



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

GIAC Firewalls, Perimeter Protection, and VPNs

Practical Assignment

Capitol SANS December 10 – 15, 2000

Carrie N. Chalmers

© SANS Institute 2000 - 2005, Author retains full rights.

Assignment 1: Security Architecture

Define a security architecture for GIAC Enterprises, a new Internet startup that expects to earn \$200 million per year in online sales of fortune cookie sayings. 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).

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.

For the purposes of this practical, I have divided the network architecture diagram for GIAC Enterprises into two networks with distinctly different functions connected by a VPN: the Production Network (untrusted) and the Corporate Network (trusted).

The Production Network is composed of the bank of Apache web servers accessed by customers via the Internet. Customers can obtain product descriptions and place orders that are confirmed through an email sent automatically back the customer. This network also consists of network management boxes, databases and backup server(s).

Filtering Router- The first security device a customer's traffic will traverse is the packet filtering router. This Cisco 2604 series router provides initial traffic filtering for inbound traffic, but also is used to prevent specific 'internal only' traffic from bleeding out to the Internet. In addition, it aids in stopping certain traffic (like pings) from reaching the firewall and degrading the firewalls performance. A VPN is run from router #2, also a Cisco 2604 to the corporate network –located on another floor of the building. (See Assignment 2: Security Policy)

Firewall #1- For this practical, I have chosen to use a Sidewinder Firewall version 5.0. This firewall is used to separate the DMZ (note that an IDS could be placed within the DMZ) from the web server bank, IDS, and smart relay. Since the web servers are a critical part of the business, the Sidewinder provides additional security that can't be accomplished using a simple packet filtering router. Although for simplicity these diagrams do not show it, two firewalls could be run in a redundant architecture. This is advisable for the production firewall in particular. Network address translation (NAT) is run on firewall #1 but is disabled on firewall #2. (See Assignment 2: Security Policy)

Web Server- These Apache web servers are accessed on port 80 for normal traffic such as product description. HTTPS using SSL (128-bit encryption) is used for actual customer transactions that may include sensitive data. This customer information, maintained on Oracle databases, is sent encrypted via a second Sidewinder firewall from the web servers. A Cisco IP Redirector will be used for load balancing between the servers (Not shown). Like the routers and firewalls, the web servers should be locked down as much as possible with all non-used services removed. The entire web server bank should be backed up on a regular basis and when ever changes are made. If the web servers were hacked, backups are a must. They allow you to compare to find malicious changes and lessen recovery time.

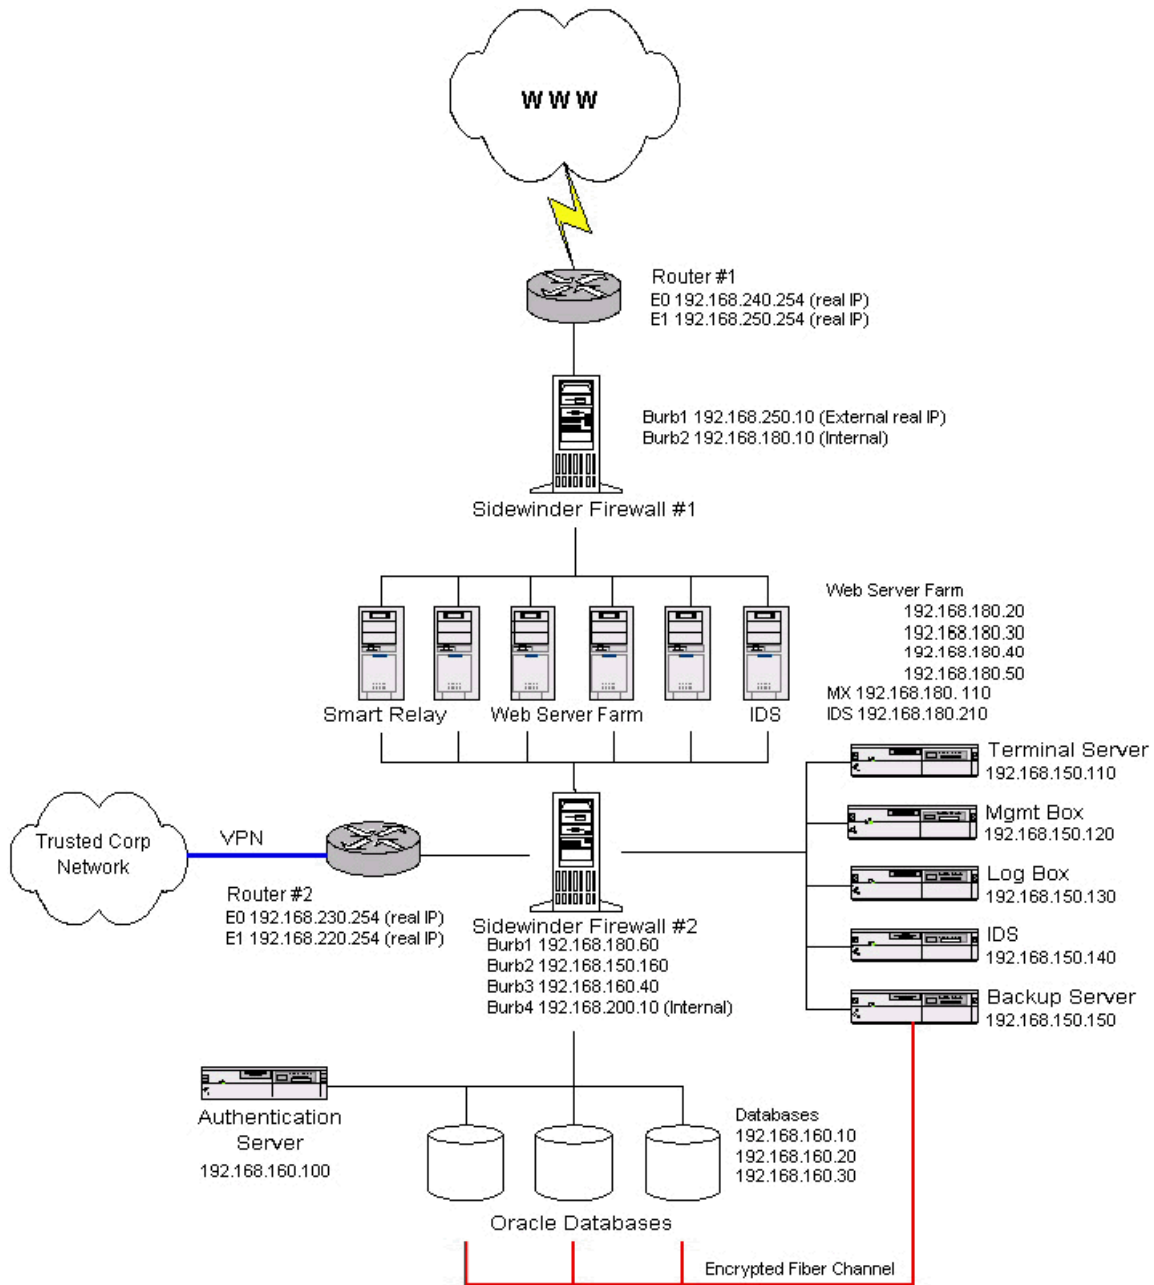
Smart Relay- The smart relay server provides outgoing only sendmail services to confirm orders placed via the Internet. This helps in cutting down the number of possible exploits that can be used to compromise the web servers or any other box on the network since incoming mail can be blocked. It is also configured to lie about its sendmail version. (See [DNS and Sendmail](#) by Hal

