



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

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

SANS GIAC Level 2
GCFW – Firewalls, Perimeter Protection, and VPNs
Practical Assignment

SANS New Orleans 2001

David M. Stokes

15 January 2005

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

Assignment #1 – Security Architecture

Task Define a security architecture for GIAC Enterprises, a growing Internet startup that expects to earn \$200 million per year in online sales of fortune cookie sayings, and which has just completed a merger/acquisition. Your architecture must specify filtering routers, firewalls, VPNs to partners, secure remote access, and internal firewalls. Be explicit about the brand and version of each perimeter defense component. Produce a diagram or set of diagrams with explanatory text that define how to use perimeter technologies to implement your security architecture.

You must consider and define access for:

- Customers (the companies that purchase bulk online fortunes);
 - Suppliers (the authors of fortune cookie sayings that connect to supply fortunes);
 - Partners (the international partners that translate and resell fortunes).
-

Architecture Description

GIAC Enterprises, a \$200+-million-per-year e-business, needs its network to provide high availability and security, while allowing key partners, suppliers, and customers to be able to conduct business with them. Customers and partners need secure communication via the Web to be able to purchase and resell the fortunes, while suppliers need multiple methods of access to the storage databases. The telecommuters and “road warriors” (sales force) of GIAC Enterprises also need secure access to the key elements of the network (like the CRM application and email). And internal users need appropriate access to the services and elements that will allow them to do their jobs successfully and swiftly. This fundamental tenets of the network are reflected in several key philosophies that can be found in the architecture (see diagram on page 2).

Key Philosophies

The key philosophies that contributed to the architecture choices I made are as follows:

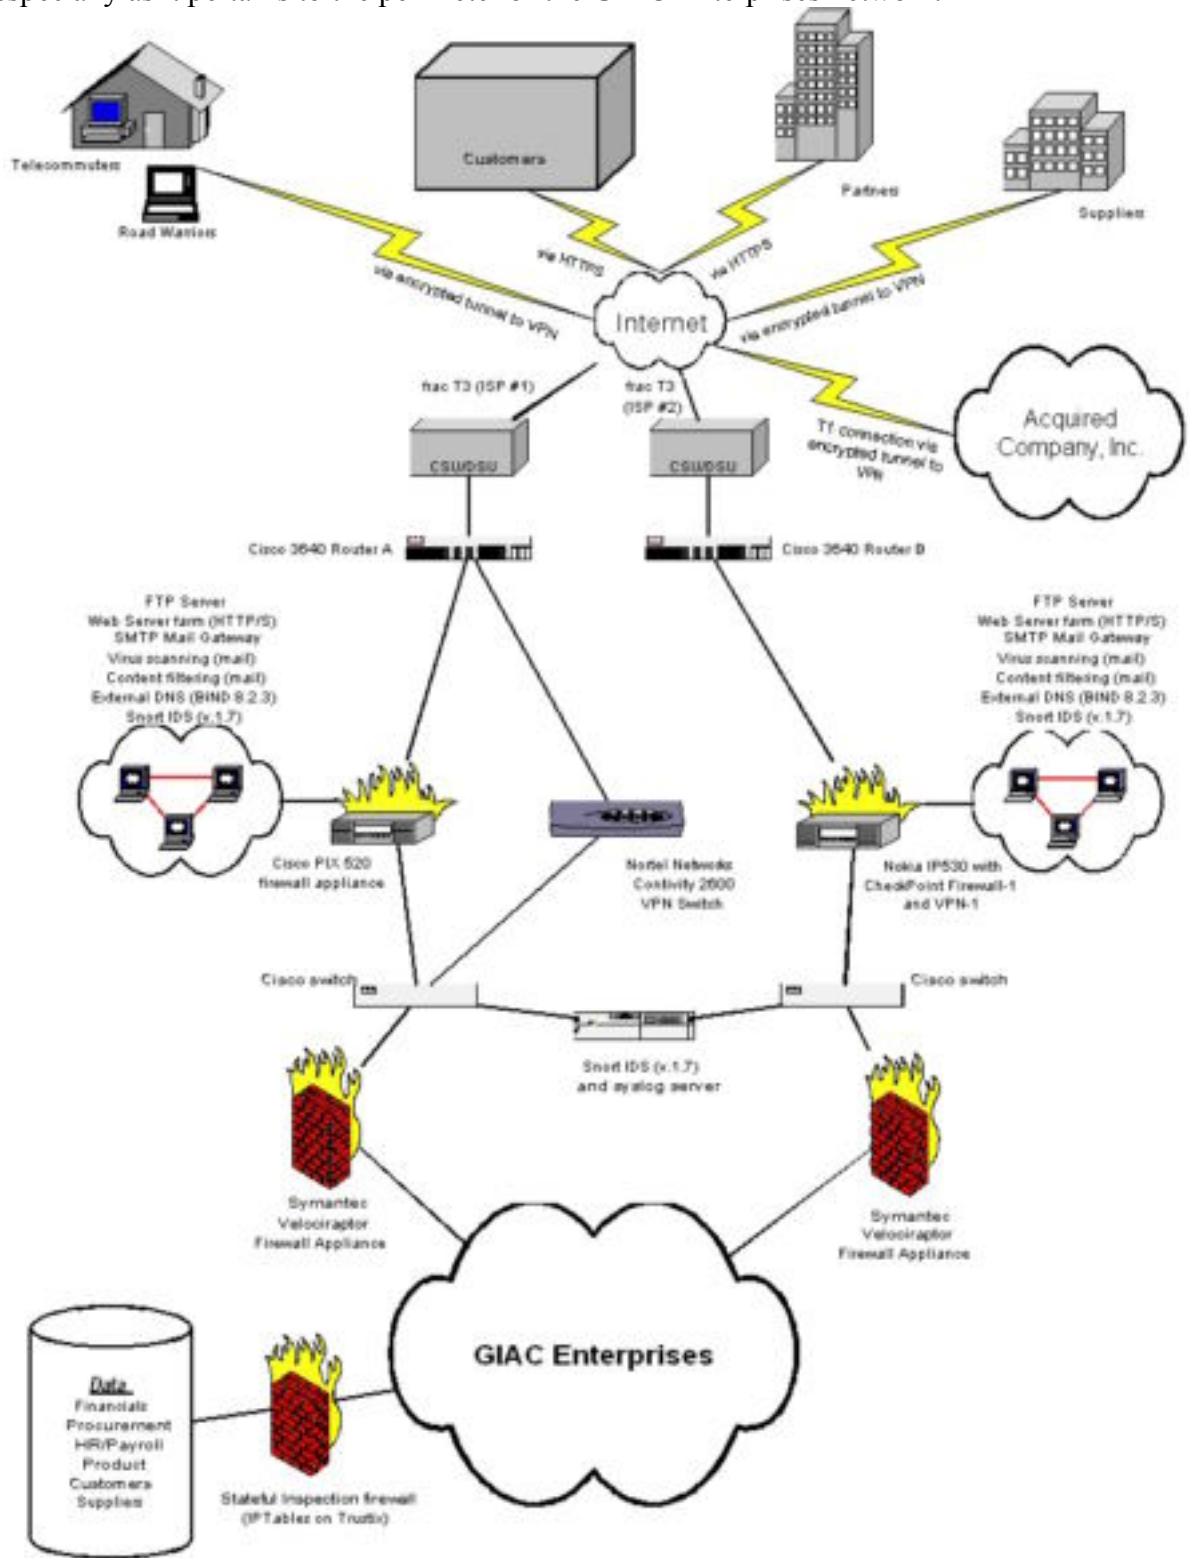
- “*Defense in Depth*”, in which multiple layers of security protect various assets based on their function and level of security classification (i.e. financial and customer data need to be protected more significantly than Help Desk data) must be employed. One should note here that this philosophy can be expensive, specifically if the company has issues with using non-open source solutions exclusively.
 - “*Access to the Right Data at the Right Time from the Right Location*” -- Since GIAC Enterprises is an e-business, the most critical success factor for any security architecture is the capacity, capability, and confidence to serve all customers no matter where (or when) in the world they are located.
 - “*Duplicate Functionality with a Duality of Solutions*” -- The added benefits of this design are (1) an increased level of security since the weaknesses of one path are not necessarily present in the other, and (2) an increased availability quotient for customers. As with the Defense in Depth philosophy, cost is high for this implementation: buying two of every key infrastructure element and related support costs (especially when the two elements are from different vendors) can send expenditures skyrocketing.
-

