Global Information Assurance Certification Paper

Firewall & Perimeter Protection
SecureZone (Version 3, Patch 2)
Richard C. Hammer
06/05/2000

Environment:
I am a subnet manager for a large organization with a Class B network broken into Class C subnets. My institution protects its perimeter with a firewall, maintains the domain name servers, and external web servers. Computer/Network security requirements differ greatly from subnet to subnet. Some groups require contractors, visitors, and students to complete their mission and want their part of the network as open as possible. Other groups require their part of the network locked down as tight as possible. It is up to each subnet manager to supplement computer/network security as needed, my subnet is one that needs to be locked down.

Task:
Last year during the Y2K frenzy I was tasked with replacing my subnet's aging linux firewall. With limited funds and time, it was imperative that the solution was not only cost effective but must be easy to install and maintain.

Firewall Requirements:
NAT (Network Address Translation)
Support 150-200 Internal Machines
Low Cost
Easy to Install and Maintain
e-mail Header Stripping
Flexible logging capabilities
Vendor must be established

I started my search for a solution at the SANS 1999 Conference in New Orleans. The vendor booths at the conferences were an extremely valuable resource for me. I was impressed with the technical people that Secure Computing, CISCO, and CheckPoint had at the conference. After returning home, I got pricing information on each product and played with demo versions of the products I could afford. I was impressed with Secure Computing's SecureZone firewall (I have heard it called the Baby Sidewinder) after I was able to install and configure the demo version in less than 30 minutes. You can read about the SecureZone firewall at the following link:

http://www.securecomputing.com/index.cfm?skey=79

After talking with Secure Computing's technical support people I convinced myself that SecureZone would meet my needs and I purchased a 200 Node license, basic technical support and the recommended hardware (Compaq DeskPro EN550 MHz platform). I really liked the idea of Secure Computing's SecureOS™ and patented Type Enforcement™ technology.

Implementation Problems:
The hardware and software arrived quickly and I started to build my new network. Once again I was impressed with how quickly I was able to get the SecureZone up and running with a basic configuration. Problems did not start to arise until I implemented my internal and external e-mail servers. During the early testing stages I noticed that the SecureZone e-mail server features did not function properly. After calling tech support it was acknowledged that there was a bug in the distribution and that it could be fixed quickly by letting them log into the system remotely. This was a real problem for me because I sit behind another firewall and my institution would not allow Secure Computing's technical support access and would not allow me to install the system outside the institutional firewall. I was not able to make the change myself because SecureZone does not allow local users access to the command prompt and Secure Computing will not allow end users to have the program that can access the command prompt remotely. Not being able to access the operating system has become one of my biggest complaints about the SecureZone firewall. I solved the problem by taking the system home and installing it on my home network. After the Secure Computing technicians did their magic (e-mail settings from generic proxy to server proxy) the SecureZone e-mail server started to forward e-mail correctly. I then noticed that the e-mail header stripping did not work properly (Internal IP address and DNS name left in e-mail header). After messing around with this for several days I gave up and I contacted Secure Computing's Technical support line again. I sent them my configuration and was notified that the header stripping feature does not work in certain cases. It was suggested that if I really needed this feature to work that I should have purchased the Sidewinder. The Y2K deadline was approaching fast and I had already invested a lot of time in SecureZone. It was just to late to change products.

Things I really like about the SecureZone Firewall:
-Cost (Under $10K)
-Easy to configure (As long as you stay within the supplied proxys)
-External interface can be configured to not reply to a ping.
-NAT works well
-Reliable – Has run without crashing for 6 months.
-Technical support is pretty good. They might not be able to fix the problem but they will work hard trying. I found them friendly and skilled.
-Does not appear to slow down network access.