Pomeranz, SANS 2000)

DNS- DNS is run on the sidewinder firewalls #1 and #3 in a split configuration with network address translation. (DNS could also be run separate from the firewall, however I am comfortable enough with Sidewinder's design to run it on the firewall itself. See the description of type enforcement in Assignment #2 firewalls.) The external exposed name servers are configured to give away as little information as possible and not to do any recursive queries. They are also configured, along with the internal name servers to give generic answers to any BIND version queries and allow zone transfers from certain IP addresses only. I will not get into how to configure BIND in this practical, but it is important to understand how DNS works and how it is related to the security of a network. DNS and Bind by O'Reilly publishing is an excellent source of information.

GIAC Enterprises Production Network

© SANS Institute 2000 - 2005, Author retains full rights

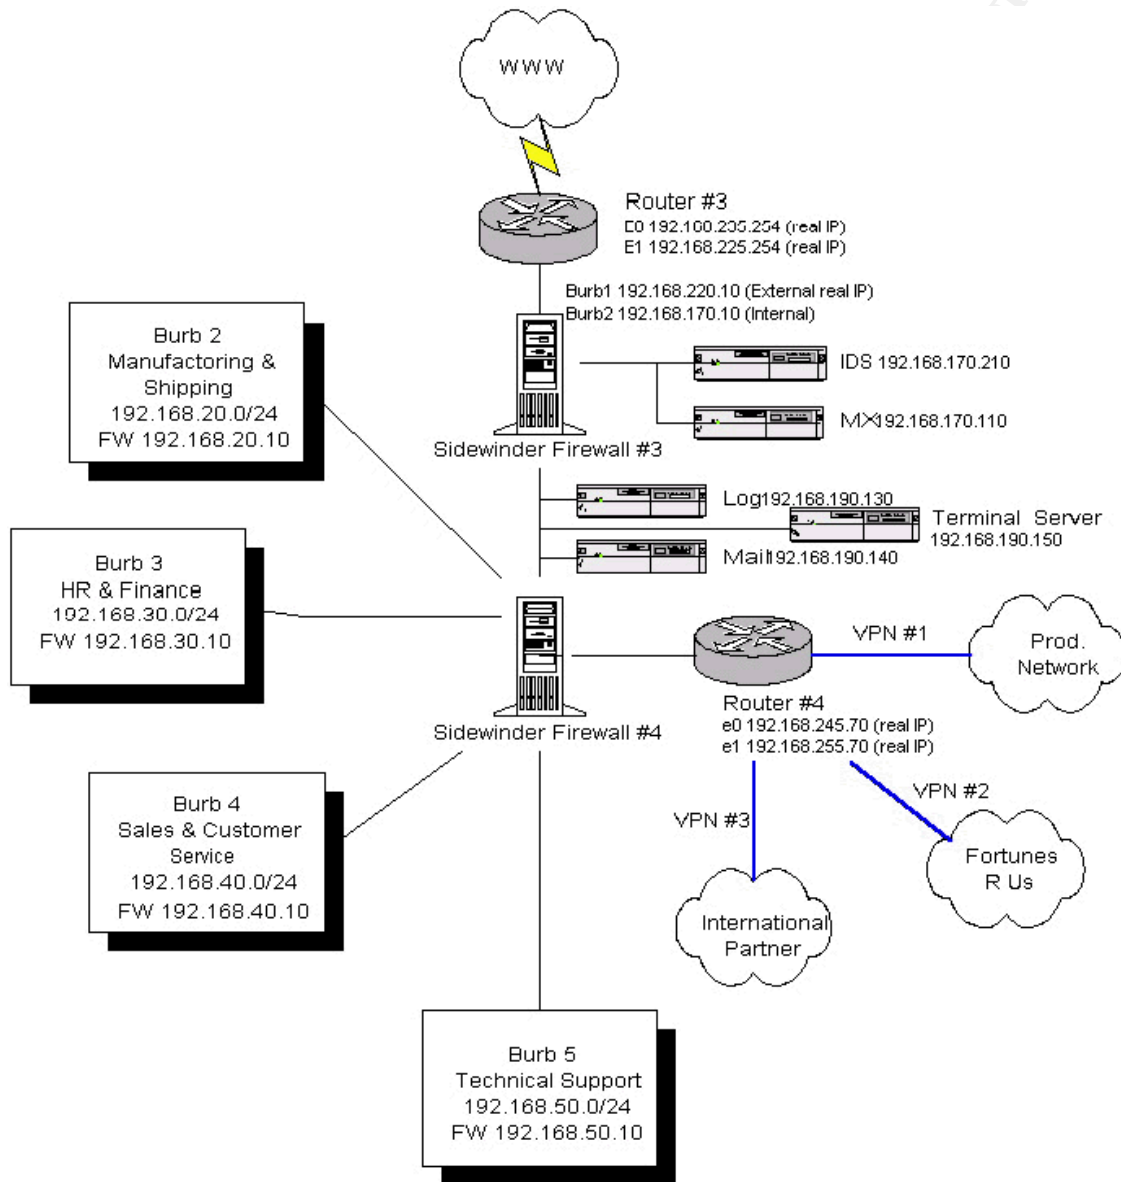


Screened Network- The screened network off burb 2 on the Firewall #2, contains many of the network management devices. The terminal server accessed via ssh by approved network personnel is used to access the routers, switches, etc. via a serial connection. With this configuration, telnet does not need to be permitted on the routers. The management box provides the same function for the databases. A syslog server collects log information from all devices that can be configured to perform this function including the routers and any intrusion detection systems on the network. The backup server is connected to the databases by an encrypted fiber channel bypassing the firewall for efficiency. Burb 2 on the sidewinder contains the Oracle databases and an authentication server for the web servers. This burb or another burb off of firewall #2 could also contain any system that needed to be accessed by multiple departments from within the trusted corporate network.

While the Production network's purpose is to allow external customers access to the

web servers while maintaining as secure an environment as possible, the Corporate network is designed to allow internal employees to perform their work most efficiently within their own departments. Each department composes its own network and includes all workstations, and systems that are primarily accessed by that department, including backup servers. Inter-department communications are strictly controlled, as is any communication with the Internet.

GIAC Enterprises Trusted Corporate Network



Filtering Router- The Corporate network's first line of defense from outside (Internet) attack is the packet filtering router #3. This 2610 Cisco router provides the same function as the router on the Web server network. (See Assignment 2: Security Policy)

Router #4 is a 7204 Cisco router providing VPN service to the production network, the two external partners. By placing the VPNs on a router located outside of one of the firewalls, the traffic can still be sniffed and filtered before entering the trusted network but is still given the protection of a VPN and packet filtering router. If the VPN went straight into the corporate network, it would become only as secure as the networks of their two partners.

