Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

SANS GIAC Firewall and Perimeter Protection Practical Assignment

SANS Network 2000 Monterey,CA.

Submitted by: Larry Koons
November 17, 2000

Assignment 1: Security Architecture

Define a security architecture that employs filtering routers, firewalls, VPN's and internal firewalls to rapidly implement the VISA "Ten Commandments" to the extent possible at GIAC Enterprises. The "Ten Commandments" are listed below:

1. Install and maintain a working network firewall to protect data accessible via the Internet.
2. Keep security patches up-to-date.
3. Encrypt stored data accessible from the Internet.
4. Encrypt data sent across networks.
5. Use and regularly update anti-virus software.
6. Restrict access to data by business "need to know".
7. Assign unique ID's to each person with computer access to data.
8. Track access to data by unique ID.
9. Don't use vendor-supplied defaults for system passwords and other security parameters.
10. Regularly test security systems and processes.

The network is depicted in Fig. 1. We will use a Cisco 7507 router running 12.1.5 IOS. This router will connect to the Internet and act as the ingress/egress filtering router. It is at this router where we will filter for spoofed addresses that have source addresses of our internal IP addresses. We will also filter out RFC1918 addresses as well as IP addresses referenced in the draft-manning-dsua-03.txt document available online at http://search.ietf.org/internet-drafts/draft-manning-dsua-03.txt . We will block IP source-routed packets from entering the network. We will block other specific TCP/UDP services from entering into the network and directed at the router that are not necessary. We will restrict incoming data to only the IP's that reside in our network.

Egress filters will be applied in the outbound direction to limit the source addresses to those of our networks. This will ensure that we are not used as a source for a spoofed attack. That is no addresses but our legitimate addresses can be the source of a packet.

IP directed-broadcasts will be turned off on all interfaces in the router to prevent smurf attacks from our network. A smurf attack is when a ping is sent to the broadcast address of a network. All devices on that network will respond and the replies will be sent to the originator of the ping. As you can see if all devices answer to a single ping the return traffic will be much greater towards the originator. This type of attack is normally used with a spoofed source address so that the return traffic goes to a third party . This large amount of traffic to the third party will in cases prevent them from doing legitimate business. This could be caused by either using all of their bandwidth or using all of the cpu resources on their servers. This is known as a Denial-of-Service (DOS) attack. Craig Huegen has written an excellent white paper on smurf and DOS attacks available at http://www.pentics.net/denial-of-service/white-paper/smurf.txt.

As stated above we will block all incoming packets with a source IP of our own internal networks. This will block anyone from outside accessing the internal systems that may allow access to them from internal addresses.
We will block traffic from RFC1918 space and the draft-manning-dsua-03.txt addresses since these should not be routing in the Internet in the first place.
IP source-routed packets will be blocked because they could be used to have a device or host return packets using selected routes. These routes may bypass firewalls and/or router filters.

We will block specific TCP/UDP services directed at the router such as echo, chargen, finger, as well as ICMP-Redirects. All of these services could be used to cause a DOS on the router itself. These are also useful in fingerprinting the network. Attackers will fingerprint a network by sending these service requests at networks and observing the responses. In this way they can tell what type

of devices are in the network which will help them in planning other attacks. Cisco has a nice description of preventing these attacks at http://www.cisco.com/warp/public/707/21.html. We will restrict incoming packets to only the legitimate IP addresses in our network. This will help alleviate the processing on the internal devices from traffic that does not belong inside the network.