Continued on next page

Assignment #1 – Security Architecture, Continued

Diagram

The diagram that follows is the high-level view of the security architecture, especially as it pertains to the perimeter of the GIAC Enterprises network.



Continued on next page

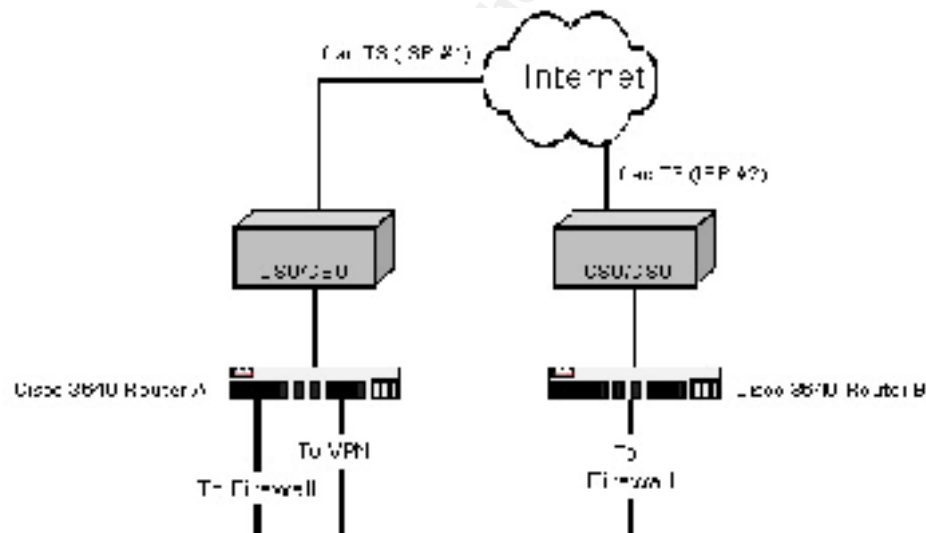
Assignment #1 – Security Architecture, Continued

Details

The specific details of the architecture in the diagram on the previous page are discussed in the blocks that follow. The details that will be highlighted are:

- Border routers
- Firewalls, both stateful inspection and application proxy
- Screened Secure Network (SSN, or also referred to sometimes as “DMZ”)
- Separate internal and external DNS servers
- Mail relay with anti-virus and content filtration capabilities
- Virtual Private Network (VPN) “devices”
- Intrusion Detection Systems
- Logging server

Border Routers The devices at the outermost portion of the network’s perimeter are the border routers.