Firewall #3- Sidewinder Firewall #3 is not only as the main defense against external attacks for the Corporate network, it also acts as a web and DNS proxy. It also screens the syslog server that collects data from the intrusion detection systems located on the screened network and on each department network. NAT is run on firewall #3 but is disabled on firewall #4 (See Assignment 2: Security Policy)

Firewall #4- While Firewall #3 deals primarily with traffic coming in from the Internet, Firewall #4 acts to separate and maintain the security of each department. It also handles and filters the traffic coming from the one internal VPN to the production network and the two VPNs to the fortune cookie saying supplier, Fortunes R Us, and GIAC Enterprise's international business partner. (See Assignment 2: Security Policy)

Screened Network- Much like the Production network, the screened corporate network houses the mail forwarder and an intrusion detection system as well as a terminal server serially connected to the router, switches, etc.

Mail- All mail will be received by the mail relay whose sole job is to store and forward the mail to the internal mail server. It will not retain any information about user accounts and will be locked down tightly with TCP wrappers, tripwire and its own firewall (ipfw) allowing only host-to-host connections on tcp port 25 and administrative ports to specific IP addresses.

DNS- DNS on the corporate network will be configured in the same way as on the production network. All queries from within the network will be handled by the internal ns server while the external ns server provides only the most generic name resolution for the domain (external address of the Sidewinder). (See DNS under the Production network configuration)

Assignment 2: Security Policy

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*

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*
7. *Explain how to test the ACL/filter/rule.*

There are many aspects of a security policy that are not explicitly asked for in this practical but that I would be remiss to omit. Therefore, I will briefly discuss a few of them and how they are implemented on this network design.

Physical Security- Physical security is often forgotten when designing a network security policy. However, it is much less trouble for a vandal to unplug a router than break into it remotely.

Physically securing the most important boxes can also help reduce any accidental service outages. For these reasons, all the boxes within the production network are placed in a secured, climate controlled environment. Any important databases or systems within the various branches can be secured locally or placed within a centralized secured area. Part of physical security is employee security. Many 'hacking' incidences are perpetrated from within. It is therefore advisable to do background checks before hiring anyone and keep aware of possible disgruntled employees. Never have a computer system dependent on one person and one on person should have 'all keys to the kingdom'.

Backups- Everything should be backed up on a regular basis. The information should be stored encrypted, and kept in a secure location away from the backed up system itself. If the computer was destroyed the data could then be restored onto a new box. Backup media should also be occasionally tested to ensure that the data is being transferred correctly.

Passwords and IDs- Each user on the network should have a unique ID and there are to be no shared passwords or accounts. This includes guest accounts. Administrators are to log onto boxes with their own IDs and then assume escalated privileges (i.e. root) only when the task requires it. Passwords should be at least 7 characters long and contain at least one upper case and lower case letter, number and special character. They also expire every 3 months and 3 incorrect logins lock a user account. When an employee ceases working for the company, all passwords known by the employee are immediately changed.

Patches and Anti-Virus Software- Keeping systems up to date with the latest patches is very important. A system administrator should check for patches and new system vulnerabilities on a daily basis. At www.securityfocus.com you can be put on the bugtraq mailing list to receive notifications about the latest news. Also look at the vendor website. Download only from well known reputable sites and before installing any patch, the patch should be tested. Anti-virus software should also put placed on the mail servers and on every user PC. As with the patches, virus updates should be downloaded as soon as they become available by the security administrator and then pushed out to the users PC's (on applicable systems).

End User Training and PCs- A system is only as secure as its users. Everyone who is going to be using a computer system, even a desktop PC, should be trained on the security implications of their actions. The most stringent ACLs cannot prevent a user from loading a floppy with a virus. A network security administrator's job becomes much easier when users are educated about simple things like not using chat services, surfing the web with caution, and not opening email attachments from untrusted sources. I would also recommend that personal firewalls be placed on desktop machines. This is especially true if any network users dial-in from home. Personal firewalls like BlackIce and Norton Personal firewall are affordable and user friendly.

The main layers of GIAC Enterprises' perimeter defense rely on packet filtering routers, firewalls, and VPNs.

Routers- All of the of border routers in my network design serve the same purpose and thus have very similar ACLs and configurations. I will cover them all in the following discussion. (The three VPNs running on router #2 and #4 will be covered separately. The first step is to harden the router itself. All services not being used should be disabled.

This can be done with the following commands:

(Global configuration mode)

<i>no ip bootp server</i>	Disable bootp service.
<i>no ip http server</i>	Disable the http management interface.
<i>no ip redirect</i>	This stops ip redirects from entering the network helping to prevent a DoS attack caused by a remote host maliciously modifying the target's routing table. (Be a good neighbor.)
<i>no service finger</i>	Disable finger responses which can be used to obtain computer information useful in breaking into a system.
<i>no snmp</i>	Disable snmp. If SNMP is enabled make sure that community names other than public and private are used and limit the IP addresses able to access this by using appropriate rules in the ACLs.
<i>no service tcp-small-servers</i> <i>no service udp-small-servers</i>	Disables small UDP and TCP servers like echo, chargen, and discard. These services, especially their UDP versions can be used to launch DoS and other attacks that would otherwise be prevented by packet filtering.
<i>no ip source-route</i>	Prevents IP source routing options from being used to spoof traffic.

(external interface)

<i>no ip directed-broadcast</i>	Prevents router from being used by an attacker as a "smurf" amplifier.
<i>no cdp enable</i> <i>no cdp running</i>	These stop Cisco Discovery Protocol from giving information about the router to directly-connected devices.
<i>no ip unreachable</i>	Stops the sending of ICMP unreachable messages for items denied in the ACLs. The less information an attacker can acquire the better. This also helps thwart udp scans.
<i>no ntp enable</i>	Prevents attacks against the NTP service. (Note- In this network design, logs are synchronized using an NTP server(s) that uses GPS system.
<i>no ip proxy-arp</i>	Disables the arp proxy services since it is not used.
<i>ip verify unicast reverse-path</i>	Checks each packet that is routed to check that the source IP does not have a route in the CEF table the points back to the same interface on which the packet arrived. It will drop any packet that does stopping SMURF attacks.
<i>no ip route-cache</i> <i>no ip mroute-cache</i>	These commands stop a router from using a cached route table. This will prevent a router from using a corrupted table.

To store the router's passwords as an encrypted string us the command *service password-encryption*. Passwords will still be transmitted over the network in clear text so use a program like

ssh to provided extra security.

At least all denied traffic hitting the router should be logged. To log to a syslog server, use the command *logging 'ip of syslog server'*. In the ACLs specify which rules are to be logged. Log as much traffic as possible in the beginning to understand what is 'normal' for the network.

A banner warning against any malicious activity is always a good idea and may be beneficial if a system is broken into and results in a court case. Use something like banner *!!!! No Unauthorized access permitted. Violators will be shot!!!!*. Banners can also be placed in the motd file. For more appropriate banners try <http://ciac.llnl.gov/ciac/bulletins/j-043.shtml> for more information on this subject.