We will use the Lucent Brick for VPN services with our partners, suppliers and remote users. This traffic will be encrypted and transmitted within a tunnel between them and the Lucent Brick. The users will be required to use the Lucent IPSEC client on their machines. This will ensure secure traffic. This traffic will terminate on servers on the internal network. The users have to authenticate with the Brick which uses a RADIUS database. There will be an audit trail on the RADIUS server. The Lucent Brick will use NAT to protect the internal IP addresses. All users will only use a single IP address to communicate to the Lucent Brick which will be the tunnel endpoint.
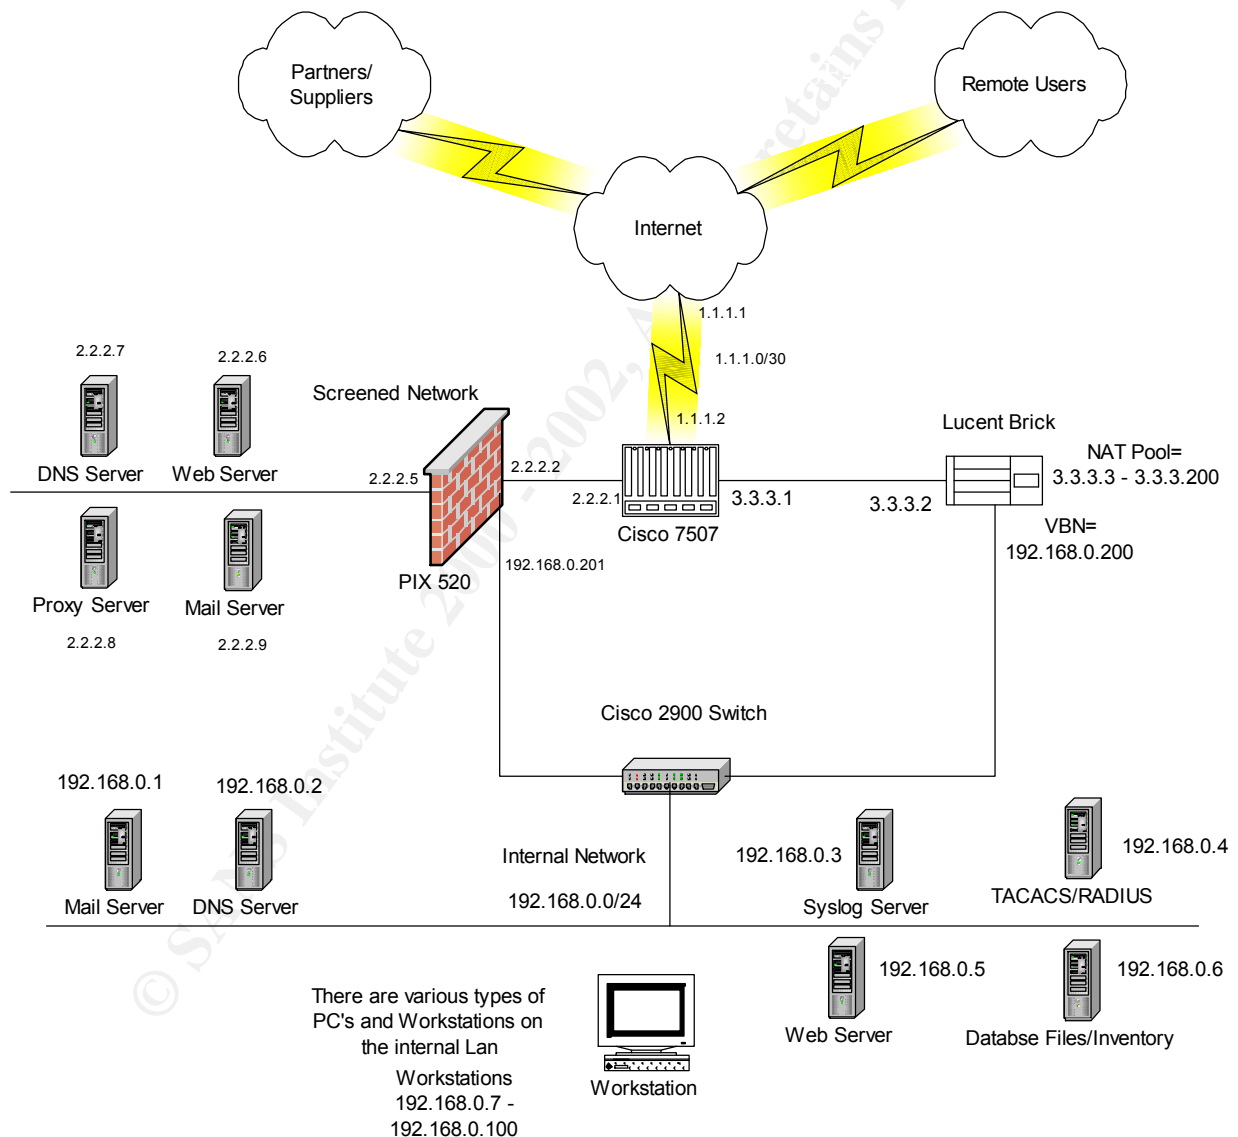

Figure 1

Assignment 2: Security Policy

For the purposes of this assignment, your security policy should be focused on implementation of requirement number 1 above "Install and maintain a working firewall to protect data accessible via the Internet." Write a tutorial on how to implement each recommended action in the filtering policy on your network. These additional filters will be in addition to the "TOP TEN" recommendations located at http://www.sans.org/topten.htm. The "TOP TEN" will not be included here, but will be assumed to be implemented.