From a security perspective, they are the first line of defense against several forms of attack. These routers should be used primarily to block both “spoofed” and private addresses (as defined in RFC 1918), as well as source-routed packets. They should also be used to control ICMP traffic, which might lead to a denial of service (DoS) attack if not otherwise managed. Unused services should be disabled, and needed services should be managed with access lists. Ingress filtering should be employed to stop malicious insiders from initiating DoS attacks with forged source addresses, and smurf and fraggle (two types of distributed DoS, or DDoS) attacks should be stopped by filtering packets sent to the broadcast addresses of the network.

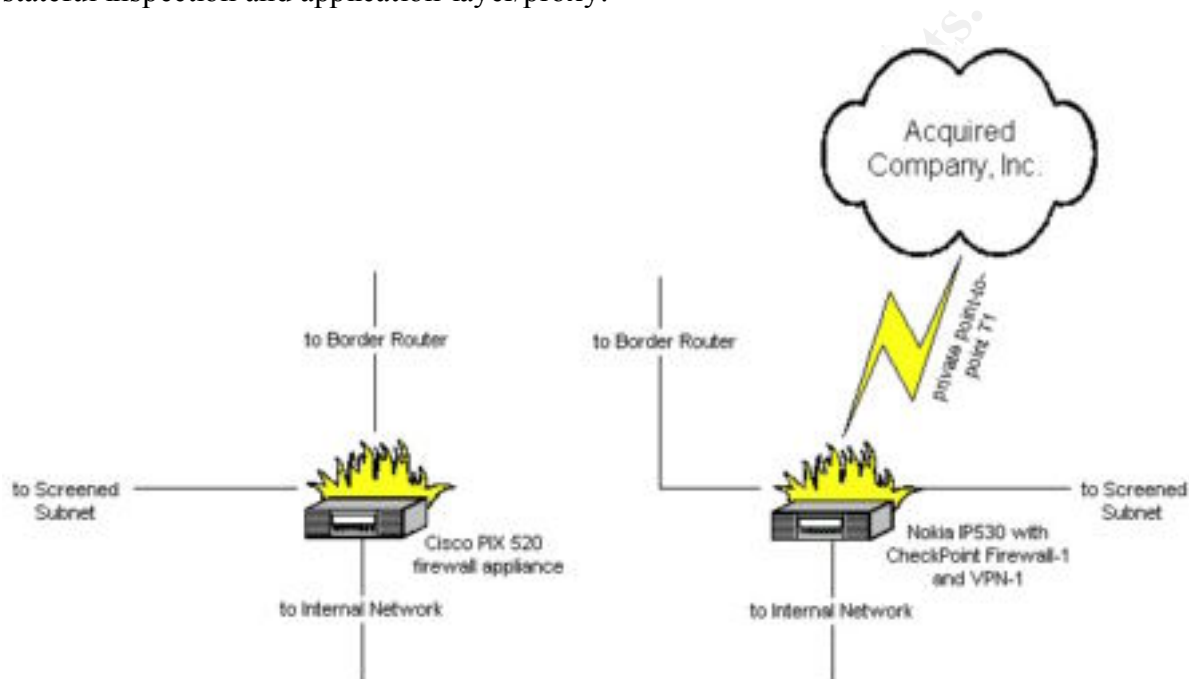
I have chosen to implement Cisco 3640 routers as the border routers based on size and capacity. They should run Cisco IOS 12.x, with appropriate patches applied. Technical configuration of the router interfaces are discussed in **Assignment #2, Border Routers**.

Continued on next page

Assignment #1 – Security Architecture , Continued

Primary Firewalls

I have indicated the use of two types (and several brands) of firewall in the architecture: stateful inspection and application-layer/proxy.



Stateful inspection (a.k.a. dynamic packet filtering) firewalls are a hybrid, mixing the speed of static packet filters (e.g., routers) with the agility of application-layer/proxy firewalls. They do more than simply filter packets; they can also track the state of all active sessions. Ports remain closed when not in use, but are opened when requested (if security policy permits) or a response is being sent.

Notice that the firewalls that I have selected are both “appliances”, meaning that when you purchase the firewall, it comes with the OS and firewall software already installed, compiled, configured, and awaiting rules (from your security policy) to be implemented. Appliances are the appropriate choice in this instance because they are easier to maintain and come with a hardened OS; the one disadvantage is that then the security afforded by the firewall is dependent on both the vendor supplying patches and fixes in a “reasonably short” period of time and the administrator applying the aforementioned patches and fixes also in a timely fashion.

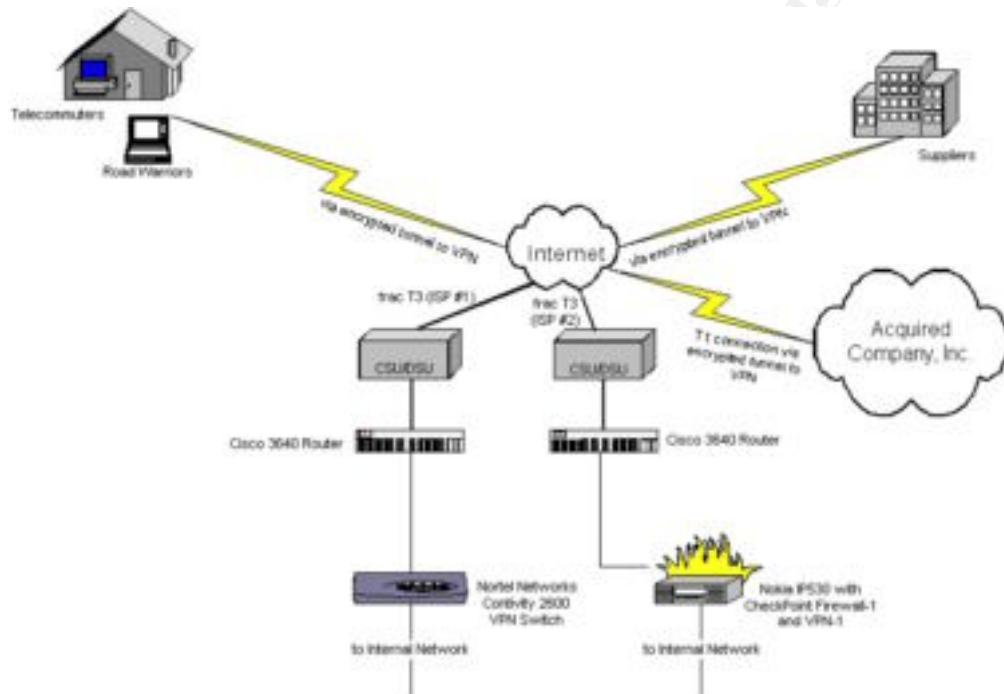
Each of the firewalls supports a screened subnet, which will allow filtered access to our public web and FTP servers, extranet applications, external DNS, mail gateway, and other key services. The absence of load balancing is apparent from the diagram, as is the notion that any failover would be rather manual. The assumption under which I am working is that management already agreed to accept this burden of risk (probably to reduce the bottom line of the project) when it was presented to them. Additional details of the security policy with regard to the primary firewalls can be found in **Assignment #2, Primary Firewalls**.

