Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

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

SANS SNAP
GIAC Firewall and Perimeter Protection Curriculum
Practical Assignment for SNAP San Jose May 8-13, 2000

© SANS Institute 2000 - 2002, Author retains full rights.

Submitted by: Jeff Behm
June 15, 2000

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

Assignment (Copied from SANS.ORG web page)
This track requires four short practical assignments. Please check your spelling and read through your wording, this is how the world will see you and we will not accept a second submission. You will be graded largely on the accuracy and educational value of your submission, but also on its appearance.

Assignment 1 – Egress filter - 10 Points
Write a one page tutorial on the reasons for or against egress filtering. Be certain to include the following:

- Syntax of the filter
- Description of each of the parts of the filter
- Explain how to apply the filter
- Explain how to test the filter

Assignment 2 – Firewall policy violations - 50 Points
List five violations of your site's firewall policy. For each log file detect:

- Show the log entry with the violation, explain all fields in the detect with a key
- Describe the rule that caught the violation including explaining the rule
- Explain the potential damage if the firewall had not stopped the attack

If your site does not have a firewall or perimeter defense system yet, or you are not allowed to use your site's logs then you may do one or more of the following.

- Select Linux firewall (such as IPChains) detects from the GIAC web page. You may purchase the book Linux Firewalls by Ziegler or use the Linux firewall web page. Now armed with this information, you should be able to complete the assignment.
- Build a firewall on a test net and attack it. There are limited time licenses on commercial firewalls for you to try. There are also freeware firewalls such as Linux firewalls.
- Use a friend's network.

Assignment 3 – Defense in depth architecture - 10 Points each
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.

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.

Assignment 4 - Create a test that demonstrates your knowledge of the subject area - 20 Points
Develop a scenario that must be solved similar to the two above and submit both your question and your answer. This assignment will be scored primarily on three factors:

- Does the submission demonstrate the student's knowledge of the subject area, so pick a problem that let's you flex your brain muscle!
- Is the solution to the problem accurate
- Is the solution well researched and list URLs, references and resources

A drawing is a requirement for this assignment.

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

Assignment 1: Discussion of reasons for egress filtering.

The case that will be studied is that of outbound http filtering. A particular company has decided that it would be in their best interest to begin to filter out "inappropriate" web sites. Additionally, their Internet bandwidth usage has skyrocketed lately and they wish to increase throughput. They have asked their IT group to come up with a solution to both problems. One solution provided was that of a web proxy/web cache server with appropriate third party URL filtering software. The problem with having the proxy/cache server inside the firewall is that users could still point their browsers directly to the firewall rather than to the proxy/cache server and effectively bypass the URL filtering software. Hence, the egress filter provides a solution to this "bypass" problem. The IT group placed an "allow outbound http access from the internal trusted networks **only to a single destination**" rule into the firewall, effectively disabling http access to the Internet (see rule 2 below). Now the problem is that even the proxy/cache server cannot perform http to the Internet through the firewall. A second rule was added to be evaluated BEFORE the "allow only one destination" rule that allows the proxy/cache server to perform http to any destination (see rule 1 below). This scenario allows the proxy/cache server to utilize its internal cache to return as quickly as possible the most actively hit pages, and only if necessary access the Internet and fetch web pages on behalf of the users. Additionally, since all users are "forced" through the proxy server, this is a perfect location to enforce the filtering of inappropriate web sites with the third party URL filtering software. Both issues regarding filtering of inappropriate sites and increasing web performance are addressed through the use of egress filtering.

Syntax of the filter

This filter is really a two-part filter. The two parts must allow the proxy server web access through the firewall while preventing all other internal machines Internet web access through the firewall. Each part of the filter contains both source and destination rules.

Source rules

Rule number	Network Source	Access	Services
1	proxy/cache server	permit	http-bypass
2	all internal networks	permit	http

Destination rules

Rule number	Services	Access	Network Destination
1	http-bypass	permit	all network destinations
2	http	permit	one internal machine

Description of each part of the filter

Part one of the filter is a firewall source rule that creates a single service (http-bypass) that includes only the proxy server and permits access to the http service, with a corresponding firewall destination rule that allows access to any destination. Part two of the filter is also a firewall source rule that creates a single service (http), includes all the internal networks of the organization, and allows access to the http service. It has a corresponding destination rule that allows access only to an internal server (The "meat" of the egress filter), but has no provisions for allowing access to external addresses for web browsing.