I will briefly explain the traffic flow for the network along with some general guidelines for the security policy. This will be in reference to perimeter security and not how servers actually perform their functions.

First, customers will access the Web server on the screened network to place orders. They can use HTTP or HTTPS if they want a secure transaction. The web server will be running SSL for secure transactions. The customer will have a choice of which method to use. Customers will be allowed to send email to the external mail server.

The external DNS server will only have addresses for the external servers and the Lucent Brick. This will hide the internal network from the Internet while allowing access to only the devices they need to get to.

The partners, suppliers and remote users will use the Lucent IPSEC client to tunnel through the Internet to reach the Lucent Brick. We will be using the security protocol ESP, the 3DES encryption algorithm and HMAC-MD5 for authentication. Once authenticated on the Brick they will be allowed to connect to the internal database/inventory machine to conduct business. They will not be allowed access to any other machines on the internal network.

Internal users will use the internal DNS. Internal users will use the Proxy Web Server on the screened network to access the Internet. This will be done using port 8080. To intiate a connection to partners or suppliers they will use the Lucent Brick and establish an IPSEC connection.

The Cisco router and PIX firewall will both use the TACACS+ server for authentication, authorization and accounting. SSH and Telnet will be allowed to access them but only from certain machines on the internal network. TFTP will be allowed from certain machines to get/put configurations.

We will assume the following IP addresses for demonstration purposes only. They DO NOT depict any real network. It will make interpreting the rulebase easier.

The ISP has allocated the following networks to us:

1.1.1.0/30 - This will be used for the wan link between our router and theirs.
2.2.2.0/24 - This will be used on the screened network.
3.3.3.0/24 - This will be used for the VPN network. The NAT pool will be 3.3.3.3 – 3.3.3.200

We will start with the Cisco 7507 router. It is here where we will place our ingress/egress filters.

Telnet to the router. Once authenticated via TACACS+ we will go to configuration mode. The following are the commands that would be entered. The prompt will be router#. The commands that you enter will be in italics.

We will first go into global configuration mode and enter commands that affect the entire router. This is where we will shut off TCP/UDP services not used. Cisco has a write up on this subject at

http://ww.cisco.com/warp/public/707/21.html - services. Cisco recommends that you turn off these services unless absolutely necessary. We will also configure the TACACS+ while in global config mode.

Router# *config t*
Router(config)#*no ip source route*
Router(config)#*no service tcp-small-servers*
Router(config)#*no service udp-small-servers*
Router(config)#*no service finger*
Router(config)#*service password-encryption*
Router(config)#*enable secret <the secret password>*
Router(config)#*aaa new-model*
Router(config)#*aaa authentication login SANS tacacs+ local*
Router(config)#*aaa authentication enable default tacacs+ enable*
Router(config)#*aaa authorization commands 15 tacacs+ local*
Router(config)#*aaa accounting exec start-stop tacacs+*
Router(config)#*aaa accounting commands 15 start-stop tacacs+*
Router(config)#*tacacs-server host 192.168.0.4*
Router(config)#*tacacs-server key <the key used with the TACACS+ server>*
Router(config)#*logging buffered*
Router(config)#*logging 192.168.0.3*

No ip source route will not allow packets that are marked to use source routing.
No service tcp-small-servers and *udp-small-servers* blocks access to the TCP/UDP services echo,chargen and discard. These services could be used as a DOS attack against the router itself.
No service finger will block finger requests that would provide some information about what users are logged into a network device. This could potentially provide an attacker with some useful information.
The *service password encryption* command encrypts the passwords in the configuration.
Enable secret is used to enter the enable password. It is more secure than using the "enable password" command. The enable secret command uses a one-way hash so the password can't be broken.
Aaa new model defines the new TACACS+ protocol for authentication, authorization and accounting.
Aaa authentication login SANS tacacs+ local instructs the router to use the TACACS+ server to authenticate, authorize and send accounting information to for all users.
Aaa authentication enable default tacacs+ enable instructs the router to use the TACACS+ server when a user tries to go to enable mode.
The other aaa commands instruct the router to use the TACACS+ server to authorize all commands entered at the command line. The router will also send all accounting information to the TACAS+ server so there will be an audit trail of what the user did when they were logged on. The accounting information could be useful during forensics of a problem in the network.
Logging buffered turns on logging on the router. There will be a buffer that will hold some of the logging that you see by entering the command *"sho log"* at the command line. The logging will also be sent to the syslog server defined by the line *"logging 192.168.0.3"*.