Continued on next page

Assignment #1 – Security Architecture , Continued

VPNs

Two distinct VPN servers are present in the architecture: a Nortel dedicated Contivity appliance, which sits in parallel with the Cisco PIX firewall, and a VPN-1 plug-in module for the CheckPoint Firewall-1.



The Nortel Contivity switch will service the telecommuters, the “road warrior” sales force, and the suppliers – all of whom will need access to web-based applications in order to talk to the business back-end databases (customers, product, etc.). The VPN-1 module will service the site-to-site connection for Acquired Company, Inc., whose employees will be able to use the same services as the employees of GIAC Enterprises.

The split in function of the two VPN servers should ease some of the burden of the servers (especially the encryption/decryption processes). I would recommend a Luna VPN Accelerator card for the Nokia, to maximize performance for a moderately heavy traffic load.


The encryption on the Nortel unit must be able to support the needs of the international suppliers, as well as those from the U.S. Management has also dictated that no IPSec clients will be configured to run with “split horizon”, meaning that they will not be able to connect directly to the Internet with the IPSec client while they are connected through the VPN and *vice versa*. For technical configuration details about the security policy, see **Assignment #2, VPNs**.

Continued on next page

Assignment #1 – Security Architecture , Continued

IDS and Central Logging Server

The need for a centralized network-level Intrusion Detection Scanner (IDS) and logging server to be part of the overall network security architecture is clearly documented by Lance Spitzner [1].



As the main network-level IDS, the box needs some dedicated processing power and, even more importantly, good logging capabilities. Hence, the marriage of the IDS with the central logging server.

As the central logging server, abundance of disk space and good I/O throughput are essential. Depending on the size of the box and the amount of logging being sent to it, my best guess would be to use mirrored disk sets (RAID 0) and maximize the number of RAID hardware controllers that the box can command. The goal to have 1 mirrored pair for each service or device between this box and the Internet may be lofty, but achievable, perhaps, with an additional investment for external RAID controllers.

Notice that no duplication of this device exists in the architecture, except some local logging and host-based IDS, in order to cross-check for ownership by crackers. However, there really is no need, as long as the box has sufficient ability to recover from any major disasters (which may mean some hot-swappable duplication within the box itself – of power supplies, fans, etc.).

The advantages of a central logging server include the ability for administrators to (1) see trends within the network environment more easily; (2) compare the protected logs housed on the hardened and armored system to a local copy for differences (which may be an indication of log tampering); and (3) be notified and/or automated responses to be carried out more quickly, based on the parsed content of the logs.

Continued on next page

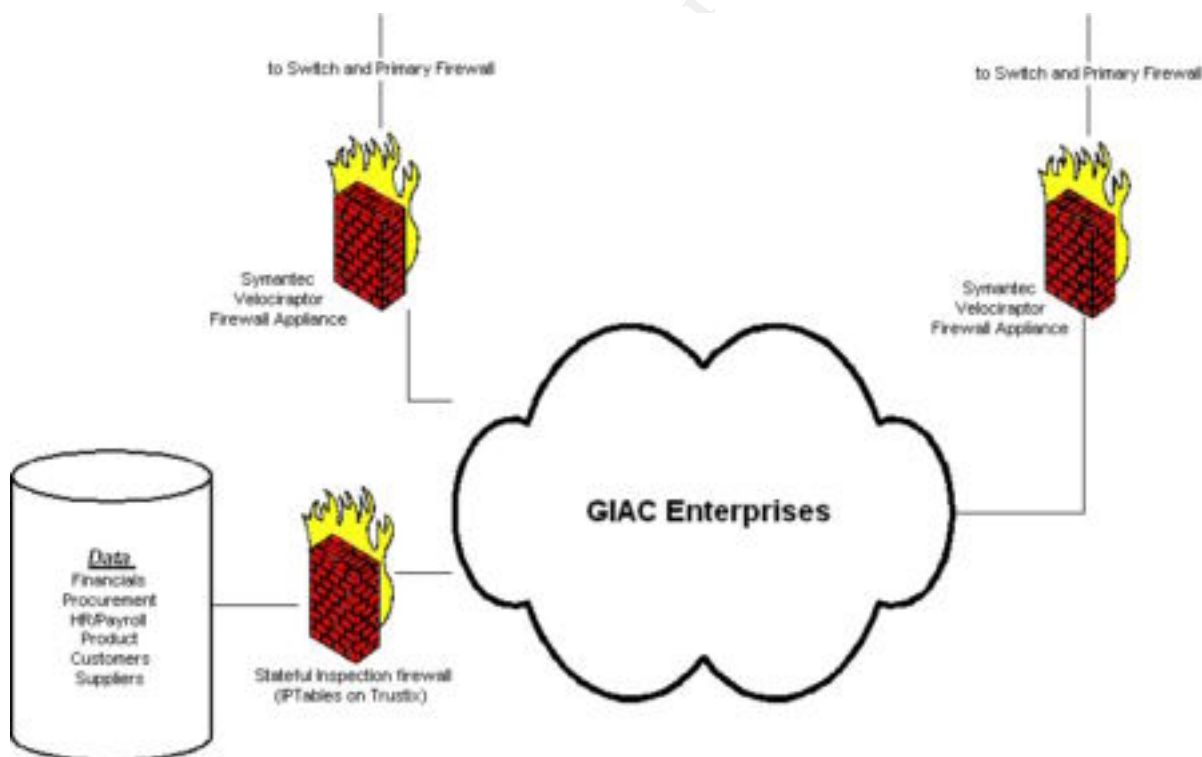
Assignment #1 – Security Architecture , Continued