How to apply the filter

Apply the filter by first creating the appropriate network sources (I.E. the proxy server, and the internal networks). Then create two services, http and http-bypass. Create two source rules that associate the proxy server with the http-bypass service and associate the internal networks with the http service. Create two destination rules that associate the http-bypass with any destination address and associate the http service with only a single internal destination. The order is very important here as the source rule for the proxy server must be evaluated before the "all internal networks" source rule, as the proxy server is on the internal network and would be denied access if the rules were in the wrong order.

How to test the filter

Configure your Internet web browser to use the proxy server and web access should be allowed. Change your web browser's configuration to use the firewall as your proxy server and web access will be denied (except perhaps to the one internal machine that was identified in Destination Rule number 2 - in any case it won't allow external web access)

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

Assignment 2: Five Violations of site's firewall policy.

Violation 1: Email relay attempt

Note: IP addresses and hostnames have been changed to protect the innocent, but all were public, routable addresses.

Line number	Date Time	hostname	process name	process id	Text of violation
1	Jun 5 10:56:22	firewall	smap	[22176]:	connect host=nodnsquery/192.168.0.4
2	Jun 5 10:56:40	firewall	smap	[22176]:	relay reject connection-remote: jeffb@abc.net connection unknown/192.168.0.4 reply 550

Analysis:
This entry comes from the Internet firewall and is an attempt to relay email from an external source to another external source using the firewall as the relay server.

Line 1 shows a remote host (192.168.0.4) successfully connecting to the smap (email) process (22176).
Line 2 shows a smap rejecting the relay attempt for the offender, jeffb@abc.com, with ip address of 192.168.0.4 and a reply of 550, relay not allowed.

Description of the Rule:
The rule that blocked this is the firewall anti-relay feature. It prevents unauthorized use of the site's Internet mail service to perform relaying activities. Mail relaying is when someone connects to your mail service and uses it to send mail to other sites outside your own network. Although not desirable in this policy, allowing relaying may be desirable, as in the case if one was hosting multiple email domains. The firewall will check the connected domain to verify if relaying is allowed or denied against the firewall policy.

Potential Damage:
Potential damage may come in the form of Denial of Service for your own site, or perhaps DoS to a downstream victim site, if one were to utilize "mass mailers" and attempt to flood your "open" email relayer with multiple, large emails. A DoS attack might still be attempted to cripple your email service, even if you do not allow relaying. At least in this case, though, your mail service would simply return "relay rejected" rather than having to process many potentially large attachments on each piece of mail.

Additionally, public relations of your site may be hurt, since the email received by downstream victims may appear to be coming from your site. The downstream victims may not have nice things to say about your site, because you are being a poor "net neighbor" due to the allowing of indiscriminate mail relaying.

© SANS Institute 2000 - 2002 Author retains full rights.

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

Assignment 2: Five Violations of site's firewall policy (Continued).

Violation 2: Improper outbound http request (attempt to go directly through firewall)

Note: IP addresses and hostnames have been changed to protect the innocent, but all were public, routable addresses.

Line number	Date Time	hostname	process name	process id	Text of violation
1	Jun 5 11:49:34	firewall	http-gw	[1233]:	deny host=nodnsquery/192.168.174.4 destination=192.168.7.80 ID=123340107

Analysis:

This entry comes from the Internet firewall and is an attempt from an internal network machine to use the http proxy on the firewall.

Line 1 shows an outbound http request being denied. Internal host is 192.168.174.4 and external (destination) host is 207.188.7.80.

ID=123340107 shows the internal identifier the firewall uses to track proxied connections.

Description of the Rule:

The rule that blocked this is the no-http rule implemented in the firewall to prevent internal machines from using the firewall directly for http requests. All http traffic directly to the firewall is denied, except for traffic originating from the internal proxy server. This rule is in place to "force" internal machines to utilize the proxy server, who in turn makes the http request (successfully) through the firewall. Why would one "force" use of the proxy server? Because, in this case, the proxy server provides an http caching function for increased performance on frequently visited sites, as well as a filtering function for those sites deemed inappropriate for business consumption by management of the company.

Potential Damage:

Potential damage comes in the form of unnecessary use of resources. If the cache server was not in place, each http request would have to be sent out to the Internet and results obtained and passed back, through the firewall, Internet router and the Internet itself. In our case, some of the http requests are filled immediately via the cache server and the request never even reaches as far as the firewall. Additionally, the management of the company believes that "blocking" of certain inappropriate sites helps protect them from lawsuits brought about by employees claiming the company "did nothing" to try to stop the inflow of this information. If the blocking was not performed, then management feels the potential for lawsuits is greatly increased.