You could test these rules by trying to connect to the router from the outside network.
As an example you could try to telnet to the router on port 7 to see if the router will echo data back to you. Telnet to the router on port 19 to check for the character generator process. These services should not work. You could also use a tool such as nmap to test the router. Nmap is available at http://www.insecure.org/ .

An example of an nmap scan to ports 7 and 19 of the router from the Internet would look like the following: nmap 1.1.1.2 –sS –v –p 7,19,23 . The result was the following:

Starting nmap V.2.30BETA21 by fyodor@insecure.org (www.insecure.org/nmap/)
Host 1.1.1.2 appears to be up
Initiating SYN half-open stealth scan against 1.1.1.2
The SYN scan took 1 second to scan 2 ports.
Interesting ports on (1.1.1.2):
Port   State   Service
7/tcp   filtered   echo
19/tcp   filtered   chargen

The "filtered" state means that nmap received no reponse to it's request. This normally means
that the request was filtered in some device. This would be true in our case since we have the
tcp-small-servers turned off.

We will now configure the router to only allow certain machines on the internal network to telnet
into the router. We will setup an access-list that restricts by IP addresses and apply the list to the
vty ports. In configuration mode do the following:
Router(config)#*access-list 1 permit 192.168.0.7 0.0.0.15*
Router(config)#*line vty 0 4*
Router(config-line)#*access-class 1 in*
Router(config-line)#*login authentication SANS*

The access-list only permits machines with the IP addresses from 192.168.0.7 through
192.168.0.15.
Line vty 0 4 defines the virtual terminals that allow you to telnet and login to the router
over an interface other than the console.
Access-class 1 in means to only allow telnet from the IP addresses defined in access-list 1.
Login authentication SANS associates the login to the "*aaa authentication login SANS tacacs+ local*"
command entered earlier. This tells the router to use TACACS+ to authenticate anyone using
the vty ports.

You can test this by trying to telnet to the router from the outside and from the inside.
You would receive the message "connection refused" from the outside and would be prompted
for a login and password if you were on one of the machines allowed in access-list 1.

We will now setup Serial0 that connects to the Internet. We will build an access-list that blocks
unwanted traffic here at the ingress to our network. These are in addition to "TopTen".
We will also turn off "directed broadcasts" and "ip unreachables". This will stop our interface
From being used as a smurf amplifier and sending backing ICMP unreachable messages.
Telnet to the router and enter configuration mode. The order in which you enter the entries is very
important. The lists are checked in order from top to bottom and pass on the first match, not the
best match. You have to make sure that the traffic you want to pass/drop is hitting the correct rule
otherwise it may have the opposite effect you wanted.

Router#*config t*
Router(config)#*access-list 100 deny ip 10.0.0.0 0.255.255.255 any*
Router(config)#*access-list 100 deny ip 172.16.0.0 0.15.255.255 any*
Router(config)#*access-list 100 deny ip 192.168.0.0 0.0.255.255 any*
Router(config)#*access-list 100 deny ip 0.0.0.0 0.255.255.255 any*
Router(config)#*access-list 100 deny ip127.0.0.0 0.255.255.255 any*
Router(config)#*access-list 100 deny ip192.0.2.0 0.0.255.255 any*
Router(config)#*access-list 100 deny ip169.254.0.0 0.0.255.255 any*
Router(config)#*access-list 100 deny ip224.0.0.0 31.255.255.255 any*
Router(config)#*access-list 100 deny ip2.2.2.0 0.0.0.255 any*
Router(config)#*access-list 100 deny ip3.3.3.0 0.0.0.255 any*

Router(config)#*access-list 100 permit tcp any host 2.2.2.6 eq http*
Router(config)#*access-list 100 permit tcp any host 2.2.2.6 eq https*
Router(config)#*access-list 100 permit udp any host 2.2.2.7 eq domain*
Router(config)#*access-list 100 permit tcp any host 2.2.2.9 eq smtp*
Router(config)#*access-list 100 permit tcp any host 2.2.2.8 established*
Router(config)#*access-list 100 permit 50 any host 3.3.3.2*
Router(config)#*access-list 100 permit 51 any host 3.3.3.2*
Router(config)#*access-list 100 permit udp any host 3.3.3.2 eq 500*
Router(config)#*access-list 100 deny ip any any*

We will now enter interface configuration mode.