Things I don't like about SecureZone Firewall:
-Can't access operating system (If you can't do it with the GUI it can't be done)
-E-mail header stripping does not work in all cases
-No auto-backup feature
-No ability to look at log files after they roll over.
-No NTP. Can't synchronize time with a timeserver.
-Auto e-mailing log files feature does not work when logs become large. I know my e-mail servers will handle the files. I have ftped the log files to my internal machine and e-mailed it through the firewall to my external e-mail server.
-Alerts do not work on single event. I set an alert for incoming ssh connections, but have never been alerted and never seen an alert in the alert log file.

Picture of my Perimeter Protection Scheme

\\

Implementation Solutions:
Except for the e-mail Header Stripping Problem I have been able to get around all these weaknesses using other systems. I run a system inside the firewall that logs all outgoing connections and I log all incoming ssh connections on the internal authentication server (with correct time stamp). I have snort installed on a machine outside the firewall looking for intrusions and logging strange incoming connections. I use SessionWall-3 to help enforce our policy and supplement internal logging. If you are planning to install a SecureZone firewall you better be prepared to use other systems to supplement your logging functions, the lack of NTP and logging flexibility requires this approach. Maybe the cost of a Sidewinder is not so high!

Policy and Network Access:
The SecureZone default configuration comes with all incoming and outgoing ports closed. During the install you can choose to open the most popular outgoing ports (telnet, smtp, www, ftp, etc.). I chose this option just to get a feel for the GUI "look and feel (No Real Syntax)."

In my particular situation I treat everything between my firewall and the institutional firewall as my DMZ. Anything I want the world to see can be put on the institutional servers outside their firewall, everything I want the institution to see I put on my servers outside the SecureZone, and anything I need to restrict access to is served up inside the SecureZone firewall. This allows me to have a very restrictive Ingress policy. I only allow ssh, ftp, and e-mail into our internal network.

I think most of us concentrate on what we allow into our networks and forget that outgoing connections pose a significant danger to our internal networks. SecureZone forces you to think about outgoing connections. Ports are closed on the SecureZone firewall unless you manually open them and there is no easy way to open all outgoing ports or even a block of ports. At first this annoyed me but it has turned out to be a blessing. The fact that you must manually open outgoing ports saved us when the PrettyPark worm hit us this year. No one in my group had requested outgoing tcp port 6667 (IRC) be opened, so our firewall stopped infected machines from connecting to the internet. See my paper at:

http://www.sans.org/infosecFAQ/prettypark.htm

My policy on Egress filtering is to open only outgoing ports required for members of my group to complete their work. You will be surprised how few outgoing ports must be opened for real work related activities. At this time I have less than 75 outgoing ports open. Users can request to have a port opened by sending me a copy of the link they need to access. In the following example I would need to open outgoing tcp port 6002:

http://www.neededlink.org/FAQ/index.htm:6002

During the initial install I took a best guess at what outgoing ports should be opened. The first month I spent a lot of time opening outgoing ports that users needed. After that first month the outgoing rule set became fairly stable. To reduce the threat of your machines being used as attack or amplification machines restricting outgoing connections is a must. There are some very good papers on the web that discuss Egress filtering.

Chris Brenton's paper:

http://www.sans.org/y2k/egress.htm

Cisco's Anti-spoofing paper:

http://www.cisco.com/warp/public/707/21.html#spoofing

Defense against DDos attacks:

http://www.sans.org/dosstep/index.htm

http://www2.axent.com/swat/news/ddos.htm

http://www.icsa.net/html/communities/ddos/index.shtml

Be a good network citizen and filter outgoing connections!

Applying a SecureZone Filter:
With SecureZone there is no real syntax, everything is done using a graphical interface. SecureZone uses the concept of regions. Each interface (Ethernet cards in my case) has a region associated with it. I only have two regions (Internal and External) plus the firewall region. To view the regions click on "GO", then select the "Regions".

The following is a snapshot of my regions:

Note: The IP addresses have been doctored to protect our network address space. I will use 172.16.16.11 as the IP address for the internal interface.