© SANS Institute 2000 - 2002, Author retains full rights.

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

Assignment 2: Five Violations of site's firewall policy (Continued).

Violation 3: Attempt to access internal DNS servers from Internet (and the internal DNS servers are not public, external DNS's)

Note: IP addresses and hostnames have been changed to protect the innocent, but all were public, routable addresses.

Line number	Date Time	hostname	Text of violation
1	May 16 14:47:44	router	%SEC-6-IPACCESSLOGP: list 104 denied udp 192.168.122.78(2320) -> 192.168.10.9(53), 1 packet
2	May 16 14:47:45	router	%SEC-6-IPACCESSLOGP: list 104 denied udp 192.168.122.78(2320) -> 192.168.10.20(53), 1 packet

Analysis:
This came from the Internet router that is the first line of defense for the internal and DMZ networks from the Internet community.

The portion of Lines 1 and 2 that show "`%SEC-6-IPACCESSLOGP:`" indicate that "A packet matching the log criteria for the given access list was detected," with the rest of the text describing the item that was logged.

Line 1 shows an inbound udp port 53 (DNS query) request being denied to host 192.168.10.9 and logged via ACL 104.
Line 2 shows an inbound udp port 53 (DNS query) request being denied to host 192.168.10.20 and logged via ACL 104.
ACL 104 has been applied to the "incoming" side of the serial interfaces of the Internet router.

The requesting machine somewhere out on the Internet (192.168.122.78) is using a "non-privileged" port (I.E. > 1023) and is making a port 53 (DNS) query to the internal host, 192.168.10.9, which is the internal primary DNS server. That fails, so access to 192.168.10.20, the internal secondary DNS server, is attempted and also denied.

Most likely cause in my analysis would be an employee taking a laptop home, dialing the local ISP, and trying to communicate with the Internet, using the laptop that is configured for a LAN connection, not a ISP dial-up connection. Proper configuration of the laptop should solve this issue, although intrusion attempts could also easily be masked in this way and should not be overlooked as "normal" poor configuration issues.

Description of the Rule:
The ACL rule that blocked this is the "access-list 104 deny ip any any log," which blocks and logs all IP (udp included!) traffic not previously explicitly allowed, per our site's policy.

Potential Damage:
Potential damage comes in the form of unnecessary information being provided to unauthorized outsiders. Having access to all information in the internal DNS would provide would-be intruders with valuable information that could be used to attempt further access. By blocking this at the Internet router, the firewall never sees this traffic.

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

Assignment 2: Five Violations of site's firewall policy (Continued).

Violation 4: Attempt to telnet from Internet to firewall

Note: IP addresses and hostnames have been changed to protect the innocent, but all were public, routable addresses.

Line number	Date Time	hostname	Text of violation
1	Jun 4 14:48:54	router	%SEC-6-IPACCESSLOGP: list 104 denied tcp 192.168.171.206(1067) -> 192.168.5.10(23), 1 packet

Analysis:
This came from the Internet router that is the first line of defense for the internal and DMZ networks from the Internet community.

The portion of Line 1 that shows "%SEC-6-IPACCESSLOGP:" indicates that "A packet matching the log criteria for the given access list was detected," with the rest of the text describing the item that was logged.

Line 1 shows an inbound tcp port 23 (telnet) request being denied to host 192.168.5.10 and logged via ACL 104. ACL 104 has been applied to the "incoming" side of the serial interfaces of the Internet router.

The requesting machine somewhere out on the Internet (192.168.171.206) is using a "non-privileged" port (I.E. > 1023) and is making a port 23 (telnet) request to the firewall, 192.168.5.10.

Most likely cause in my analysis would be a dial-up Internet user attempting to gain telnet access to the firewall, perhaps in an attempt to gain further access into the internal network. The first line of defense, the Internet router, blocked the access and the firewall never saw the telnet request.

Description of the Rule:
The ACL rule that blocked this is the "access-list 104 deny tcp any host 192.168.5.10 eq telnet log," which blocks and logs any attempt to connect to port 23 (telnet) on the firewall from the Internet. It is our feeling that any telnet access to the firewall from the Internet is much too risky to allow, even if allowed from a specific IP address, as IP addresses may be forged and unauthorized access obtained.

Potential Damage:
Potential damage comes in the form of unauthorized access to a vital link in the security infrastructure of this site. If access to the firewall was obtained (especially telnet!), then access to the rest of the network would be sitting at the front door of a would-be intruder. Data that could be accessed, stolen, modified, etc, includes, but is not limited to, financial data, engineering data, HR data, etc. Access to those types of data could be detrimental to the success of the company.