Router(config)#*int serial0*
Router(config-if)#*no ip directed-broadcast*
Router(config-if)#*no ip unreachables*
Router(config-if)#*ip access-group 100 in*

We have now applied the access-list 100 inbound on serial 0 . This will block all traffic
except for the traffic that we have explicitly allowed. We have allowed HTTP and HTTPS to
the web server, SMTP to the mail server, DNS to the DNS server and IKE,ESP and IKE to the
Lucent Brick to used for the VPN tunnels. Everything else will be blocked. You may find that
you have to make adjustments to this list as your business needs change. A good rule to follow
is "only allow what is necessary and block everything else." You can check what type of traffic
is coming into your network by issuing the command "*sho access-list 100*". This will show you
the number of times that each line has been hit by a packet. An output would look like the following:

 deny ip 10.0.0.0 0.255.255.255 any (1234 matches)
deny ip 172.16.0.0 0.15.255.255 any
deny ip 192.168.0.0 0.0.255.255 any
deny ip 0.0.0.0 0.255.255.255 any
deny ip127.0.0.0 0.255.255.255 any (20 matches)
deny ip192.0.2.0 0.0.255.255 any
deny ip169.254.0.0 0.0.255.255 any
deny ip224.0.0.0 31.255.255.255 any
deny ip2.2.2.0 0.0.0.255 any
deny ip3.3.3.0 0.0.0.255 any
permit tcp any host 2.2.2.6 eq http (1898765 matches)
permit tcp any host 2.2.2.6 eq https
permit udp any host 2.2.2.7 eq domain
permit tcp any host 2.2.2.9 eq smtp
permit tcp any host 2.2.2.8 established (455968675)
permit 50 any host 3.3.3.2
permit 51 any host 3.3.3.2
permit udp any host 3.3.3.2 eq 500
deny ip any any (4567)

The numbers are the number of times that a packet satisfied the requirements of that line.
If you start to see the lines with deny on them getting big numbers, you would want to contact
your provider and request them to block that traffic from reaching your network. It will be up to them
to trace these packets back to their origination.

We will now build access-lists to apply to e0 and e1. This will be applied inbound and will only allow
source addresses from the connected network. This will stop us from sending spoofed packets out
to the Internet. You could build an access-list of your valid networks and apply it outbound on the Serial
interface and it would work just as well. I prefer to stop traffic at the ingress point into the router. If you

applied it to the Serial interface outbound the router would have to accept the packet from E0, analyze it, route it to S0 and then apply the access-list. When the access-list is applied to the incoming interface, the first thing the router does is evaluate the incoming packet against the access-list. If it fails the access-list the router drops the packet and no further processing is required. This could become important in a busy router as processing power becomes a factor. Enter configuration mode on the router.

Router(config)#access-list 101 permit ip 2.2.2.0 0.0.0.255 any
Router(config)#access-list 101 deny ip any any
Router(config)#access-list 102 permit ip host 3.3.3.2 any
Router(config)#access-list 102 deny ip any any

Router(config)#*interface e0*
Router(config-if)#no ip directed-broadcasts
Router(config-if)#no ip unreachables
Router(config-if)#ip access-group 101 in
Router(config-if)#exit
Router(config)# interface e1
Router(config-if)#no ip directed-broadcasts
Router(config-if)#no ip unreachables
Router(config-if)#ip access-group 102 in

Again you can look at the access-lists and tell if you are receiving any unwanted traffic by using the "*show access-list* " command.