To apply a rule, you open the "Network Access Page" and a flowchart of the network rules will be displayed. On my network I have four access rules, sshin (ssh into the firewall), E-mail (SecureZone E-mail Server), FTP-in (restricted source IPs, redirected to single internal machine and limited to certain users), and outgoing rules.

The following is a snapshot of the network access page with "outgoing rules" selected:

The yellow shaded areas are the parts of the flowchart related to the "outgoing rules". To view the flowchart for "sshin" you would double click on "sshin" and the flowchart associated with that rule would be displayed with yellow shading. The first node of each rule flowchart is the name of the rule. The second node of each rule flowchart is the service(s) that the rule supports. You can have one or many services in a rule. "sshin" has only one service associated with it (ssh- tcp port 22). "outgoing rules" has ~75 services associated with it. The next part of the rule is the "From-To" node. This is the part of the rule that shows which direction the rule flows. "sshin" flows from the External region to the Internal Region (Ingress) and the "outgoing rules" flow from the Internal region to the External region (Egress). To display the services that are associated with a rule, double click on the "Services" node of the flowchart. The services associated with the rule are listed under the "Selected" column. SecureZone comes with many services already defined. SecureZone understands services like e-mail and ftp and supply special proxies or servers for these services.

If the service you need is not listed you can create a new service by clicking on the "New Service" button. You will need to supply and name, port, and protocol (tcp or udp). Once you create a new service all you need to do is find it in the "Available" list and click on the "Add" button. Note: once you create a new service there appears to be no way to remove it from the "Available" list.

The next thing you need to do is select the direction you want the network traffic to flow. For "outgoing rules" I selected Internal to External. For rules like e-mail the SecureZone has built in servers. For my "E-mail" rule I selected "Internal to Firewall" and "External to Firewall" since e-mail is allowed to flow in both directions. The network flow for rules "sshin" and "FTPin" is "External to Firewall", with destination redirects to internal authentication servers. When you double click on the "From-To" node of the flowchart you can configure the network flow for the rule.

"From-To" node for "outgoing rules"

SecureZone allows you to rewrite source and destination IP addresses as well as redirecting packets to different IP addresses. In my outgoing rule set I rewrite the original internal IP address with the address of my firewalls external region IP address. To do this with SecureZone you add a "Rewrite SRC" node to your flowchart. When you double click on the "Rewrite SRC" icon you have the option of rewriting the original IP address with whatever you want the external world to see.

"Rewrite SRC" for "outgoing rules"

You will notice that I rewrite all internal IP addresses with 172.16.16.15 (not the real external IP) and leave the ports alone. I have never tried to rewrite the original source port so I really don't know if that works. An "*" is a wildcard so you can use it for all internal IP addresses and all ports. The check mark allows the connection (X denies connections). There are other options that I have never used that can limit things like time of day, maximum sessions, users, and type of authentication. In theory you could add alert nodes "!" to your flow charts and you would get an alert anytime the rule hits the alert. I have an alert node on my "sshin" rule but have never gotten an alert. To save changes that you have made to a rule you need to double click on another rule and you will be prompted to save your changes.

Testing Rule Sets:
It is very important to make sure that any change you make does what you want it to do and does not create any unwanted side effects. Making sure that the change you made works is easy, all you need to do is access the service that you tried to open. If you can access the service it works! Making sure you have not comprised the security of your firewall is a trickier situation. When I make changes to my firewall rule sets I run through a series of tests to make sure everything else stayed the same.

I have a honeypot outside my firewall that should never see network traffic. If I want to verify that a port on my firewall is closed I try and access that port on my honeypot from inside the firewall. If the honeypot sees the traffic I know the port is open. A simple example would be to telnet to the honeypot using the port number switch.

Note: IP addresses are not real – 172.16.15.1 (External interface of firewall) and 172.16.15.23 (honeypot).