© SANS Institute 2000 - 2002 Author retains full rights.

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

Assignment 2: Five Violations of site's firewall policy (Continued).

Violation 5: Attempt to connect to external POP3 mail server through firewall

Note: IP addresses and hostnames have been changed to protect the innocent, but all were public, routable addresses.

Line number	Date Time	hostname	Text of violation
1	Jun 6 09:38:48	firewall	unix: securityalert: tcp if=qfe0 from 192.168.49. 52:1218 to 192.168.200.18 on unserved port 110

Analysis:
This entry comes from the Internet firewall and is an attempt from an internal machine to connect to a remote machine on tcp port number 110 (pop3 mail service). Because of the "if=qfe0," one can determine that the request is from an internal machine, as the qfe0 interface is on the "inside" network. Also, it was apparent the request originated from the internal network because the IP address in the "from" portion was an internal IP address (although not obvious here, due to the sanitized IP addresses).

Line 1 shows an internal machine (192.168.49.52) using a "non-privileged" port (I.E. > 1023) attempting to connect to the port 110 (pop3) process on remote machine 192.168.200.18. The "unix: securityalert:" portion of the line implies that there is a potential attempt to breach the security of the firewall, because there was nothing in place to "properly" handle this request, so it is considered to be a security "alert."

Description of the Rule:
The rule that blocked this is not so much a "rule," but is due to the fact that the firewall has NO processes listening on the specified pop3 port, 110. Since there is not a process listening for port 110 connections on the firewall (I.E. the pop3 proxy is not enabled on the firewall), the request is logged and dropped by the firewall.

Potential Damage:
Potential damage may come in the form of information coming into the organization without being "screened" by an email scanning utility. This may allow viruses to be introduced into the organization that may have otherwise been caught by the screening software and cause loss of information, downtime of critical servers, etc.

Additionally, public relations of your site may be hurt, if your email server falls victim to a virus introduced into the internal network in this way. Say an internal user receives a virus-ridden email from the external pop3 service, and subsequently forwards to all his buddies internally. Depending on the severity of the virus, it may cause DoS and/or downtime issues with the mail server if it is of the type that resends itself to everyone in the "Address List." This downtime may cause bad publicity for the organization, but may have been caught if the only email entering the organization is through approved (and screened) channels.

© SANS Institute 2000 - 2002 Author retains full rights.

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

Assignment 3: Defense in Depth Architecture

Sub-Assignment A: Dual Internet Connections

One of the best ways to minimize the effect of DDoS is to make sure your own systems do not participate in a DDoS attack on someone else's network. If everyone did that, there would be no place for hackers to stow their agents in anticipation of launching them, once enough second-hand victims had been found. Because we know not everyone will perform those action we must ensure that the perimeter routers on each connection filters both inbound (from the Internet) and outbound (from internal networks) traffic. The outbound filters should ensure that the source IP address is truly from the internal networks, and has not been spoofed to try to mask where the attack is coming from. This won't stop outbound traffic, but will make tracing back the attack possible to the internal machine, and can serve as an early-warning that your subnet is being used in a DDoS attack onto someone else's network.

In this specific case, the Internet servers could be separated from the "regular" internal networks by using one Internet connection to serve only the Internet services (Web, FTP, DNS, etc. - I.E. inbound initiated). The second connection would be used to serve only the internal users who need access to Internet services (Web browsing, other Internet access - I.E. outbound initiated). This would keep the Internet servers from being exploited and used to deny service to the internal machines and would also eliminate the need for a firewall to protect the (non-existent) internal networks, because on the first connection, there would be no internal networks to protect (I.E. It would be "all DMZ" on the first Internet connection). Since no outbound requests would be initiated, the only traffic that would need to pass through the router would be in response to an incoming request.

On the second Internet connection, where all systems would be located on the internal networks, a perimeter router and a proxy-type firewall would be put into place and there would be no systems between the perimeter router and the firewall (I.E. No DMZ network). The only traffic that should be expected on this perimeter router should be initiated from the internal networks, so the router/firewall should be configured to only allow outbound traffic (with replies).

An intrusion detection system could be employed to at least detect when DDoS is occurring. Although it may not prevent DDoS attacks, at least this configuration would inform the site that DDoS is taking place on the site, or is being launched from the site, and allow administrators to take appropriate steps to stop the DDoS traffic.