We now move to the configuration of the Pix firewall. The Pix firewall is setup with three interfaces. The outside, inside and screened interfaces. Each interface will have a security level associated with it. The inside will be the highest level at 100. The screened interface will have a security level of 50 and the outside interface will have a security level of 0. Security level 100 is the highest and 0 is the lowest. The Pix will NOT pass traffic from a lower security level to a higher security level unless explicitly configured to do so. You configure this by building access-lists and applying them to the interfaces. Our Pix is running 5.2.3 code. Earlier versions of Pix did not use access-lists and were much more difficult to configure. In our configuration we will only allow necessary traffic between interfaces and will block everything else. The Pix will use the TACACS+ server for authentication, authorization and and accounting. The Pix will send syslog message to our syslog server at 192.168.0.3.

We will configure the firewall with the interfaces and access-lists as described below.
The inside interface will have a security level of 100 (highest), and have access-list in2out applied.
The outside interface will have a security level of 0 (lowest), and have access-list out2in applied.
The screened interface will have a security level of 50 (middle), and have access-list screened2out applied.

 The inside interface will have access-list in2out applied to it that allows only HTTP traffic to the proxy Server on port 8080, POP3 TCP/110 to the external mail server and TCP/UDP on port 53 to the DNS server. We will also allow telnet to those machines. The inside interface will also allow telnet to the Pix and the Cisco 7507 for maintenance and operations. Users on the inside will use the Proxy server on the screened interface to access the Internet. This will prevent attackers from contact directly with any inside machines. All traffic will be originated from the inside to the proxy server. In the event the proxy server gets compromised it is unlikely that it could infect inside machines. By running host-based security software, such as Axent ESM, we will be able to tell if the proxy server has been tampered with. Axent ESM is a host based application that checks for known vulnerabilities on various platforms, such as Unix, Windows and Solaris. ESM has a tracking feature that will update periodically the vulnerabilities found on a machine and send a report to designated parties. Machines are setup in domains and can have distinct security policies run against a domain or individual machines. More can be found at http://www.axent.com/.

The Pix will only allow telnet from the 192.168.0.7-192.168.0.15 hosts on the inside. We will only allow
HTTP and HTTPS traffic from the Internet on the outside interface to the web sever and proxy server.
Email will be allowed to the mail server and UDP DNS requests will be allowed to the DNS server.
No DNS zone-transfers will be allowed. DNS zone-transfers are an easy way for an attacker to quickly map
your network. This may enable them to use exploits against your machines. We will also allow the 7507
router
to send syslog messages to the syslog server.

access-list in2out permit tcp any host 2.2.2.8 eq 8080
access-list in2out permit tcp any host 2.2.2.9 eq 110
access-list in2out permit tcp any 192.168.0.7 255.255.255.240 eq 23
access-list in2out permit udp any host 2.2.2.7 eq 53
access-list in2out deny ip any any

access-group in2out in interface inside -- This applies the access-list to the inside interface

access-list out2in permit tcp any host 2.2.2.8 eq 80
access-list out2in permit tcp any host 2.2.2.8 eq 443
access-list out2in permit tcp any host 2.2.2.6 eq 80
access-list out2in permit tcp any host 2.2.2.6 eq 443
access-list out2in permit tcp any host 2.2.2.9 eq 25
access-list out2in permit udp any host 2.2.2.7 eq 53
access-list out2in permit udp host 2.2.2.1 host 192.168.0.3 eq 514
access-list out2in deny ip any any

access-group out2in in interface outside – This applies the access-list to the outside interface

access-list screened2out permit tcp host 2.2.2.6 eq 80 any
access-list screened2out permit tcp host 2.2.2.6 eq 443 any
access-list screened2out permit tcp host 2.2.2.8 eq 80 any
access-list screened2out permit tcp host 2.2.2.8 eq 443 any
access-list screened2out permit tcp host 2.2.2.9 eq 25 any
access-list screened2out permit udp host 2.2.2.7 eq 53 any
access-list screened2out permit tcp any eq 23 192.168.0.7 255.255.255.240
access-list screened2out deny ip any any

access-group screened2out in interface screened – This applies the access-list to the screened interface.

You can check the syslog server or the log on the Pix for activity on the Pix. This should be done on
a daily basis at a minimum, with three times a day normal.
You can test the rulebase by connecting a machine inside a run a nmap scan using an illegal source
address. The command would like this :
nmap x.x.x.x –e <the interface to use> -S y.y.y.y –v –sP –PI.
This nmap command would ping the IP address x.x.x.x using the source IP of y.y.y.y. This is called
spoofing your source address. This would check the rulebase to make sure that only the addesses you
specified were allowed.