Example for closed port:

```
[internal machine]$ telnet honeypot 8010
Trying 172.16.15.23...
telnet: Unable to connect to remote host: Connection refused
[internal machine]$
```

Looking at the snort alert file I see no connections made to the honey pot so the "Connection refused" message came from the firewall.

Example for open port:

```
[internal machine]$ telnet honeypot 8000
Trying 172.16.15.23...
Connected to honeypot
Escape character is '^]'.
Connection closed by foreign host.
```

The snort alert log shows the connection to the honeypot was made!

[**] tcp traffic to honeypot [**]
06/01-11:35:19.819088 172.16.15.1:8414 -> 172.16.15.23:8000
TCP TTL:64 TOS:0x0 ID:57507 DF
S*** Seq: 0x5E36E393 Ack: 0x0 Win: 0x2000
TCP Options => MSS: 1460 NOP WS: 0 NOP NOP TS: 12893795 0

You could test all outgoing connections by running a port scanner and use the honeypot alert file to map all open connections. A similar concept can be used to check incoming connections. Have a machine on the inside of the firewall running snort and keying in on connections that make it into the firewall from a know source IP. Running NAT requires you to key on connections coming from an external machine because you would not be sure where the traffic is being redirected inside the firewall if something is broken.

I have Sessionwall-3 installed inside my firewall and use it to monitor incoming traffic. I find Sessionwall-3 to be a very good tool for enforcing my security policy. Sessionwall-3 comes with basic intrusion detection, virus detection and url monitoring/filtering. It displays information in a readable form and unlike Securezone it has good logging and report generating capabilities. You can also focus Sessionwall-3 to look at a specific host or server. I have ISS Internet Scanner installed on a machine outside my firewall and scan the external IP address of the firewall and focus SessionWall-3 to view on all traffic that it sees from the scanning machine. The only traffic I have ever seen make it inside my firewall from an external scan is port 22. The ftp and e-mail port scans do not make it past the SecureZones internal interface.

Sessionwall-3 snapshot showing ssh incoming traffic (IP addresses and names doctored):

Since SecureZone's logged features are so weak I have come to reply on Sessionwall-3 and snort for enforcing my security policy.

You can find out more information on these programs at the following links:

Snort:
www.snort.org

SessionWall-3 (now called eTrust Intrusion Detection):
http://www.cai.com/solutions/enterprise/etrust/intrusion_detection/

ISS Internat Scanner:

http://www.iss.net/customer_care/resource_center/product_lit/

Firewall Policy Violations:

Recognizing network traffic that is not normal to your environment is an important aspect of administering a firewall. I want to know about any traffic that is not expected or normal. I have mentioned the SecureZone logging weakness numerous times and will show you a couple cases where it just does not work at all. Please note that all the IP addresses have been changed. I will use 172.16.15.1 as the IP address for the external interface on the firewall.

Example 1 - telnet attempt into the firewall:

Here is an alert that my honeypot logged the other day. The honeypot is setup to alert on undesired network traffic to the firewall. In this case someone tried to telnet into the firewall. It turns out that the user of this system just forgot and used telnet instead of ssh. I searched the SecureZone log files for this attempted connection and found no log entry. I can only assume that the SecureZone logs only connections that succeed and drops the others. I double-checked my Sessionwall-3 log to make sure that no telnet session from this client entered the firewall. Successful telnet connections into my internal systems would open up all sorts of security concerns, passwords being sniffed along the wire, full access to our internal network, and most of all the firewall policy or function did not work properly. It is reassuring to know that telnet sessions are being stopped at the firewall entrance.