The hardware used here would include the following:

- A perimeter router for each connection.
- A proxy-type firewall on connection 1 to protect the internal networks by using proxied connections to the Internet.
- Optional intrusion detection clients/servers to detect when intruders are attempting to break in.
- Optional "probe" clients on each connection to test for vulnerabilities on the other connection.

The router on connection 1 would be configured to block all inbound traffic that was not a response to an internally generated request. It should also be configured to block all outbound traffic that has a source address outside the "known" internal address space.

The proxy firewall on connection 1 would be configured to proxy all communications from internal systems to external systems to mask the internal IP addresses from the Internet community. It would also provide a second line of defense for the internal networks should the perimeter router be compromised.

The router on connection 2 would be configured to block all outbound traffic that was not a response to an externally generated request. There may be a few other items that are valid to allow, such as DNS queries to externally located DNS servers, or email generated via cron scripts on the DMZ servers, etc. These items should all be blocked initially and allowed on a case by case basis.

See the Network Layout diagram for a graphical display of the proposed network configuration. Note the optional systems (intrusion detection and "probers") are not shown, as they don't directly affect resistance to DDoS, but they may provide the ability to detect and thwart DDoS attacks, if properly monitored.

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

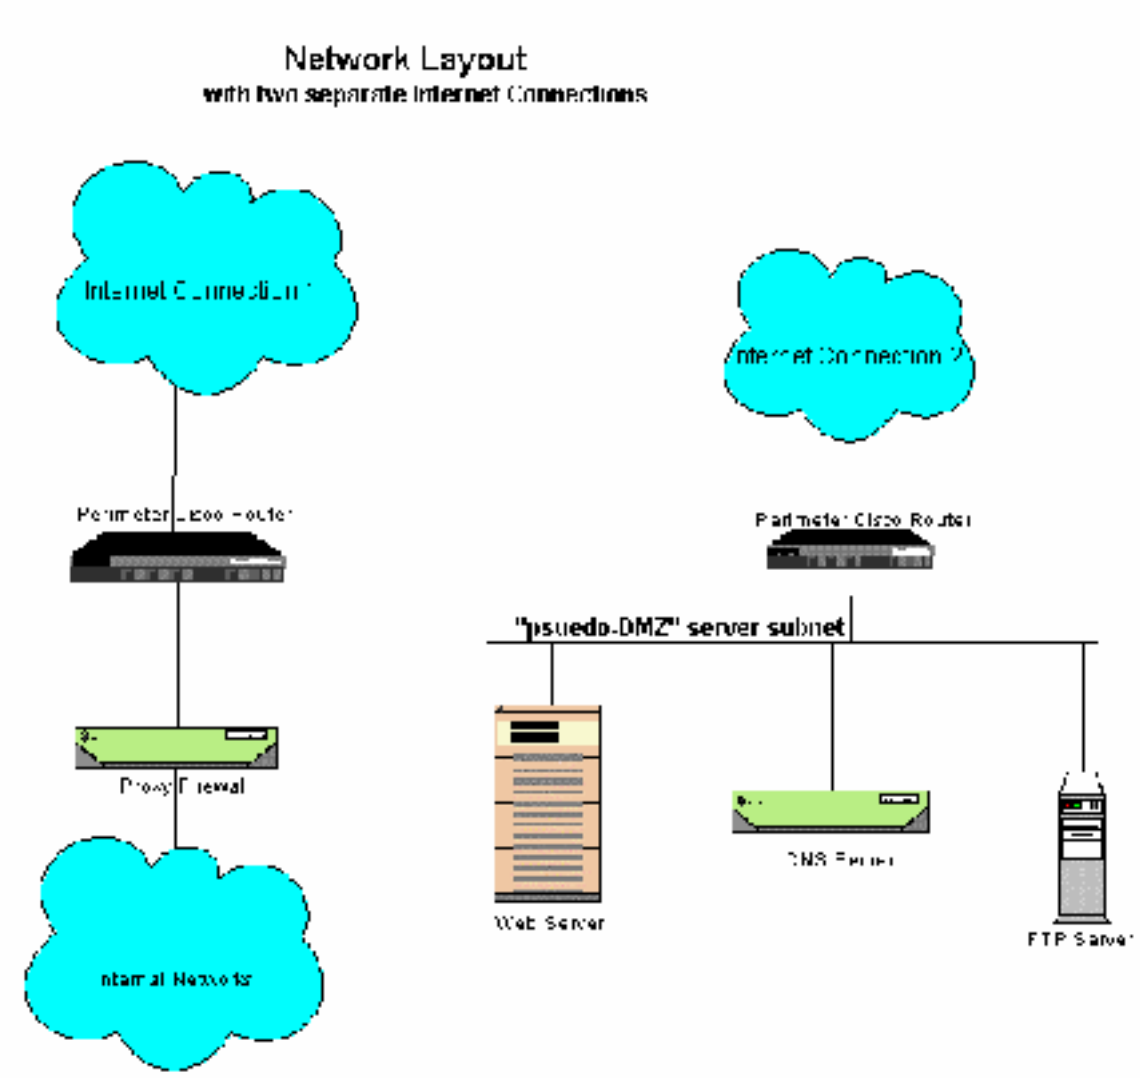

This configuration provides the following:

- Perimeter defense via ACL's in each perimeter router as the main line of defense against the Internet community.
- For connection 1, proxied connections to those services that are desired, via the proxy firewall as the second line of defense.
- For connection 2, the DMZ servers cannot be used to "attack" the internal networks, because no internal network exists on this connection.
- For connection 1, since there are no DMZ servers to plant the DDoS attacking software, the risk is minimized.
- For connection 2, since there are only DMZ servers there is no internal network to "terrorize," so again the risk is minimized. This connection could be used, however to launch attacks to other unprotected networks elsewhere.

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

Assignment 3: Defense in Depth Architecture (Continued)

Sub-Assignment B: Provide protection with equipment already purchased

First, we need to protect the entire site from the Internet using perimeter protection. The Cisco router should be connected to the Internet as well as the DMZ network. The router will be used to filter all unwanted traffic from the Internet and will allow only approved traffic to enter, which will then be handed off to the proxy firewall box for further processing. The proxy firewall box should be configured to only allow those services desired, as defined by the site's security policy, and will log and drop all other attempts to enter the internal networks.

There should be very few denials logged on the firewall box, as the router should be gleaning most of the unwanted traffic before the firewall even sees it. Adequate monitoring of the router and proxy firewall logs should be performed to ensure they are configured as desired. If one sees things at the firewall that the router should have caught, it would be prudent to revisit the router configuration to ensure something was not missed and to correct the problem, if one exists.

Next, we proceed to further protect the research and accounting networks. We'll start with the research network. Plug one of the appliance-type firewall appliance boxes (ATFB) into the main router of the company (and configure the main router appropriately to accept it). Use the ATFB to perform NAT to hide the IP addresses on the research network from the other internal networks. Also, configure the ATFB to restrict incoming (to the research network) traffic to only what should be allowed into the research network, and to allow all outbound traffic (unless it is desired to restrict outbound traffic). Effectively, the research network is the LAN side (I.E. the trusted network) of the ATFB, and the rest of the internal network(s) are considered the WAN side of the ATFB (I.E. the untrusted networks).

Repeat the above paragraph for the accounting network. The two ATFB's may be configured differently depending on the differing requirements of each respective network.

See the Network Layout diagram for a graphical display of the proposed network configuration.

© SANS Institute 2000 - 2002, Author retains full rights.