It is also possible to configure rate limiting for SYN packets in order to slow down DoS attacks. Before doing this, it is important to measure the amount of SYN packets during a normal state.

```
interface e0
  rate-limit output access-group 102 XXXXXX YYYYYY ZZZZZZ conform-action
  transmit exceed-action drop
  rate-limit output access-group 101 XXXXXX YYYYYY ZZZZZZ conform-action
  transmit exceed-action drop
```

```
access-list 101 permit tcp any host eq www
access-list 102 permit tcp any host eq www established
```

XXXXXX is the maximum bandwidth

YYYYYY is the value that is between 50% and 30% of the SYN flood rate burst normal
ZZZZZZ and burst max rates.

(Egress Filter)

Egress ACLs on all routers are functionally the same and therefore very similar. Applied to the inside interface of the router, these rules listed below prevent GIAC Enterprises Network from being used in a DDoS attack by blocking spoofed addresses. Only valid IP addresses on the separate networks are allowed out to the Internet. In this case the only IP addresses allowed outside are the external (bur1) interface of Firewall #1 and #3.

```
access-list 102 permit 'real IP address from GIAC FWs' 0.0.0.0
access-list 102 deny any log
interface e1 (internal interface of router)
ip access-group 102 in
```

(Ingress Filters)

The Ingress ACL on the production server router #1,#2, and router#3 will block everything except traffic that is explicitly allowed. Since only DNS (53 UDP only), Https (443), Http (80) and out going SMTP (25) should be traversing router #1, these are the only ports that will be opened. This is done by writing allow rules for the fore mentioned traffic, ending with the rule *access-list 30 deny all all log* . (All denied traffic that hits the router should be logged. On the syslog box, simple scripts can be written to filter out noise traffic that is deemed as harmless.) The router has been configured this way to provide an extra layer of security for a network that is very attractive to vandals. This leaves the Firewall completely free to examine the remaining traffic for non-legitimate packets, i.e. tunneling malicious traffic through port 80. As this is a very simplistic ACL so I will not go into specifics.

For the purposes of this practical, the ACL on the corporate network router #3 will allow all traffic except what it specifically denies. (In reality, I would configure this router like router #1 with a *deny all all* statement at the end of the ACL.)

It is important to block spoofed addresses from entering your network. All addresses coming from private addresses spaces should be block.

```
access-list 101 deny 10.0.0.0 0.255.255.255.255 log
access-list 101 deny 172.16.0.0 0.15.255.255 log
access-list 101 deny 192.168.0.0 0.0.255.255 log
```

Also block your own address spaces from entering via the router.

```
access-list 101 deny 'your network IPs' 'mask' log
i.e. access-list 101 deny 192.168.180.0 0.0.0.255 log
```

Deny the loopback address.

```
access-list 101 deny 127.0.0.0 0.255.255.255 log
```

And illegal addresses

```
access-list 101 deny 0.0.0.0 255.255.255.255 log
access-list 101 deny 255.255.255.255 0.0.0.0 log
```

(Packets from the outside going to the network's broadcast addresses can also be blocked at the router.)

Ident can be used by an attacker to gather information so blocking it is a good idea.

```
access-list 101 deny tcp any any eq 113
```

Block ports that are associated with certain exploits. While this might cause some legitimate traffic to have problems entering the network temporarily, it could save an administrator big headaches later. The router is also a good place to block that annoying exploit de jure.

```
i.e. Ultors Trojan      access-list 101 tcp deny any any eq 1234
Deep Throat            access-list 101 udp deny any any eq 2140
Back Orifice           access-list 101 udp any any eq 31337
```

Any service can be blocked as well. I advise blocking those services in particular that are known security vulnerabilities and are popular. This cuts down on traffic to the firewall. These include those netbios ports, ICQ and terminal emulators like PC Anywhere.

To provide added security for the name servers, blocking tcp DNS at the router is advisable.

```
access-list 101 tcp any any eq 53
```

The last rule in this type of packet filtering router will be

```
access-list 101 permit any any
```

To enhance security, telnet will not be allowed on the routers. Instead, administrators will ssh to a terminal server which is directly connected to the routers. The terminal server itself is password protected and uses ipfw to allow traffic from specific IP addresses only.

Firewalls- For the firewalls, I have chosen Sidewinder version 5.0. These firewalls come pre-NSA hardened and provide a feature called type enforcement (TE). TE scrutinizes activity on the firewall and controls communication between the Internet and internal network. Security attributes are assigned to every login, process, and file. The firewalls are partitioned into domains. Each process and file can only reside and run within the domain in which it was created. Also every logon has some role and its domain determines its associated privileges. Root access only applies to a single domain. This means that an attacker must devise a new break-in strategy for each domain and they are unable to carry their 'tools' along with them because TE prevents compiling and installing foreign executables onto the OS. Different brands of firewalls could be used throughout the network in hopes of limiting the possibility of attack through shared vulnerabilities on the firewall. However, it is important that the firewall administrator feels comfortable with all the firewalls being used. It is more difficult to exploit a bug in the software than it is a poorly conceived rule set. Lastly, all firewalls are not to be remotely administered but only at the terminal. Rules for remote administration will not be seen in the following examples for this reason.

To avoid repetition, I will cover only the rules on firewall #1 and #4. Sidewinder rules are order dependent, so careful consideration should be taken in rule placement. It is first necessary to define the names that will be used in place of the burb numbers on both firewalls.

Firewall #1 Production Network

- External - Burb 1 access to the Internet
- Internal - Burb 2 containing the web servers and smart relay

Firewall #4 Corporate Network

- External - Burb 1 access to firewall #3 and the mail server
- Ship - Burb 2 Manufacturing and Shipping departments
- Finance - Burb 3 Human Resources and Finance
- Sales - Burb 4 Sales and Customer service
- Tech - Burb 5 All technical support including NT admins, networks, etc
- VPN - Burb 6 VPN point of entry

The network objects must also be identified.

Firewall #1

- Web - web server farm
- Mail - the smart relay

Firewall #2

- VPNUK - traffic from our International Partners using the VPN
- VPNFor - traffic from Fortunes R Us
- Prod - production VPN
- Mail - mail server
- Fordb - Fortune Cookie saying database
- Saledb - Sales database

Sidewinder, like most firewalls ends with a *deny all all* statement. "All" is denoted by a "*" in an ACL.

Please note that firewall #1 is running both NAT and DNS services.

Firewall #1 Access Control Rules

Name	Service	Agent	Action	Src Burb	Source	Des Burb	Des
http_in	http	Proxy	Allow	External	*	Internal	Web
http_out	http	Proxy	Allow	Internal	Web	External	*
https_in	https	Proxy	Allow	External	*	Internal	Web
https_out	https	Proxy	Allow	Internal	Web	External	*
dns_out	dns	Proxy	Allow	Internal	*	External	*
smtp_out	smtp	Proxy	Allow	Internal	Mail	External	*
deny_all	*	*	Deny	*	*	*	*

The first two rules allow for web traffic to be proxied from the Internet to the web servers and visa versa. Rules 3 and 4 provided the same function for SSL. These rules are placed first since most of the traffic coming through this firewall will trigger on one of these rules. The fifth rule allows for DNS to be proxied from the internal to the external burb. (See DNS under assignment #1 for more information). DNS traffic from the External burb is never allowed to access the internal name server. The last rule allows for mail to be proxied from the smart relay to the external burb only. The router's ACL will deny any mail sent to the internal network. Although not shown in the access controls rules table, all traffic through the firewall is logged.