This may seem like a little bit of overkill, but I don't believe you can ever be to secure.
We secured the Pix about as tight as it can be. We have only allowed the services and ports that
are necessary to operate while blocking all other traffic. This based along with well secured hosts
should help alleviate any known vulnerabilities. Of course the Pix as well as the hosts will
 have to be updated as new attacks become known. You can subscribe to many lists to receive security
updates such as http://www.securityfocus.com/. There you can sign up for the platforms you need. You can
also sign up with Microsoft at http://www.microsoft.com/ and Cisco at http://www.cisco.com/.

We now move to the Lucent Brick. This will terminate our VPN tunnels from partners, suppliers and

remote users. The users will be using the Lucent client that establishes a tunnel between their machine and the Brick tunnel address. The Lucent client will use IPSEC using ESP, 3DES and HMAC-MD5. The Brick will authenticate the users with the Radius server before allowing the tunnel to establish. Once authenticated by Radius the client and Brick will setup the encrypted session. The Brick will use statically defined NAT addresses for the inside machines. At this time the clients will only be able to connect to the database/inventory machine.

So the rules in the Brick will allow radius authenticated users to establish a tunnel. At that point they will only be allowed to access the one machine. As business needs change this policy can be adjusted. The internal users will be allowed to make connections to the outside clients through the Brick using a NAT address and setting up a tunnel. All traffic between the outside world and the internal machines will be

encrypted and inside a tunnel. This should help prevent anyone from "sniffing" the data along any external path the data takes. Figure 2 is a display of the Lucent GUI.

Figure 2

ASSIGNMENT 3: Audit your security architecture

The audit will be conducted in three phases as follows.
Phase one will be to plan the assessment. Phase two will be the implementation of the plan and
Phase three will be recommendations for improvements or alternate architectures.

Phase One:

The first thing we will do is meet with the operations and security people at GIAC Enterprises.
It is important to meet with different groups within the company since they will probably all
have a different view of what the security policy is and how it should be enforced. It might even
help them understand each other's view of security. We will start with a discussion of what business
GIAC is in. How is the operation run. What hours of operation do they have. How many people
have access to the systems that are being protected. What type of data is being used. How is the data
accessible and how do they access it. These are all questions to get an idea of what the operation is.
We will need network diagrams in order to look at the design of their network. Once we know how
their data is used, we can start to evaluate how to protect it.

We will start our tests in a layered fashion. That is, we will test the router, firewall, hosts and Brick
independently. After each layer we will check the logs on the devices.
We would run some tests off-hours so that we don't interrupt their daily business. We would also
monitor some systems during business hours so we can get a baseline of what type and how much
traffic there is normally. While there during business hours we would also observe the physical
security of the systems as well as how the systems are being operated. That is, are workstations left logged
on while employees go to lunch or break, are consoles to servers left on while not being actively used.
These are two very bad practices because it allows anyone access to company data. We will nmap to scan
machines from both inside and outside to look for vulnerabilities. We will load an Axent ESM client on
important machines . We will run Axent NetRecon on the networks to check for vulnerabilties. We will run
password cracking programs to check for weak programs. We will install some IDS machines to watch
traffic patterns.

Phase Two:

To start we will run nmap from outside the network to see what systems are visible. We may have to do
this from another site if physical access is not available on-site to another network. We will use x.x.x.x in
the examples to represent the network IP's. We would run a separate scan for each network. We will run
the following nmap commands:
nmap x.x.x.x-x –sP -v - This is a simple ping sweep to find what machines are up.
nmap x.x.x.x-x –sT –v – This is a TCP connect scan that will discover what ports are open on machines.
nmap x.x.x.x-x –sU –v – This is a UDP scan to check for open UDP ports
nmap x.x.x.x-x –O –v – this is OS fingerprinting. Depending on the results nmap will determine what OS is
running on the machines. Attackers use this to identify machines and then use known exploits against them.

There are many more nmap scans that could be run. Nmap is available at http://www.insecure.org/.

Now we will test the router by trying to telnet to it from all segments. The only successful telnet should be
from the machines that are defined in access-list 1 on the router. Also the user will have to have a valid
account on the TACACS+ server . All others should fail. Try to telnet to different ports on the router such
as 7(echo) and   19(chargen). These should not answer. Try to run a finger to the router. It should not
answer. Try to ping the broadcast address of the networks connected to the router. The "*no ip directed-*

broadcasts" should prevent any response. Try to do some SNMP gets and sets to the router addresses. These should all fail. If not then you need to check the router configuration.