Snort Alert message showing the telnet attempt:

```
[**] telnet to firewall [**]
05/29-08:19:13.901209 172.16.12.231:1135 -> 172.16.15.1.:23
TCP TTL:126 TOS:0x10 ID:57604  DF
**S***** Seq: 0x113E72   Ack: 0x0   Win: 0x2000
TCP Options => MSS: 1456 NOP WS: 0 NOP NOP TS: 0 0 NOP NOP SackOK
```

The snort rule that caught this follows:

```
Alert & Log         Destination IP
  |     Any source IP          |    Destination Port
  |          |                 |           |
alert tcp any any ->  172.16.15.1/32 23 (msg: "telnet to firewall";)
       |          |                    |                 |
       |    Source Port                |    message to print in log files
    tcp only                      Netmask
```

Example 2 – Ping scans

One of the things that I like about the SecureZone firewall is the ability to configure the external interface to not reply to a ping. For that matter it will not acknowledge any connection to a port that is not opened. This makes it very difficult for a hacker to locate my network entrance with scanning tools. There is no reason that I can think of for anyone to ping my firewall except me. Once again I relied on my honeypot to monitor the firewall for alert information. The rule is easy to configure, log any ICMP_ECHO that is directed at the firewall's external interface. I tried logging all ICMP that came to the interface but that gave me too many false positives since ICMP_ECHO_REPLYS come back to the firewall when someone inside the firewall pings an external machine.

Snort alert log entry. It turned out to be an institutional scan of my subnet.

[**] icmp echo to fire [**]
05/01-16:50:48.557109 172.16.100.101 -> 172.16.15.1
ICMP TTL:64 TOS:0x0 ID:61755
ID:8477 Seq:0 ECHO

ICMP snort rules that I use:

```
Alert & Log          Destination IP                  message to print in log files
  |     Any source IP          |     Destination Port              |
  |           |                |           |                       |
alert icmp any any -> 172.16.15.1/32 any (itype: 8 ;msg: "icmp echo to fire";)
      |           |                |               |
      |      Source Port           |     itype=8 is ICMP_ECHO_REQUEST
 icmp only                    Netmask
```

Note: You need to include any for the ports even though ICMP does not use ports.

Note 2: I also have ICMP rules for redirect (itype=5), timestamp request (itype=13), information request (itype=15), and address mask request (itype=17).

Ping requests from strange IP addresses are usually a sign that something bad is going to follow. It is certainly an easy way to map a network. I try to get as much information on the source of any ping request that comes to my firewall or my honeypot.

Example 3 – ssh into firewall.

I like to keep track of the ssh connections that come into the network. I setup SecureZone to alert me when I get an ssh connection coming into the firewall but it has never worked. You can see the alert "!" in the following SecureZone flowchart rule:

I have talked with Secure Computing's technical support people about the alert problem

and there seems to be a problem logging single events as alerts. They had intended it to be used only for repeated events or attacks. Needless to say I rely on Sessionwall-3 and the authentication server log files to keep track of incoming ssh connections. I am interested in both successful and unsuccessful connections. If I see ssh connections from a user that I know is onsite it might be something I need to investigate. Repeated unsuccessful connections coming from the same machine or by the same user is also something to investigate. Earlier in the paper I showed you how SessionWall-3 displays ssh connections so I will concentrate on the authentication server's log entry here. I have a script that runs every night that sends me e-mail with all ssh connections and attempted connections. This combined with SessionWall-3's visual display allows me to keep track of incoming connections. I do a similar thing with the incoming ftp server.

Here is what a failed connection looks like. The user must have mistyped his/her password since I saw a successful ssh connection less than a minute later.

May 1 16:56:11 aserver sshd[26067]: log: Connection from 172.16.9.100 port 1023
May 1 16:56:16 aserver sshd[26067]: fatal: Connection closed by remote host.

Knowing what connection are expected is a benefit of being a subnet manager. I think once you get more than 200 systems or users on your network you have a hard time knowing what connections are expected. I usually know who is onsite and who is on travel. That helps me to decide if a connection is reasonable or not.