Firewall #4 has a much more complicated rule set. However, it is not using NAT. For the sake of

simplicity, the technical administrators will be allowed to access entire burbs. Administrative objects could be set up to allow only specific IP addresses within each source burb to access specific IP addresses on each destination burb. We must assume that the external VPN traffic is going to a system on a particular burb. In this example Fortunes R Us is assessing, via the Shipping and Manufacturing, burb a fortune cookie database using sql. Our International partner is accessing the fortune cookie database and a sales database using sql as well. The last assumption is that everyone on our corporate network is allowed to surf the web.

Firewall #4 Access Control Rules

Name	Service	Agent	Action	Src Burb	Source	Des Burb	Des
sql_in	sql	proxy	Allow	VPN	VPNFor	Man	Fordb
sql_in	sql	proxy	Allow	VPN	VPNUK	Man	Fordb
sql_in	sql	proxy	Allow	VPN	VPNUK	Sales	Saledb
deny-vpn	*	*	Deny	VPN	*	*	*
http_out	http	proxy	Allow	*	*	External	*
https_out	https	proxy	Allow	*	*	External	*
smtp_out	smtp	proxy	Allow	*	*	External	Mail
smtp_in	smtp	proxy	Allow	external	Mail	*	*
dns_out	dns	proxy	Allow	*	*	external	*
prod_admin1	ssh	proxy	Allow	Tech	*	VPN	Prod
prod_admin2	ssh	proxy	Allow	Sales	*	VPN	Prod
tech_ssh	ssh	proxy	Allow	Tech	*	External	*
tech_ping	ping	proxy	Allow	Tech	*	*	*
deny_all	*	*	Deny	*	*	*	*

The first four rules on this table are used to control the access available from the VPNs to GIAC Enterprises internal network. The first three rules allow the two VPNs to access the systems that they require (as mentioned above). Since the rules are order dependent, the fourth rule denies all other actions taken by traffic originating from the VPN. Our International partner and the fortune cookie saying supplier will not be able to use our Internet connection, mail server, or access the production network.

All future rules discussed do not apply to incoming traffic on the VPN burb. The next two rules allow for all on the network to access the Internet via the proxy on firewall #3. Any filters i.e. porn or sex, placed on web traffic should be placed on the web proxy on firewall #3. Mail is accessible by all burbs to only the mail server itself. DNS is also accessible to all burbs. The last set of rules allow for management of the corporate network and production network. Administrators on the Tech burb are allowed to access via ssh both the production VPN and also the terminal server, mail and IDS box located on the external burb. In addition, they are permitted to ping all boxes on the network (corporate and production) for testing purposes. The Sales and Finance burbs are also allowed to access their information on the production VPN. (Their traffic can be more tightly managed to a specific system by the ACL on firewall #2.) Although not shown, rules should be in place allowing the technical staff to administer systems on the different branch's burbs. (Without knowledge of the systems located in these burbs, it is impossible to put effective rules into place). The final rule denies all other traffic going through the firewall. Any traffic hitting a deny statement is logged and sent to the syslog server.

VPNs- The three VPNs are used to provide extra security to reach the production network and the two partners companies accessed that must access GIACs network via the Internet. Because there is not address translation on firewall #4, it is important to have IP addresses assessing the VPN that will not conflict with the IP addresses on the Corporate network. The VPN use Cisco IPSEC with a MD5 hash, DES 3 encryption and rsa generated keys. (Note that the level of encryption used in an international VPN is dictation by both US exportation laws and the laws of the country at the other end of the VPN.) RSA keys can be generated using the global

command *crypto key generate rsa [usage-keys]*. Both the packet header and the payload will be encrypted to provide extra protection for sensitive business information. IP addresses allowed in through the VPN should also be limited only to those that need access. Using laptops to access the network via a VPN is strongly discouraged due to the possibility of laptop theft. The VPNs will use self-generated keys for the time being because there are currently no regulations regarding certificate authorities. Keys will be regenerated every 6 weeks. The following example is from VPN #3. Comments are made in blue.

```
!
crypto isakmp policy 1
  hash md5
  authentication rsa-encr
  group 2
!
!
crypto ipsec transform-set FORTUNEUK esp-des esp-md5-hmac
!
!
crypto key pubkey-chain rsa
  addressed-key 192.168.50.50 (address of interface at the other end of the VPN)
  address 192.168.50.50
  key-string (shared key)
    305C300D 490DG03V HJ20D43L 01010105 00034B00 BMD58021 EAF064D1 00CP074B
    8D098861 1475CE1C B2E8FD98 BA12I068 FB10FFED 35DF6A83 06092A86 48M6F70S
    C72141C1 F5FBC963 7RE09963 66C7A6E1 630203MK F03219KU P379403F 0001
  quit
!
crypt map FORUKMAP 10 ipsec-isakmp
set peer 192.168.50.50
set transform-set FORUKSET
match address 110 (Incoming IPs must match ACL 110)
!
!
!
interface FastEthernet 1/0
  description Fortune UK Ltd Entrance
  ip address 192.168.175.1 255.255.255.0 (router's external interface IP address)
  no ip redirects
  no ip directed-broadcast
  no ip route-cache
  no ip mroute-cache
  no cdp enable
(Put all the ACLs you would on any other router. See router configurations above)
  crypto map FORUKMAP
!
!
!
access-list 110 permit ip (source IP network) (destination IP network)
192.168.20.0 0.0.0.255 192.168.120.0 0.0.0.255
```

The best way to test that your security policy, ACL's VPNs, etc. are working is to perform a security audit; conveniently discussed in the next assignment.

Assignment 3: Audit Your Security Architecture

You have been assigned to provide technical support for a comprehensive information systems audit for GIAC Enterprises. You are required to audit the Border Router and 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 Border Router and Primary Firewall are 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.*

Assessment Plan

The best way to test all the ACLs and firewall that make up your perimeter protection is to run simulated attacks against the network. Bringing in an assessment team to run a scan like ISS would be ideal, however these teams are very expensive. A cheaper alternative is to use tools like Satan (www.fish.com/~zen/satan/satan.html), saint (www.wvdsi.com/saint/index.html), nmap (www.insecure.org/nmap/), or nessus (www.nessus.org). For this assessment I will be using the last two. The costs are therefore minimal with only the time –estimate 5 days- of one network security administrator and one firewall administrator. Spare boxes or laptops can be used to run the scans.

The first thing to consider is the purpose of this assessment. In this case, we want to determine that 1. The rules on the firewalls and ACLs on the routers are working correctly and in accordance with the policy. 2. The routers, firewalls, etc. are logging to the syslog servers correctly. 3. The intrusion detection systems are placed for maximum benefit. 4. The boxes themselves are secure 5. Adequate plans are in place for an organized response to security incidences and 6. Systems administrators are monitoring their system properly.