IDS and Central Logging Server, continued

The choice of Snort, an open source, easily updated, configurable IDS is a good one. User community support exists, as do plug-ins for parsing log files (Snortsnarf, just to name one). It runs on UNIX or Linux, and it is free.

Internal Firewalls

These firewalls allow for segmentation of services that may reside behind them, as in the case of the protected corporate data. That data (like payroll, product, and customer credit card information) is considered “sensitive” and “confidential” – it warrants special protection from the rest of the internal network and the company’s employees (except those with authorization).



They also add an additional layer of filtration to VPN traffic.

The lack of duality in my choice of solution here – Symantec’s Velociraptor appliance – stems from a desire to provide continuity for the administrators who will implement the rules of the corporate security policy on these machines. Although the rule bases may be different, based on the traffic that will flow through each, familiarity with one provides familiarity with all.

Continued on next page

Assignment #1 – Security Architecture , Continued

Screened Security Network (SSN)

On the SSN, an array of services will be provided: FTP, HTTP, HTTPS, SMTP, external DNS, and subnet-based IDS.

The external DNS is the “public” half of the split DNS structure; the internal DNS is found in the GIAC Enterprises internal cloud.

The Web server farm will have to interact securely with the Data subnet through 3 firewall layers, and then respond via HTTPS back to the customer. For the HTTPS service, we will use a Certificate Authority (CA) like VeriSign, RSA, or Baltimore.

Good practice might dictate here the employment of a reverse proxy scenario. I might recommend instead a tool like Tripwire, that will check for and alert web personnel to changes on any given web page.

Virus scanning and content filtering of email is important to complete *before* the email message and its attachment(s) get forwarded inside the firewall. The list of vendors and wares in this space is growing daily; SANS provides the names of several on its Network Security Roadmap 2001 [2].

Additional Notes

Some stresses and strains on the architecture of which one should be aware before implementation:

- IT staff need to be fluent in security measures and configurations, as well as being specifically knowledgeable with regard to (too?) many different types of firewall, VPN, and IDS solutions.
 - Rule sets for devices that provide the same service will look similar, but maintenance of the appliances and the software will become quite a chore if too few people are assigned to the task.
 - Diversity of device has a price.
 - On the other hand, consolidation of logging to a single, protected server is wise (to keep it out of attackers’ hands), but is also both a single point of failure and a potential performance bottleneck.
-

Assignment #2 – Security Policy

Task Based on the security architecture that you defined in Assignment 1, provide a security policy for AT LEAST the following three components:

- Border Router
- Primary Firewall
- VPN

You may also wish to include one or more internal firewalls used to implement defense in depth or to separate business functions.

By 'security policy' we mean the specific ACLs, firewall ruleset, IPSec policy, etc. (as appropriate) for the specific component used in your architecture. For each component, be sure to consider internal business operations, customers, suppliers and partners. Keep in mind you are an E-Business with customers, suppliers, and partners - you MAY NOT simply block everything!

(Special note VPNs: since IPSec VPNs are still a bit flaky when it comes to implementation, that component will be graded more loosely than the border router and primary firewall. However, be sure to define whether split-horizon is implemented, key exchange parameters, the choice of AH or ESP and why. PPP-based VPNs are also fully acceptable as long as they are well defined.)

For each security policy, write a tutorial on how to implement each ACL, rule, or policy measure on your specific component. Please use screen shots, network traffic traces, firewall log information, and/or URLs to find further information as appropriate. Be certain to include the following:

1. The service or protocol addressed by the ACL or rule, and the reason these services might be considered a vulnerability.
2. Any relevant information about the behavior of the service or protocol on the network.
3. The syntax of the ACL, filter, rule, etc.
4. A description of each of the parts of the filter.
5. An explanation of how to apply the filter.
6. If the filter is order-dependent, list any rules that should precede and/or follow this filter, and why this order is important. (Note: instead of explaining order dependencies for each individual rule, you may wish to create a separate section of your practical that describes the order in which ALL of the rules should be applied, and why.)
7. Explain how to test the ACL/filter/rule.

Be certain to point out any tips, tricks, or "gotchas".