Example 4

I had mentioned earlier that the SecureZone had protected our internal network from the PrettyPark worm. The combination of having the outgoing tcp port 6667 closed and SessionWall-3's visual representation allowed me to catch and notify my institution that PrettyPark was a real threat. Once again, no entries the SecureZone log showed that there was a problem. The SecureZone did log the connection once I opened the port up to get the signature of the outgoing connections. I will not talk anymore about this firewall policy breach here. You can read about this at the following link:

http://www.sans.org/infosecFAQ/prettypark.htm

Example 5 – Personal firewalls

I know this is a perimeter protection paper but I would like to give a plug for personal firewalls and how they enhance security. If you do not have outgoing proxies on your firewall you are vulnerable to Trojans connecting your internal machines to the internet without your knowledge. Personal firewalls can reduce this threat. I use a personal firewall program both at work and at home. These programs will prompt you when a program on your system tries to access the network and blocks computers trying to access ports on your system. These are handy features that enhance security. If a Trojan infects your system you and your firewall administrator might not notice the outgoing connection. The personal firewall gives you a chance of knowing. I have to admit that

my personal firewall has never gone off at work, but it is constantly blocking connections to the PCs on my home network. The personal firewall program that I use at home is ZoneAlarm. It works pretty well and it is free.

Here is a message that my home ZoneAlarm firewall recorded:

Time: 5/29/00 15:52:46
The firewall has blocked Internet access to your computer (POP3) from 206.182.235.227 (TCP Port 3591).
This was definitely a port scan. I looked in my home server log and found the same scan.

/var/log/messages entry from my home server:

messages:May 29 14:55:23 homeserver popper[10855]: @[206.182.235.227]: -ERR POP EOF received
messages:May 29 14:55:23 homeserver popper[10856]: @[206.182.235.227]: -ERR POP EOF received

I did a reverse nslookup on the IP address and the machine was not listed. I then started messing around trying to find out who owned this machine. I finally got lucky and telneted to this machine's port 25. Send mail was very cooperative and told me the systems DNS name. A forward nslookup on the machine name gave the correct IP back. I sent e-mail to root@strangemachinename.com and it was kicked back with the following message: "Sorry, no mailbox here by that name. vpopmail (#5.1.1)". I then opened my browser and loaded the page served up by that IP address and sent e-mail to the webmaster, asking "why they scanned the pop port on my PC." Needless to say I have not heard back from them. I will leave finding the DNS name of this machine to you.

You can find ZoneAlarm at the following link:

http://www.zonelabs.com/

Final Comments on SecureZone:

I understand that Secure Computing is going to discontinue support for the SecureZone firewall next year. I think they already have! This is a real shame I think with a couple small improvements SecureZone could find a real niche in the market place. It is extremely easy to configure, reliable and SECURE. Adding a command prompt will fix most of the things that I have complained about. You could run NTP and you could use scripts or swatch to maintain your logfiles. You could setup cron jobs to backup your files and/or move them to other machines. The header stripping is supposed to be fixed in patch 3 (which is now available). With swatch you would not even mind that the alerts do not work for single events. I would REALLY like the SecureZone firewall if these things were fixed!

Scenario #1

Submit a detailed design for a site with dual connections to the Internet that is optimized to be resistant to DDOS attack. Include a description of the hardware and configuration. A drawing is a requirement for this assignment. Please keep in mind that the main goal of this assignment is to allow you to demonstrate what you have learned in the course, there may not be a "perfect" answer to this problem.

\\\

My Solution:

Description:

It is unclear from the description of the scenario if I have control of the border router(s). If I control of the border router configuration I would configure them to not reply to ICMP ping requests and block all other ICMP requests coming from the internet. If the router supports extended access lists, I would not allow incoming tcp and upd packets for services that I know are not needed for my network. The router will be configured to only allow outgoing packets that have valid source IP addresses for my internal IP space and will not route reserved source IP addresses or private (RFC1918) address space. Source routing will also be denied. This should reduce the chances of my network systems being used as attack or amplification devices.