To fulfill the assessment's purpose, several penetrations tests need to be done. The first step is to research what company information is openly available to the public. Then, on both networks –corporate and production- nmap and nessus will be run from a box located on the outside of the gateway routers. The same tests are run between the routers and the firewalls. Similar scans are done on the internal firewall and against each 'sub' network with particular attention being paid to the screened network. Passwords can be checked to see that they are following policy by using tools like Jack the Ripper or l0ftcrack. Before doing any type of penetration testing or password cracking make sure to get written approval from upper management. This will provide some protection if legal issues arise. It is advisable to notify users of any password cracking attempts and their purpose ahead of time to prevent misunderstandings. Most of the scans and tests can be run during working hours but any tools that actually attempt exploits against the system should be done during periods of least impact to the user community. Such programs have been known to bring down systems and cause network disruptions. GIAC Enterprises does not use modems, but special attention should be paid to any dial-in access. Tools available to scan phone lines include ToneLoc and THC-Scan.

Implementation

Research should be done ahead of time to see how much information is available freely to the outside. Lookup your own IP addresses in ARIN (www.arin.net) and use Nslookup, dig, etc to find out how much information your ISP in particular is giving out via DNS. Nslookup and dig should also be used to insure that GIAC's own DNS has been configured to give out minimal information including its BIND version. Try sending an email to a non-existent user to see if any internal machine names are returned in the bounced email headers. A lot of information can be gleaned from a companies website such as employees and the branch they work in, phone numbers, computer systems, etc. This provides perfect information for social engineering

attacks. Scrutinize all information before publishing it on a website.

The first testing is against the gateway routers using nmap. I will be running the assessments against routers #1 and #3. Since nmap uses pings to establish the existence of a box on a particular address, it is the perfect way to test whether the router is dropping ICMP packets as designed. Nmap can then be used in a stealth mode using unescorted ACKs in place of pings since they should be passed by the router and by the firewall on open ports not running a proxy. Fingerprinting has also been turned on to attempt to identify the operating system. As expected, the results indicate that router #1 is filtering all traffic except that going to port 80, 443 and 53.

```
# Nmap (V. nmap) scan initiated 2.53 as: /usr/local/bin/nmap -sS -v -O -oN /tmp/nmap.133.010119
192.168.240.254
```

Interesting ports on (192.168.240.254):

(The 1519 ports scanned but not shown below are in state: closed)

Port	State	Service
53/udp	open	dns
80/tcp	open	http
443/tcp	open	https

If the scans show any port open that is not specified in the security policy the routers ACLs need to be reexamined. In this case, the open ports are acceptable.

Router #3's results show that it is far more open, but that it also is blocking ICMP packets effectively. Tests can then be continued from outside the router, or as follows from behind the router. This tests firewalls more stringently. Since the firewalls have also been configured to drop pings coming from the external interface (burb1) unescorted ACKS are again used. The same scans should be run against each firewall and the results compared against the security policy. The results of these scans produced finding in accord with the policy.

Because the mail relay, web servers and DNS servers are in more exposed positions, special attention should be paid to these boxes. The following nmap and nessus scans were run against the mail relay itself.

```
#Nmap (v. nmap) scan initiated 2.53 as: /usr/local/bin/nmap -sS -PT25 -v -oN nmap.sv -O
192.168.170.110
```

Warning: No TCP ports found open on this machine, OS detection will be MUCH less reliable

All 1523 scanned ports on dopey (192.168.170.110) are: filtered

Too many fingerprints match this host for me to give an accurate OS guess

TCP/IP fingerprint:

T5(Resp=N)

T6(Resp=N)

T7(Resp=N)

PU(Resp=N)

```
# Nmap run completed at Wed Jan 17 17:52:19 2001 - 1 IP address (1 host up) scanned in 196 seconds
```

This scan not only demonstrates that the mail relay box is tightly configured with tools such as TCP wrappers and ipfw, but also that it is logging to the syslog server correctly. For example, this log was generated during the nessus scan.

```
Jan 17 17:48:33 dopey /kernel: ipfw: 40511 Deny TCP 192.168.yy.zz:35304 192.168.170.110:6003
```

The following is the results of the nessus scans that triggered this log message.

Nessus Scan Report

SUMMARY

- Number of hosts which were alive during the test : 1
- Number of security holes found : 3
- Number of security warnings found : 3

TESTED HOSTS

192.168.170.110 (Security problems found)

DETAILS

+192.168.170.110:

. List of open ports :

- o smtp (25/tcp) (Security hole found)
- o general/udp (Security warnings found)
- o general/tcp (Security hole found)

. Vulnerability found on port smtp (25/tcp) :

The remote SMTP server did not complain when issued the command :

```
MAIL FROM: root@this_host
RCPT TO: /tmp/nessus_test
```

This probably means that it is possible to send mail directly to files, which is a serious threat, since this allows anyone to overwrite any file on the remote server.

NOTE : ** This security hole might be a false positive, since some MTAs will not complain to this test and will just drop the message silently. Check for the presence of file 'nessus_test' in /tmp ! **

Solution : upgrade your MTA or change it.

Risk factor :
High

. Vulnerability found on port smtp (25/tcp) :

The remote SMTP server did not complain when issued the command :

```
MAIL FROM: root@this_host
RCPT TO: |testing
```

This probably means that it is possible to send mail directly to programs, which is a serious threat, since this allows anyone to execute arbitrary command on this host.

NOTE : ** This security hole might be a false positive, since some MTAs will not complain to this test, and instead will just drop the message silently **

Solution : upgrade your MTA or change it.

Risk factor :
High

. Vulnerability found on port smtp (25/tcp) :

There is a buffer overflow when this MTA is issued the 'HELO' command issued by a too long argument.

This problem may allow an attacker to execute arbitrary code on this computer, or to disable your ability to send or receive emails.

Solution : contact your vendor for a patch.

Risk factor :
High

Similar scans were also run on the web servers.

```
# Nmap (V. nmap) scan initiated 2.53 as: /usr/local/bin/nmap -sS -PT80 -v -oN nmap.la 192.168.180.20
```

Interesting ports on donald (192.168.180.20):

(The 1517 ports scanned but not shown below are in state: filtered)

Port	State	Service
22/tcp	closed	ssh
80/tcp	open	http
123/tcp	closed	ntp

TCP Sequence Prediction: Class=random positive increments
Difficulty=21413 (worthy challenge)

Sequence numbers: A99A649F A99BD7DC A99D6D09 A99E39B4 A99F2D49

Remote operating system guess: FreeBSD 2.2.1 - 4.0

The following is part of the nessus scan of the web server. Yet again TCP wrappers has been used. Please note that the vulnerable port found is blocked at the firewall.

Nessus Scan Report

SUMMARY

- Number of hosts which were alive during the test : 1
- Number of security holes found : 1
- Number of security warnings found : 1

TESTED HOSTS

192.168.180.20 (Security problems found)

DETAILS

- 192.168.180.20 :
. List of open ports :
o general/udp (Security warnings found) o general/tcp (Security hole found)
- . Vulnerability found on port general/tcp :

It was possible to disconnect the remote host by sending it an ICMP echo request packet containing the string '+ + + ATH0' (whithout the spaces). It is also possible to make the remote modem hangup and dial any phone number.

Solution : add 'ATS2%5' in your modem init string.