Continued on next page

Assignment #2 – Security Policy, Continued

Border Routers Using the first 9 steps outlined in Brett Eldridge’s “Building Bastion Routers Using Cisco IOS” [3], I include the following configurations with discussion after each step. Step 10, “Protect hosts behind the router”, is really configuration that belongs on the firewall, as it is the main “protector” of the hosts located internally and on the SSN. Step 11 is a testing and verification step which will be discussed at the end of the block.

```
!  
! Step 1 - Password Protection  
!  
service password-encryption  
enable secret my.password
```

As far as the Corporate Security Policy is concerned, these first two commands are just best practices. The first encrypts the password file (albeit not strongly), and the second encrypts the privileged EXEC mode password with an MD5 hash.

```
!  
! Step 2 - Limit remote access (with “ultra-paranoid” config)  
!  
access-list 99 deny any  
line vty 0 4  
  access-class 99 in  
  exec-timeout 0 1  
  login local  
  transport input none
```

Here, the Corporate Security Policy follows the best practices again, granting remote access to no one. Problem with this policy is that support personnel who have to fix a problem after standard business hours will be unable to make modifications (or even simply restart the router) without physically being in the office. Boiled down, travel time for support personnel means increased downtime; and downtime means lost revenues (and profit!). But until revenues are proven to be lost, management would rather take the downtime.

```
!  
! Step 3 - Limit local access  
!  
line con 0  
  login local  
  exec-timeout 2 0  
line aux 0  
  login local
```

Limiting the local access in this way is more security from unwanted insiders than attackers outside the company. Even if the router is secured physically, people with “keys” to the secured facility are not supposed to be able to just walk up to the router console and gain access without knowing the password.

Continued on next page

Assignment #2 – Security Policy, Continued

Border Routers, continued

The border router configurations continue from the previous page.

```
!  
! Step 4 - Display login banner  
!  
banner login #  
        WARNING: Only authorized access is permitted!  
#
```

Legally, the reason for having a login banner is so that a judge will not throw out a criminal (or civil) suit against an alleged attacker simply because you did not tell him explicitly that there was “No Trespassing” on your private property. In reality, the only people who should ever see this message are the ones who are trying to login to the system locally.

```
!  
! Step 5 - SNMP  
!  
no snmp
```

Corporate Security Policy dictates that no SNMP should be allowed from the border routers due to the nature of SNMP. Of course, if you have a Data Center with operators 24x7, then that policy makes some sense. If you are relying on something to tell you when it is having a problem, SNMP might be one of the better choices – but it also might require significantly more configuration.

```
!  
! Step 6 - Logging data  
!  
no logging console  
! Uncomment line logging to firewall you are connected to  
logging xxx.xxx.246.2  ! external interface of Cisco PIX firewall  
!logging xxx.xxx.246.3  ! external interface of Nokia firewall
```

These commands tell the system not to log to the console and then to send log data out to the external interface of the Cisco PIX firewall (in the case of Router A; the Nokia firewall for Router B). The firewalls have rules telling them to direct UDP syslog packets to the central logging server.

I should mention something of NTP, since without the router’s internal clock being attuned to the same time as the rest of the network, forensics can be significantly more difficult. However, at the present, the company does not have any internal NTP servers; and since configuring NTP to search externally for the correct time could provide an opportunity for an attacker to spoof an NTP server and potentially carry out, at minimum, a DoS attack, management has decided that the risk is too great to bear.

Continued on next page

Assignment #2 – Security Policy, Continued**Border Routers, continued**

The border router configurations continue from the previous page.

```

!
! Step 7 - Other protection mechanisms
!
!           Global commands
!
no ip source-route
no service tcp-small-servers      ! Default with IOS 12.x
no service udp-small-servers      ! Default with IOS 12.x
no service finger
no ip bootp server
no ip http server                  ! Default, but be paranoid
no ip domain-lookup
no cdp run                        ! cdp = Cisco Discovery Protocol
no ip unreachable

```

Most of these protection mechanisms are simply best practices, many have been identified by more than one source (Lance Spitzner [4], Cisco [5], Brett Eldridge [6], Frank Keeney [7], and numerous GCFW candidates in their practicals). See any one or all of them for a more detailed description of each line. I categorize the lot under “Turn off all unnecessary services, as well as the necessary ones which are insecure”.

```

!
!           Interface-specific commands
!
interface Ethernet0               ! External interface of router
no ip directed-broadcast          ! Prevent smurf attacks - Step 9
no ip proxy-arp
no ip redirects
no cdp enable                     ! duplicated effort with no cdp run
no ntp enable

```

Shown above only for interface Ethernet0, these best practice commands should be applied to each interface of the router, including Ethernet1 (our external interface) and Ethernet2 (our internal interface on Router A to the Nortel Contivity VPN switch).

```

!
! Step 8 - Anti-spoofing
!
! Beginning of access-list 101
!
! Deny RFC1918 addresses (and by using 172.16.x.x for internal
! and 192.168.x.x for DMZ addresses, we also filter any spoofed
! addresses of machines located behind firewall)
!
access-list 101 deny      ip 192.168.0.0 0.0.255.255 any log
access-list 101 deny      ip 172.16.0.0 0.15.255.255 any log
access-list 101 deny      ip 10.0.0.0 0.255.255.255 any log