2 Nokia Firewall-1 systems will be installed as the firewall systems. I will implement VRRP (Virtual Router Redundancy Protocol) and Firewall-1's synchronization feature for added redundancy. Firewall-1 allows external Web and DNS requests to be redirected to specific web and DNS servers located inside the firewalls. E-mail will be redirected to the internal e-mail server. These servers will be hardened, with all other services removed. Anything I want the world to access will reside on these hardened machines. The external DNS server will only supply records that the outside world needs to know about the network. I will configure my network IDS to flag all traffic that I do not expect to see on these two machines. The IDS will also act as a honeypot. No services are being run on the honeypot except the IDS. It can be assumed that any traffic going to the honey pot is suspicious. I will treat these machines as my DMZ for this scenario.

The internal DNS will supply information on all other machines in the network. The

internal WWW server will serve up internal web pages. Access from the internet will require strong authentication through the firewalls. I will not allow DNS zone transfers through the firewalls or between DNS servers. I would only open needed outgoing ports on the firewall. Depending on the level of security required on the internal network, I might implement outgoing proxies with authentication for access to the internet. Another internal appliance firewall might also be desired to protect machines processing very sensitive information. Personal firewall programs will be installed on internal PCs processing sensitive information to reduce the threat of Trojans connecting internal machines to the internet.

Scenario #2

A site has two critically important internal subnetworks, research and accounting, that require a high degree of protection. The site is connected to the Internet. An employee that has since left secured budget approval for one Cisco router, one proxy firewall and two appliance type firewalls with 2 10/100 NIC's, capable of performing in a bridging nature (similar to SunScreen), and this equipment has been ordered and has arrived and cannot be sent back. Submit a detailed design for the most effective protection. A drawing is a requirement for this assignment.

My Solution:

I would configure the border router to not reply to ICMP ping requests and block all other ICMP requests coming from the internet. If the router supports extended access lists, I would not allow incoming tcp and upd packets for services that I know are not needed for my network. The router will be configured to only allow outgoing packets that have valid source IP addresses for my internal IP space and will not route reserved source IP addresses or private (RFC1918) address space. Source routing will also be disabled. This should reduce the chances of my network systems being used as attack or amplification devices.

The DMZ for this system is the region between the border router and the proxy firewall. The external DNS and web servers are installed in the DMZ. Anything I want the world

to access will reside on those machines. They will be hardened with all unneeded services and programs removed. I will configure my network IDS to flag on all traffic that I do not expect to see on these machines. The IDS will act as a honeypot as well. No services are being run on the honeypot except snort. It can be assumed that any traffic going to the honey pot is suspicious.

Internal company DNS, WWW, e-mail, and common servers will reside in the region between the Proxy firewall and the appliance firewalls. This is where all shared resources will be served up. The scenario does not say whether the two groups share a common e-mail server of have separate servers. With my solution you can choose either situation. I would serve up the company only web pages with the www server in this region. Other common servers might exist in this region but they should not include extremely sensitive data. Access from the internet will require strong authentication through the proxy server. If real tight security is required I would implement outgoing proxies with authentication to access the internet. I will not allow DNS zone transfers through any of the firewalls.

I would keep the research and accounting networks completely separate with very limited access between the two. I would use NAT on the two appliance firewalls so I could hide the address space for the two protected network. I would tightly restrict services going into each of the appliance firewall areas, maybe only shh and e-mail. File transfers between the two subnets can be done using servers inside the common region or via e-mail. I would only open needed outgoing ports on the two appliance firewalls. I would require personal firewall programs to run on all internal PCs to reduce the threat of Trojans connecting machines to the internet.

Scenario #3

A site has two internal subnetworks, research and facility. The research subnet requires a high degree of protection. The facility subnet supports a large number of visitors, researchers, graduate students, and temporary contract employees. These people need remote access to the internal network when they are off site. Strong authentication is not an option since the visitors come from many different organizations and colleges. The facility network needs to be treated much like a University network. All sensitive computing is done in the research subnet and data security is the number one network priority. The facility network needs to be as secure as possible while still allowing visitors remote access to the facility systems. It is your job to design a network that meets the needs of both networking environments.