Risk factor :

Medium

I also attempted to telnet to the web server through port 80. While it was possible within the 192.168.180.0 network, the attack was blocked by the firewall.

The last boxes that need particular attention are the DNS servers (running on the firewalls). BIND is not easily configured, but is essential. To make matters worse, DNS a prime target by vandals for information gathering and attacking. There are some simple precautions that can be taken to tighten down DNS as briefly discussed already. Dig can be used to check that DNS zone transfers have been denied. This is done by adding *axfr* to the end of a dig request originating from an outside dns server for GIAC's domain name. If DNS is properly configured to allow only zone transfers from your secondary ns servers, there should be a negative result returned from the DNS server. It this case we are denied by the router and firewall. It is therefore necessary to do this test from the same network, however it yields the same negative results. To prevent cache poisoning, recursive queries should also not be allowed

from the external ns. This can be tested by doing a nslookup query from the name server being tested. Recursion is turned off if it does not come back with a final resolution response. To ensure that the name server is not advertising the version of BIND running, it should also be tested using dig.

```
# dig -t txt -c testing VERSION.BIND ns.giac.com
```

The answer to this query depends on what has been placed in the "named.conf" file on the ns server. I recommend that it be configured to reflect something generic like "BIND version 'wouldn't you like to know'." This should be done for every name server on the network.

Analysis

The results of the perimeter assessments should be discussed with management and system administrators, and a date set to close all the findings. Another scan should then be conducted 1 to 3 months later depending on the severity of the findings. Vulnerabilities with a high threat level should be verified as closed as soon as possible. Regular scans are then conducted every 3 to 6 months unless major changes are made to the network.

The perimeter design actually proved to be quite robust when attacked from outside the routers and/or firewall. Many of the boxes themselves are also quite secure. One recommendation is to configure as many of the systems as possible to not send resets to unescorted ACKs on ports that they are not listening on or for un-established connections. I would also recommend that all banners be configured to not divulge the operating system. Some systems can even be configured to lie to any query about version or OS. Vandals can be tricked into wasting time attempting hacks for the wrong OS allowing the administrator more time to catch an attacker in the act and prevent any real damage. Placing a honeypot in the DMZ (Behind router but in front of the firewall) makes for interesting reading and helps a security administrator understand how attacks occur.

As mentioned in assignment #2, I would also recommend placing an ACL on router #3 similar to that on router #1. Denying everything except for what is explicitly permitted makes it more difficult for an attacker because it adds another layer of security the vandal must traverse. Lastly, although performing your own perimeter defense assessment is affordable and should be done on a regular basis, it is also strongly recommended that outside persons be allowed to conduct similar tests. Internal personnel are often too subjective as to the security of their network and may unintentionally overlook design flaws. It also adds to the attack knowledge base, providing for a more thorough security audit.

Assignment 4 - Design Under Fire (25 Points)

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:

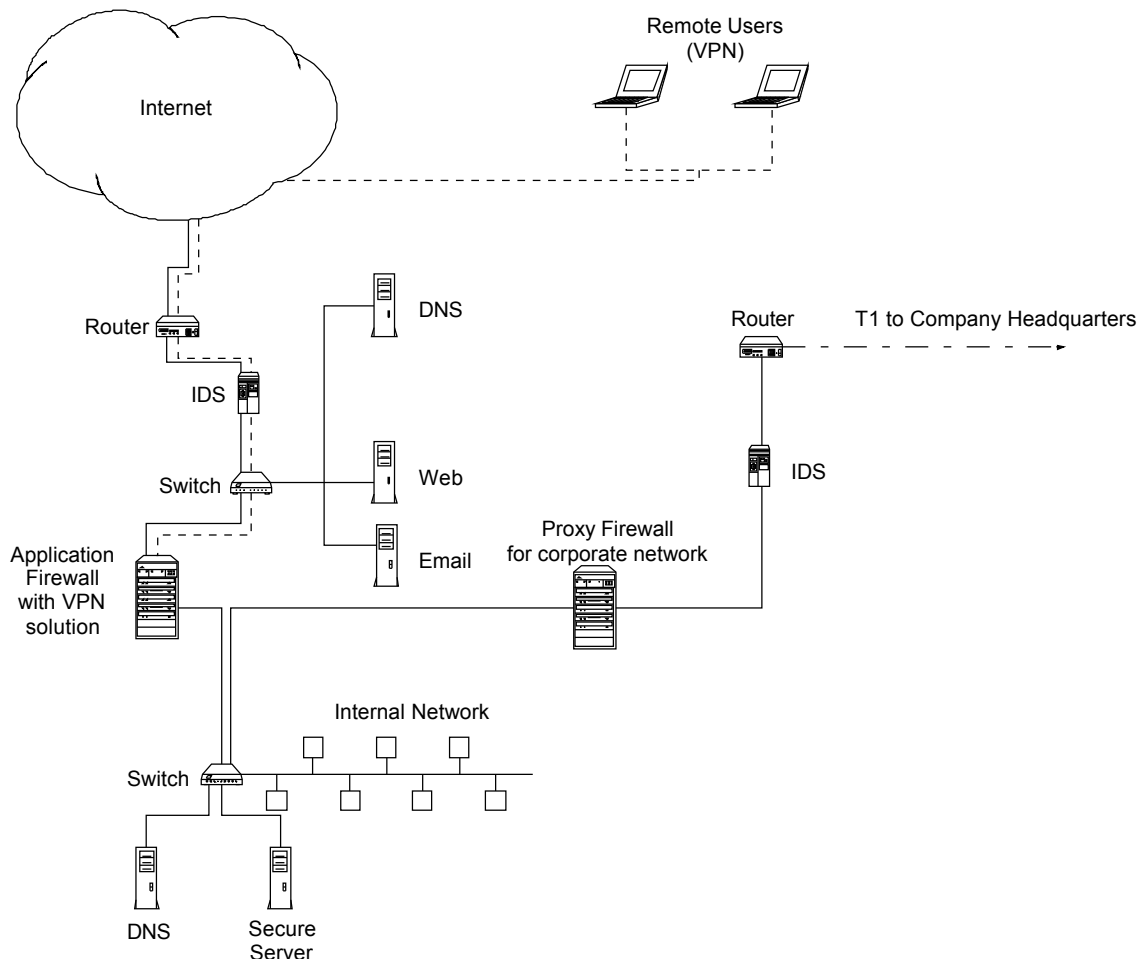
1. An attack against the firewall itself. Research vulnerabilities that have been found for the type of firewall chosen for the design. Choose an attack and explain the results of running that attack against the firewall.
2. A denial of service attack. Subject the design to a theoretical attack from 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.

I have chosen to use Chris Stevenson network design. While he does state that the firewall is a Checkpoint, he has not given specifics about the routers, web server, or email server. For

this reason, I will need to make assumptions as to his security design and security policy. I will cover several possible attacks and their consequences for each of the three scenarios. The attack method used would inevitably depend on what the attacker discovered about the network through scans, and various data mining attempts.

http://www.sans.org/y2k/practical/Chris_Stevenson.doc.

