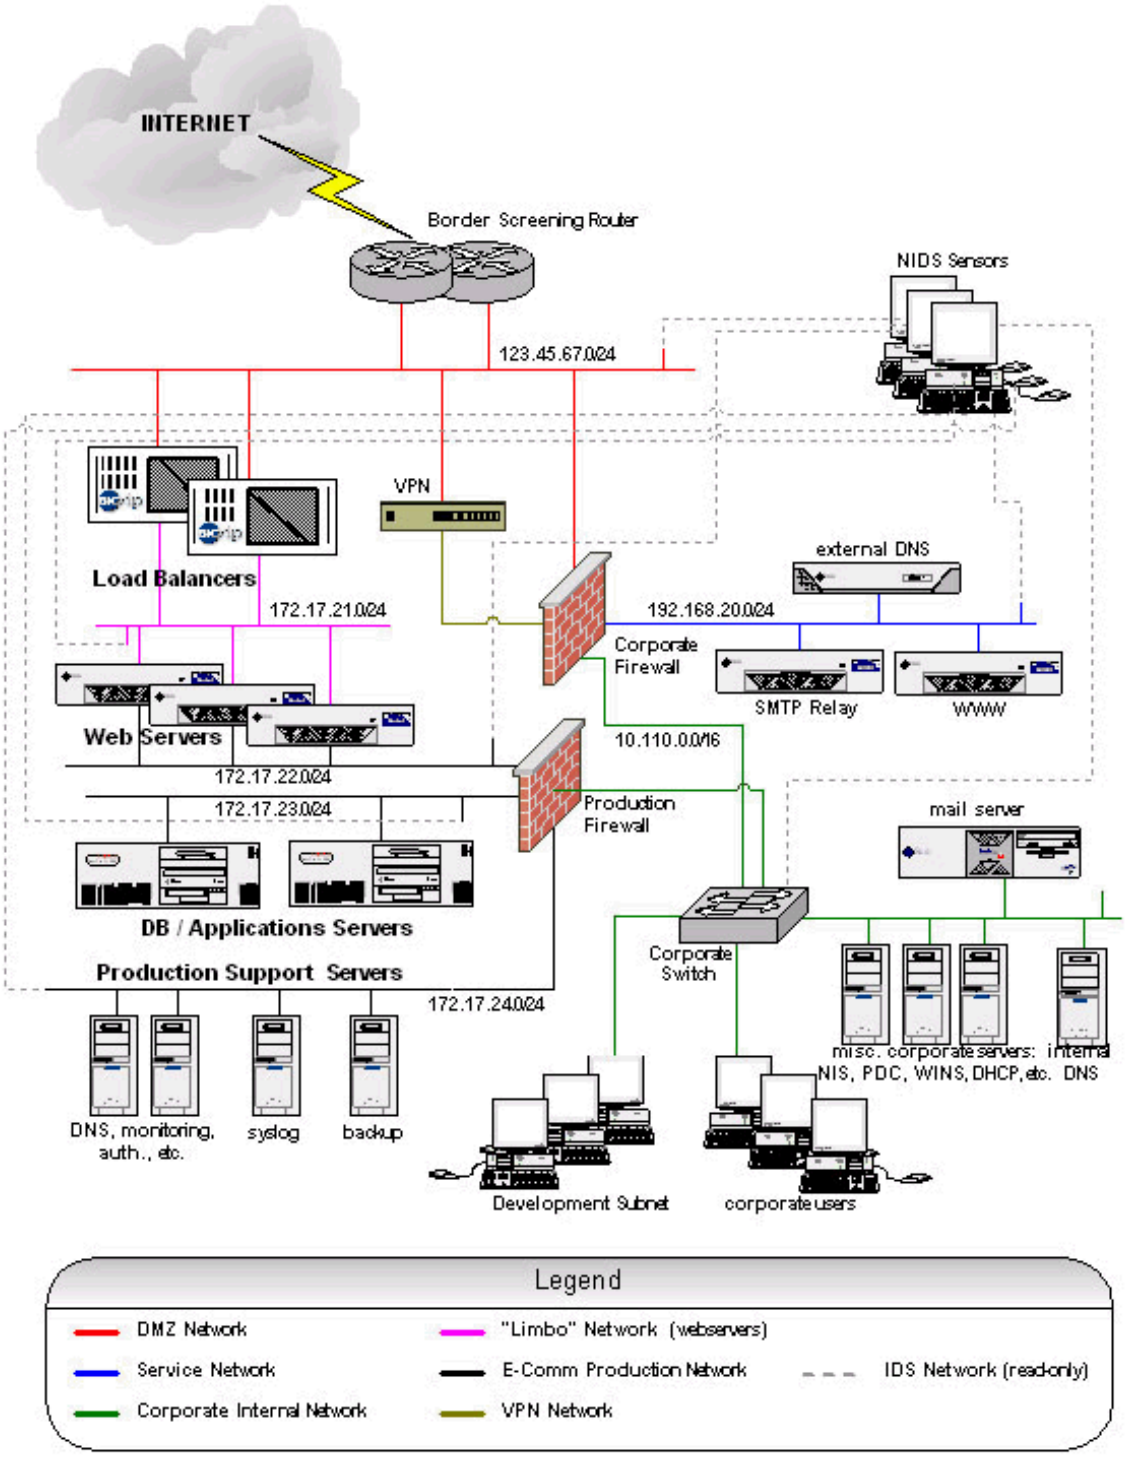

Attack 1 & 2:

The first attacks are directed against the firewall itself. In this case, the external firewall is Checkpoint Firewall-1 version 4.1. Since this attack requires multiple attacking hosts and is against a vulnerability on the firewall itself I combined parts 1 and 2.

This attack involves the SYN Defender Gateway on the firewall. When this option is set to protect internal hosts from a SYN attack it will hold queue the SYN requests and check to see if the internal host is accepting traffic on that port. If the firewall is overwhelmed with a SYN flood from many hosts the queue will fill up and result in a firewall denial of service attack. To protect against this we can set the SYN protection to a passive gateway that passes the SYNs on and will not slow the firewall during a SYN flood.

Attack 3:

The last attack is to compromise an internal system. I would pick one of the web servers in the 172.17.21.0/24 subnet. The reason for choosing this server is because web servers can easily exploitable through the application and this one is only protected by a screening border router. My goal would be to bring the server down or change the content.

First I would scan the router and host with nmap to find any ports that may be open. If the border router is only blocking spoofed addresses and nothing else, then all active ports on the web server would be open for attack. For example, if the web server is an NT server running IIS, there could be the NetBIOS ports, FTP, Gopher, and the IIS administrative service open. These ports all have know vulnerabilities (most fixed by patches, but not everyone keeps up-to-date!).

I would attempt to discover the OS and web server software through the HTTP responses that frequently respond with the names and version numbers. This will allow me to lookup specific exploits for the web server software or OS. I would attempt each exploit in the hopes that I will find an unpatched hole. If this is unsuccessful, I could try to discover exploits of any other software on the machine such as Cold Fusion, an application that is notorious for not being properly secured.

With enough time and knowledge of the OS and software and the luck of a behind-the-times administrator, a hole will likely be found and the server will fall in one way or another.