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

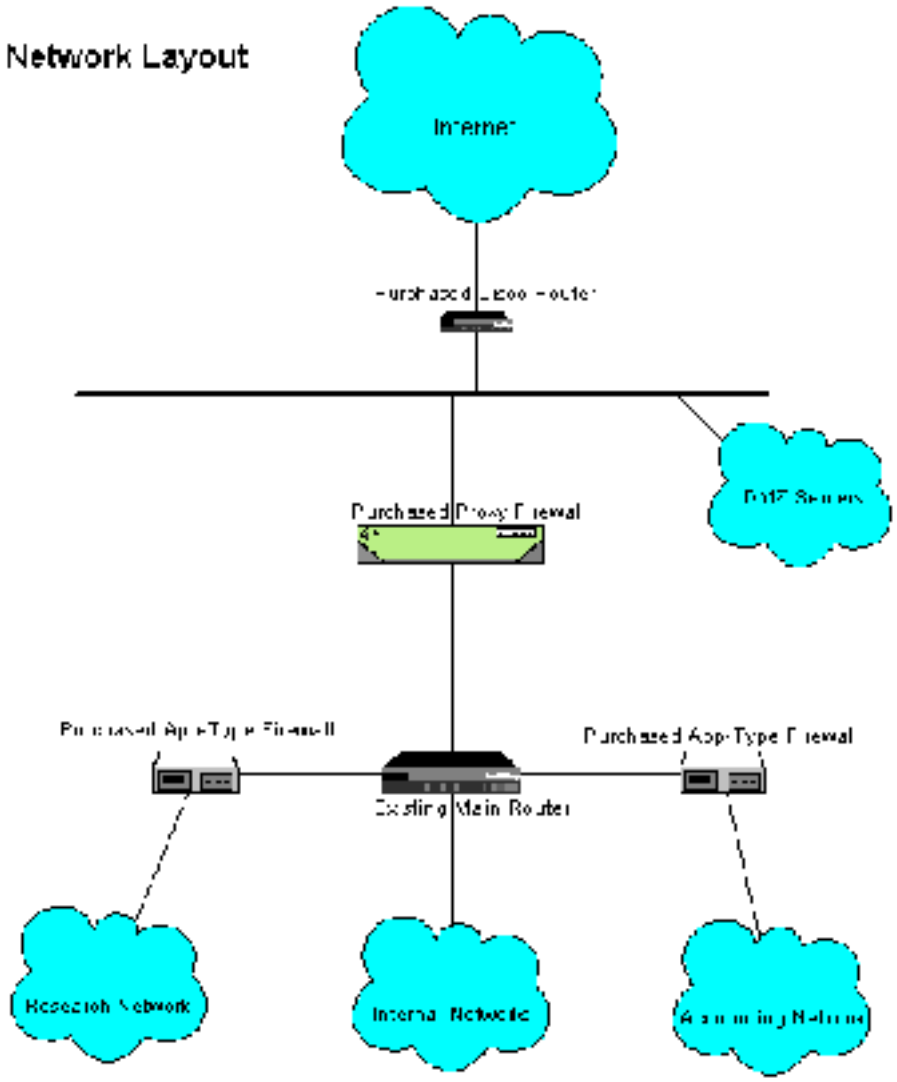

This configuration provides the following:

- Perimeter defense via ACL's in the Cisco router as the main line of defense against the Internet community.
- Proxied connections to those services that are desired, via the proxy firewall as the second line of defense.
- Protection for the research and accounting networks from the rest of the internal networks (which, by the way, are considered "hostile" in the eyes of research and accounting!)

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

Assignment 4: Create a Test demonstrating Subject Area knowledge

Q: A company has purchased a second Internet connection. The company uses its original connection in the usual manner, sending/receiving email, web browsing, ftp access, company-specific web serving, etc. The second connection was purchased to provide web-hosting services to external companies. The company desires to keep the two connections separate, although the administrators, who reside on the internal network and connect to the Internet using the first connection, need access to the other sites to perform day-to-day maintenance of the external web hosting equipment. Using the Internet, and without using dial-up modems or direct WAN connections to the equipment on connection 2, provide a secure network design that allows access from connection 1 to all equipment on connection 2.

A: One solution would be to connect an appliance-type firewall box (ATFB) on the second connection that provides VPN support. It would need to talk to a machine on connection 1 that also provides VPN support, like a firewall with VPN, a Windows2000 VPN server, etc., but one must ensure the two VPN endpoints are compatible (I.E. can talk the same encryption methods). Optionally, the connections could each have an ATFB installed that can interact with each other to ensure compatibility, but can sometime require additional cost that may be avoided by utilizing an existing firewall or other VPN capable piece of equipment.

The ATFB on connection 2 can be used to guard the web-hosting servers, as well as be used to provide access from the first site to this site, via a VPN tunnel. The tunnel could be created using an encryption method that provides full encryption of each packet (I.E. ESP or AH protocol) to help mask the true identity of the packet, since this traffic will be traversing the Internet. Ensure that the security associations on both endpoints are in agreement, and are strict enough to match your desired level of security.

This setup also provides an added benefit of having the ability to build machines that can be placed on each connection that can "probe" the other site to check for and report on vulnerabilities found on the other site. These machines may be used to test the intrusion detection system that should be placed on each connection.

See the Network Layout diagram for a graphical display of the proposed network configuration.

© SANS Institute 2000 - 2002, Author retains full rights.