We will then run nmap from the inside. We will run this against all networks. After completion we will check the logs on the machines we scanned to see if they recognized the fact the nmap scanned them. Host based IDS programs should detect certain scans and log them. Axent ESM, TCP Wrappers http://www.porcupine.org/pub/security,BlackICEDefender http://www.networkice,com/Products/BlackICE and BOF http://www.nfr.net/bof are a few.

We will then test the Pix firewall. Try to telnet to the Pix from all segments using different source addresses. The only ones allowed are the internal machines. Try to do SNMP gets and sets to the Pix. Again you should be unable to get a SNMP response. Try to test each rule in the Pix by testing from each segment to the other segments. Use your browser to do HTTP to the wrong machines. Try DNS queries and DNS zone transfers to the DNS machine and the other hosts. An example of a DNS zone transfer would look something like this:

nslookup

>ls –d giac.net

This should fail since zone transfers were not allowed through the Pix. Try to ping the Pix as well as the devices on the other segments. Do this from all segments.

Test the Brick in the same way as the firewall since it is a firewall/vpn machine. Place a sniffer on both sides of the brick to make sure that traffic is really encrypted into and out of the Brick. Try to setup an IPSEC connection from outside using an invalid login. Make sure the tunnel does not come up. Bring up an IPSEC session and try to access machines on the inside that you are not allowed to get to.

Now to test the hosts. First and foremost turn off all unneeded services. Sit down with the administrators and discuss what services are actually used. Next would be to check and ensure that all patches have been installed. Patches can be found at the various vendor web sites. These vendor web sites will also have new security vulnerabilities listed with the patches. This is one of the biggest reasons why systems get compromised. It is IMPERATIVE to stay up to the latest patch levels.

Run an nmap to discover what ports are open. Find out if these ports are being used. I would suggest to load and run Axent's ESM software. It will test for the most common vulnerabilities on hosts of all platforms. ESM is all GUI interface and easy to use. The GUI did not look good in this report so I will omit it. You can get a 30 day trial from http://www.axent.com/. It will also show where to look for information to repair any vulnerabilities. We will also run tests by command line such as trying to telnet to hosts. We will run crack or another password checker to ensure that all accounts have passwords and they are strong passwords. Passwords should be changed regularly.

Crack is available at http://www.users.dircon.co.uk/~crypto.

L0pht Crack is available at http://l0pht.com/.

Install Tripwire on the hosts and check the logs for suspicious activity or if any files have changed. Tripwire is available at http://www.tripwire.com/.

Re-run the scans and attempts to access the hosts and make sure Tripwire picks that up.

Place a sniffer on all segments and look for the top ten talkers on each segment. Verify that they are legitimate. Archive this data over several days so you have a snapshot of the normal traffic in the network.

Log in to various hosts and try to access other hosts. Do traceroutes and make sure data flows along the paths it is supposed to. Try sending use the mail server as a mail-relay to ensure it is not.

Phase Three:

After running the audit against this network check the results wth other security people. Verify that the results are indeed what was expected. After checking this network it was found to be pretty secure. That is, only specific traffic was allowed to the machines . This traffic was validated by checking access-list and host logs.

There are some recommendations that I would make here. I would install permanent IDS systems such as Securitywizards Dragon or ISS RealSecure. I would place one on each segment that has hosts on them. I would load a host based IDS on every host . Axent ESM is great for this. I would assign a person or group

12

to check logs at least daily. These people should understand what traffic is normal on the network. There needs to be a baseline of what is normal on the network. This group should subscribe to different security lists and maintain a list of vulnerabilities. There should be a list of all patches installed and what machines they are installed on. I would also use a router behind the Pix and Lucent Brick to add an extra layer of security. Proper access-lists could be installed there. I would have several sniffers available and periodically check to see who the top talkers are on each segment. Look for abnormalities.

Depending on what users are behind the Pix firewall, I may add additional firewalls to partition off different groups such as accounting, marketing and other groups. This will add another layer of security.

You can never have enough layers. The most important thing I would recommend is that there be an ongoing training program so that employees are aware of the security risks in what they do.