```

Continued on next page

Assignment #2 – Security Policy, Continued**Border Routers, continued**

The `border router configurations continue from the previous page.

```

!
! Deny packets with localhost, broadcast, multicast (class D),
! and "reserved for future use" (class E) addresses
!
access-list 101 deny      ip 127.0.0.0 0.255.255.255 any log
access-list 101 deny      ip 255.0.0.0 0.255.255.255 any log
access-list 101 deny      ip 224.0.0.0 31.255.255.255 any log
access-list 101 deny      ip 240.0.0.0 7.255.255.255 any log
!
! Deny test-net and end node autoconfig, respectively
!
access-list 101 deny      ip 192.0.2.0 0.0.0.255 any log
access-list 101 deny      ip 169.254.0.0 0.0.255.255 any log
!
! Filter out all ICMP
!
access-list 101 deny      icmp any any
!
! Deny packets without ip address (0.0.0.0) or with first octet
! zeros.  BEWARE: May also filter out packets from many BOOTP/DHCP
! clients.
!
access-list 101 deny      ip 0.0.0.0 0.255.255.255 any log
!
! More spoofing prevention:  Insert IP address of external
! router interface ip address
!
access-list 101 deny      ip host xxx.xxx.243.10 any log
!
! Allow only IP packets that have made it this far into our network
!
access-list 101 permit    ip any xxx.xxx.243.0 0.0.0.255 log
!
! Finish with denial of anything that fell through the cracks
!
access-list 101 deny any any log
!
! Apply access-list 101 to external interface
!
interface Ethernet0
 ip access-group 101 in

```

After denying many, many different types of IP (and ICMP) packets entrance into our network, the final 2 lines of the access list (prior to applying to the external interface) allow any additional IP traffic destined for our network in and deny everything else.

The amount of logging here might seem excessive, but with a central logging server with lots and lots of disk space to fill, I didn't feel the need to be stingy. Besides, after the first real attack, forensics will be much easier with more extensive logging.

Continued on next page

Assignment #2 – Security Policy, Continued**Border Routers, continued**

The `border router configurations continue from the previous page.

```
!
! Step 9 - Mitigating Denial of Service Attacks
!
!           Egress filtering (stops malicious insiders)
ip access-list extended egress
  permit ip xxx.xxx.243.0 0.0.0.255 any
  deny ip any any log
!
interface Ethernet1
  ip access-group egress in
```

The egress filter is one method of protection from participating in a Distributed Denial of Service (DDoS) attack. Packets leaving the network must have a source address that belongs to our network. All others are dropped and logged.

Finally, it should be apparent from the configuration that border routers play a significant part in the total security of the network. The access lists that you build to protect your network assets are linear: always make sure that the last line of the access-list is `deny ip any any log` and the lines directly before that last line permit desired packets into/out of the network. The rule of thumb for building access lists is something like, “Deny specifically; permit only what is necessary; and then deny all else.”

Primary Firewalls

The rulebase for the primary firewalls will be very similar, as they are mostly two distinct filtered access points into the same network with identically configured SSNs. I choose to illustrate the rulebase with a table structured similarly to the CheckPoint Firewall-1 GUI. After presenting the entire rulebase, I will discuss the lines in order.

No	Source	Destination	Service	Action	Track
1	Fw-admin	Firewall	FireWall1	Accept	Long
2	Any	Firewall	NBT ident	Reject	
3	Any	Firewall	Any	Drop	Long
4	Mailserver	NOT Internal	SmtP	Accept	Long
5	Dns-server	NOT Internal	Domain-udp	Accept	Long
6	Internal	Webfarm	http https	Accept	Long
7	Internal	Mailserver	SmtP	Accept	Long
8	Internal	FTPserver	ftp	Accept	Long
9	Internal	Screened-Service	Any	Drop	Long
10	Internal	Any	Any	Accept	Long

Continued on next page

Assignment #2 – Security Policy, Continued**Primary
Firewalls,
continued**

The table illustrating the rulebase for the primary firewalls continues from the previous page.

No	Source	Destination	Service	Action	Track
11	Any	Webfarm	http https	Accept	Long
12	Any	Mailserver	Smtpt	Accept	Long
13	Any	FTPserver	ftp	Accept	Long
14	NOT Internal	Dns-server	Domain-udp	Accept	
15	Webserver	Internal-Oracle	SQL*Net	Accept	Long
16	Screened-Service	Any	Any	Drop	Alert
17	Any	Any	Any	Drop	Long

This rulebase is very simple, and about 90% complete; the remaining 10% are missing due to absent requirements. For example, without understanding the communications of the web applications, how am I supposed to know exactly how they need to interact through the firewall with the Internet and the internal network? (See Rule #15.)

Again, the same principle found in constructing the Border Router ACLs is found in the construction of the rulebase for the primary firewalls: the rules are linear, meaning that the packet is inspected and matched against each line of the rulebase (starting with No. 1, and working down to No. 17). Therefore, make sure that your rules do not contradict one another; the first rule that the packet matches is executed. And the last rule of all must be the “catch-all”; if the packet gets there, then it is dropped.

Again, as with the Cisco IOS configuration, lots of logging in this setup. NBT is chatty, so we decided not to log that. Everything else is logged long, at least for the moment. I suspect that as time progresses, and the rulebase gets more complex, some rules will get logged short and others not at all. That being said, the explanation of the rules themselves follows:

The first rule says that if you are using the FireWall1 admin client to connect to the firewall, allow it to continue (and log it long).

The second rule rejects any NBT or ident services attempting to connect to the firewall.

The third rule drops all other packets attempting to use any service on the firewall (and log long).

The fourth and fifth rules allow sendmail and the external DNS server, respectively, to send packets into the Internet (as long as they do not try to connect to the internal network!).

Continued on next page

Assignment #2 – Security Policy, Continued

Primary Firewalls, continued

The explanation of the rulebase continues from the previous page.

The sixth, seventh, and eighth rules allow for internal clients to use HTTP (secure and non-secure), FTP, and SMTP services on the SSN (and log long).

Rule 9 drops all other attempts from the internal network to get to the SSN; and rule 10 allows all non-SSN traffic from internal users out the front door.

Rules 11-14 permit all web traffic to the web farm; all FTP traffic to the FTP server; all SMTP traffic to the sendmail server; and all DNS traffic to the DNS server (all of which are located on the SSN).