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

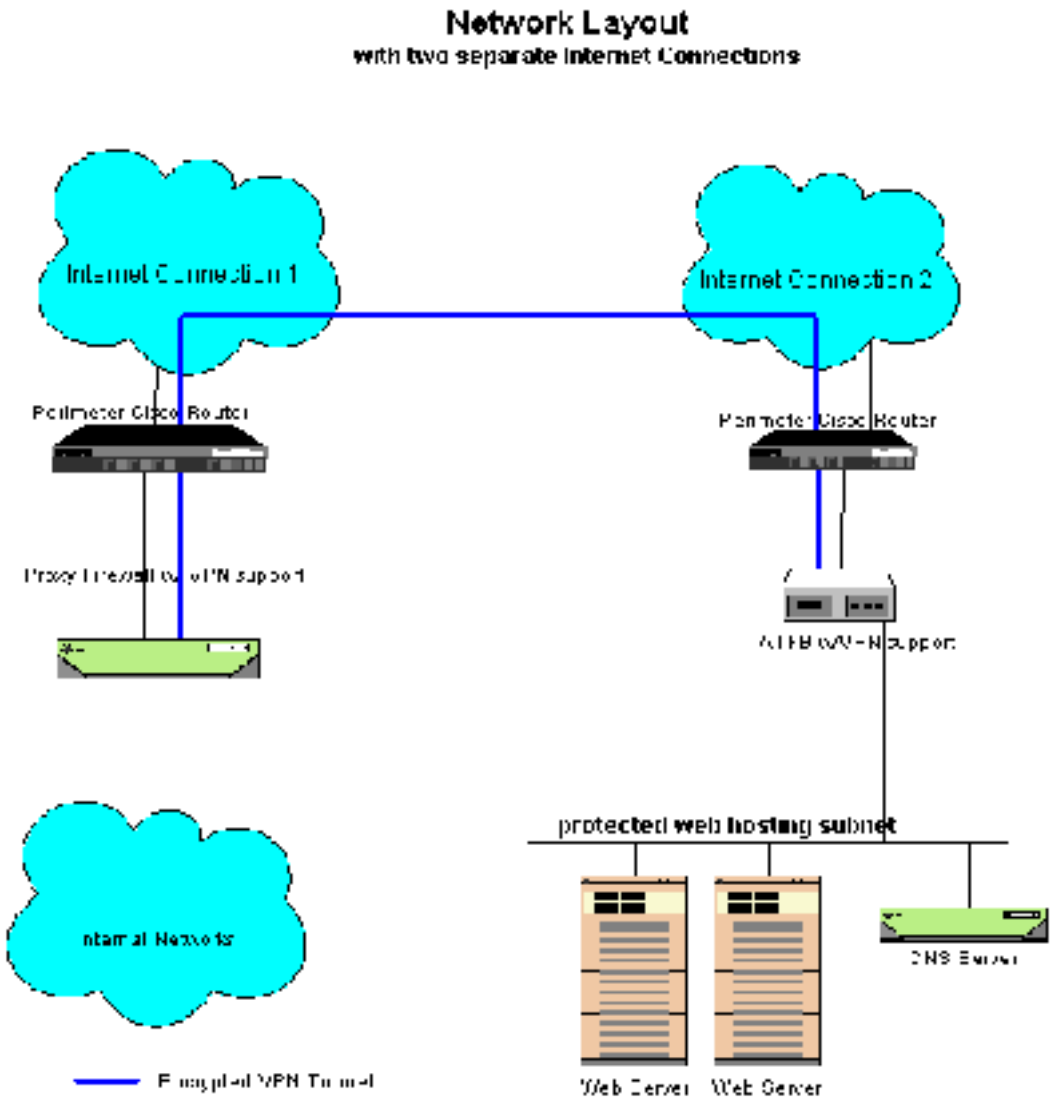

This configuration provides the following:

- Perimeter defense via ACL's in the Cisco router as the main line of defense against the Internet community.
- Proxied connections for those on connection 1 that need access to Internet services, via the proxy firewall as the second line of defense.
- Encrypted VPN tunnel, which is not as secure as a direct connection between the two sites, but provides a fairly secure, cost-effective solution.
- Ability for the administrators on connection 1 to access servers on connection 2 via an encrypted VPN tunnel. This tunnel is created between the proxy firewall (w/ VPN support) on connection 1, to the ATFB (w/ VPN support) on connection 2.

URL's/References/Resources:
The site, http://www.sonicwall.com/Firewall-PRO/vpn_overview.html, gives an overview of SonicWall's ATFB w/ VPN support. These boxes can create a VPN tunnel with each other, or with compatible products, such as MS Windows 2000 server with VPN support. *"SonicWALL VPN provides an easy, affordable, and secure means for businesses to connect all offices and partners together."* Of course, this is directly from the vendor's page, but refer to http://www.sans.org/y2k/firewall.htm to get an independent observer's opinion of the product.

Virtual Private Networks: Planning and Implementation - A Bird's of a Feather presentation by George Freeman presented at the San Jose SANS conference.

I also relied heavily on my own actual experience with SonicWall and VPN over the Internet to formulate this question and answer.

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.