My Solution:

\\\

I would configure the border router to not reply to ICMP ping requests and block all other ICMP requests coming from the internet. The router supports extended access lists, and would be configured to deny incoming tcp and upd packets for services that are not needed for either environment. The router will be configured to only allow outgoing packets that have valid source IP addresses for the internal IP space and will not route reserved source IP addresses or private (RFC1918) address space. Source routing will also be disabled. This should reduce the chances of my network systems being used as attack or amplification devices for hackers.

I would install a Checkpoint's Firewall-1 at the entrance to my network. The DMZ for this site is the region between the border router and the Firewall-1. I would install the external DNS server and external web server in the DMZ. Anything I want the world to access will reside on those machines. All unneeded services and programs will be removed from these hardened systems. I will configure my network IDS to flag on all traffic that I do not expect to see on those two machines. The IDS will act as a honeypot as well. No services are being run on the honeypot except snort. It can be assumed that any traffic going to the honey pot is suspicious.

Internal DNS, WWW, e-mail, and common servers will reside just inside the Firewall-1. This is where all facility wide resources will be served up. The facility e-mail will be stored on the e-mail server in this region; research e-mail will be forwarded into the research e-mail server. I would serve up the facility wide web pages with the www server in this region. Other common servers might exist in this region but they should not include sensitive data. Access from the internet will be via ssh and file transfers can be made using scp. This should allow facility users remote access to their data and facility systems if they maintain a valid account.

I would install a Secure Computing Sidewinder firewall on the entrance to the research subnet. I will treat the facility network as an untrusted network. Remote access will require ssh with strong authentication. Only ssh and e-mail will be permitted into the research subnet from outside. NAT will be used to hide the address space for the research subnet. A network IDS will be installed outside the research firewall. It will flag all services going in but ssh (port 22) and e-mail (port 25). The IDS will act as a honeypot as well. No services are being run on the honeypot except snort. It can be assumed that any traffic going to the honey pot is suspicious. All outgoing services will go through a proxy and require authentication and personal firewall programs will run on all PCs. This will reduce the threat of Trojans connecting research machines to the internet. All research workers will get their e-mail on the internal research e-mail server. E-mail coming from outside will be forwarded into the research e-mail server. The internal www server will serve up research only web sites.

Split DNS will be used for the research network. The facility DNS server will have no information about the research network. The DNS server in the DMZ will only have

information on servers that outside users need for access into the network and the DNS server in the facility subnet will have information on facilty servers and clients. DNS zone transfers will not be allowed through any of the firewalls.

Scenario Solving Conclusions:

It is fun to solve network scenarios, but without intimate knowledge of a site's computer environment they are impossible to solve correctly. The solutions can be drastically different depending on funding, level of support expertise, user profile, management, remote access needs, and sensitivity of data being stored. The real value in solving scenarios is going to the vendor sites and finding out what type of solutions are available. The real challenge is trying to match a product to an organizations ability to fund and maintain the solution.

Here are some sites that I visited while researching these scenarios:

Nokia and Checkpoint's Firewall-1 Product at:

http://www.checkpoint.com/products/firewall-1/index.html

http://www.nokia-networks.com/nn__firewall_software.shtml

Secure Computing's Sidewinder firewall at:

http://www.securecomputing.com/index.cfm?skey=232

CISCO products at:

http://www.cisco.com/

Axent products at:

http://www.axent.com/Axent/Public/

SLM Software Inc. (formerly MilkyWay Networks) at:

http://www.milkyway.com/home.html

Global Technology Associates, Inc. at:

http://www.gta.com/Pages/products.html

A couple of great Network Security Sites:

http://www.cert.org/
http://www.sans.org/
http://www.icsa.net/
http://www.denialinfo.com/
http://home2.freegates.be/bchicken/index2.html
http://www.pimmel.com/deepfiles.php3
http://www.l0pht.com/