Attack 1- The firewalls employed in this design are both Checkpoint FW-1s. Unfortunately, the designer of this architecture did not state what version the firewall was running only that one of the FW-1 is also providing VPN service. The first obstacle to attacking the firewall directly is the packet filtering router. This router is blocking IP spoofing attacks originating from within the network and filtering ICMP packets. I would first run a tool such as nmap in stealth mode (to attempt to avoid detection) using an unescorted ACK to discover the ports open on the servers on the screened network and the firewall. This can also give information about the operating system and version of the firewall. With this information, an attack method can be chosen. (I would prefer not to attack the firewall itself but instead to bypass the firewall altogether. This could be done by introducing a Trojan onto a computer system located on the other side of the firewall. It could be hidden in an email attachment, placed on the computer via a compromised sendmail or web server on the screened network. It would also be quite easy to trick an employee into downloading the Trojan onto his computer, ie "really cool game, download now." This Trojan could then establish a connection to my computer at a preset time on a port that is allowed out of the firewall such as 80)



Since this assignment requests an actual attack against the firewall itself, I have chosen a denial of service attack that takes advantage of a VPN-1 vulnerability on version 3.0 and 4.0. On a Firewall-1 a packet is first sent through an internal stack maintained by the firewall before it is passed onto the operating system's native stack, and checked against its ACL. If this connection is allowed it is added to the connection table with a timeout set to 60 seconds. It is raised to 3600 seconds when the remote host responds with an ACK. If a connection is initiated with an ACK packet, and it is allowed by the firewall's ACL, the timeout is set to 3600 seconds and does not care if the remote system responds. Sending multiple ACK packets as root from an internal system on port 80, for example, will quickly fill up the connections table and place the firewall in a 'failed closed' state. (I have already covered several ways to gain access to an internal box. Alternatively, all this is made easier by getting possession of one of the laptops using the VPN.) A simpler attack is to force the firewall (version 4.0 or 4.1) to use 100% of its available processing time by sending illegally fragmented packets to or through the firewall. The firewall will use its processing time logging these packets. Both of these attacks were found on Bugtraq and can be prevented on properly patched firewall-1s. These are only two of the many recent attacks found against checkpoint FW-1.

Attack 2- I have chosen to use TFN2K as my DDoS attack against this network because Unix, Solaris and Windows boxes that are directly or in directly connected to the Internet are susceptible. TFN2K daemons residing on clients are controlled by the master –with a spoofed IP- via TCP, UDP, ICMP, or by all three methods randomly. The agents in turn have the same flexibility and can launch flood attacks against the target by TCP/SYN, UDP, ICMP/PING, Broadcast packets or alternate between the attack methods. The TFN2K daemon will not reply to the clients which sends 20 encrypted command packets, interspersed with decoy packets sent to random IP addresses, to the daemon. It is impossible to say how this network configuration would respond to this attack for certain. But, even with a router configuration that filters ICMP packets, it would most like result in a dramatic slow down or 100% processor utilization consumption on the firewall. In addition it should easily cause a denial of service on the email, DNS, and web servers.

There is no way to completely prevent a TFN2K DDoS attack. One of the best methods is to ensure that systems on a network do not become clients. Place rules on routers and firewalls that prevent spoofed IP addresses from entering or leaving the network (as was done on this router). Use an application proxy firewall or keep non-proxy services to a minimum. This should help prevent unsolicited traffic inbound from the Internet. Block ICMP packets if possible at the router or firewall, if this is not possible, at least block Ping traffic. Deny UDP and TCP traffic except on specific ports used by the network.

There are a few things that can help slow down a network against any DoS attack. Always keep up-to-date on software patches. There are many older DoS service attacks that have been fixed in newer software releases. Have multiple connections to the Internet. This will prevent a DDoS attack from completely shutting down business. Tune software so that it can handle increased amounts of traffic and set timeouts for network connections. Distribute the load across multiple servers or sites i.e. use round robin DNS or a device like a Cisco director. Also use redundant routers and firewalls. Filter out all incoming RFC 1918 address spaces. Configure rate limiting for SYN packets as shown in assignment #2 routers. Lastly, use the *ip verify unicast reverse-path* command to stop smurf attacks. One of the most important things that should be done is to gather as much information as possible during the attack. This can be done using a program like tcp dump with the command *tcpdump -l interface -s 1500 (MTU size) -w capture_file*. This information should then be given to the local FBI office as soon as possible. (www.fbi.gov/contact/fo/fo.htm)

Attack 3- Since the DNS server, Web server and Email server are only protected by a packet filtering router, they are some of the easiest targets for attacks. I have chosen to attack the web

server with the assumption that it is running IIS. In addition, web servers offer many interesting opportunities to disrupt normal business, access a protected network, launch other attacks, and discover sensitive data etc. I have chosen attacks that can easily be done by remote users by modifying the URL sent to the IIS server. The router provides no protection for the web server in these types of attacks..

The first attack I have chosen is very simplistic. IIS 3.0 supports server side scripting using "Active Server Pages" .asp files. These are not mean to be visible to the user and can contain sensitive information such as passwords. By appending a . to the end of a URL, a remote user can display .asp files. For example <http://www.international.com/default.asp> becomes <http://www.international.com/default.asp.> Microsoft claimed to have patched this bug but in fact, this exploit is still possible on machines that have been 'patched'. By simply replacing the '.' before asp with %2e, (<http://www.international.com/default%2easp>) an attacker can save the contents of the .asp file to disk for later viewing.

More sensitive information can be obtained on IIS 4.0 because the ISAPA application webhits.dll that provides hit-highlighting for Index Server can be used by an attacker to get access to files on the same logical drive. This could include log files, customer databases and Active Server Pages. The hit-highlighting allows a web user to have a document returned with the original search terms highlighted on the page. The name of the document is passed to the .htw file with the CiWebHitsFile argument. Webhits. opens the file highlights and returns the resulting page. Because the user has control of the CiWebHitsFile argument passed to the .htw file they can request almost anything. Also, the source of ASP and other scripted pages can be revealed too. Webhits.dll will follow double dots allowing an attacker is able to gain access to files outside of the web virtual root as long as they know how to build the URL. For example, <http://charon/iissamples/iissamples/ooop/qfullhit.htw?CiWebHitsFile=../../winnt/system32/logfiles/w3svc1/ex000121.log&CiRestriction=none&CiHiliteType=Full> allows a vandal to view the web access logs. This attack can even be done if no .htw files exist on the system. This is achieved by getting inetinfo.exe to invoke webhits.dll and getting webhits.dll to access an existing file. Crafting a special URL can accomplish this.

Finally, if no useful information can be gleaned from these ventures, a simple DoS attack can disrupt business. One amongst many I have found affects an NT 4.0 server running IIS version 3.0. A remote user can stop web services by sending a URL that contains a certain number of characters i.e. <http://GIACwebserver.com/?something=XXXXXXXXX...> The attacker must use Netscape Navigator and the URL must contain a CGI value and value pair of a certain length, around 8K. A tool called IIServerSlayer has even be developed that aids in determining the correct length bring the server down. These are only a few ways to cause havoc on and IIS server. If nothing else, the examples I have chosen should demonstrate that the security measure places on each system are just as important as a networks perimeter protection