Rule 15 attempts to allow SQL*Net traffic generated by the web applications in the web farm, destined for the Data Subnet, to pass into the internal network.

Rule 16 drops all other packets from the Screened-Service network (SSN) to any destination using any service. This rule prevents people who may attempt to perform mischievous acts after taking over one or more of the servers on the SSN, hoping that GIAC Enterprises will receive the blame for the mischief.

Rule 17 is the “catch-all” rule: Any packet that did not match one of the previous 16 rules should be dropped and logged long. The logging on this rule allows the administrators to see where additional rules may need to be added to the firewall’s rulebase, or where attackers attempted to break through the firewall into the internal network.

Note that although the rulebase is functionally complete, meaning that it has implemented all the rules set forth in the Corporate Security Policy, some room still exists for improvement, particularly in the area of performance boosting. Some of the rules might be arranged into more of a decreasing order, from most packets processed to least packets processed. The reordering would speed things along, considering that the closer to the top a rule is, the less processing power is required should the packet match the given rule. For example, instead of being at rule 11, if the web traffic rule was at rule 4, the significant amount of web traffic that any e-commerce web-based dot-com business that is doing well has would experience less seven less rules through which to process, thereby speeding up the entire operation of the firewall.

VPNs

Two distinct types of VPN appliances are used in this architecture: the Nortel Networks Contivity 2600 VPN switch and the Nokia IP530 with CheckPoint FireWall-1 and VPN-1. Discussion from **Assignment #1** has identified the Contivity box to be used for remote access by telecommuters, road warriors, and suppliers, and the Nokia to be used for site-to-site connectivity between GIAC Enterprises and Acquired Company, Inc.

Continued on next page

Assignment #2 – Security Policy, Continued

VPNs, continued

The discussion regarding the setup and configuration of the VPN devices continues from the previous page.

The setup for each box is unique. As the Nokia is using CheckPoint's VPN-1 solution, and knowing that it is somewhat integrated with the FireWall-1 product, I believe that the setup and configuration will be similar to the process described by Chris Brenton for the case study that used VPN-1 as remote access for users [8].

The Corporate Security Policy identifies the following key elements for VPN configuration:

- *All users from Acquired Company, Inc. (ACI) should have access to all GIAC Enterprises servers and information, according to authorized clearance levels.* In other words, the VPN should not prevent ACI employees from doing whatever their job requires on any server, or with any particular information, in GIAC Enterprises network. Of course, the assumption here has to be that ACI's network is as secure as the GIAC Enterprises network. If not, either GIAC Enterprises needs to lend them some network security support staff, or the VPN box will need to be tightened more securely than the Corporate Security Policy indicates.
- *All users from GIAC Enterprises should have access to all ACI servers and information, according to authorized clearance levels.* Vice versa of the first tenet.
- *Telecommuters and road warriors (sales force) need access to servers and data (again, according to authorized clearance levels) from their laptops and desktop systems via dial-up, DSL, or cable modem.*
- *Suppliers will connect with the same hardware via the same channels, but their access will be restricted to the internal Data subnet for depositing new fortunes (used by the custom-built supplier application to keep our product stored in a place where it is under a watchful eye).*

The choices I've made regarding configuration of the Nortel box are as follows:

- ESP encryption with DES, 3DES, or MD5 algorithms (3DES may have export restrictions to countries where our suppliers reside)
 - ISAKMP negotiation for key exchange and security association
 - Split horizon implementation
-

Assignment #3 – Audit Your Security Architecture

Task

You have been assigned to provide technical support for a comprehensive information systems audit for GIAC Enterprises. You are required to audit the Primary Firewall described in Assignments 1 and 2. Your assignment is to:

1. **Plan the assessment.** Describe the technical approach you recommend to assess your perimeter. Be certain to include considerations such as what shift or day you would do the assessment. Estimate costs and level of effort. Identify risks and considerations.
2. **Implement the assessment.** Validate that the Primary Firewall is actually implementing the security policy. Be certain to state exactly how you do this, including the tools and commands used. Include screen shots in your report if possible.
3. **Conduct a perimeter analysis.** Based on your assessment (and referring to data from your assessment), analyze the perimeter defense and make recommendations for improvements or alternate architectures. Diagrams are strongly recommended for this part of the assignment.

Note: DO NOT simply submit the output of nmap or a similar tool here. It is fine to use any assessment tool you choose, but annotate the output.

Assessment Plan

To assess the perimeter, I will need a team of four analysts to implement the plan described in the following paragraphs and document the results.

The technical approach will be to have three team members (testers) test each of the interfaces of the primary firewalls, border routers, and VPN servers for compliance with the corporate security policy and for vulnerabilities. The fourth member of the team (researcher) will be responsible for researching exploits for and patches/upgrades to remove known vulnerabilities for each of the devices, their OSes, and appropriate software. The testers will use the exploits that the researcher uncovers to verify that the vulnerabilities continue to exist.

Once the vulnerabilities have been identified, and the researcher has identified patches or upgrades that remove the vulnerability, a second test will be to have 2 different testers run a vulnerability scanner against all three parts of the network – externally available interfaces, the internal network, and the SSN. For the systems themselves, my tool of choice is Nessus Security Scanner (version 1.0.7a), available at <http://www.nessus.org/>. Nessus Security Scanner is an open-source solution that has recently been named a winner in the category of Vulnerability Assessment Tool for *Network Computing's* 7th Annual Well-Connected Awards held in Las Vegas on Monday, May 7, 2001. In addition, *Network Computing* has also given it highest grade on their vulnerability scanner report card (see article available at <http://www.networkcomputing.com/1201/1201flb1.html>). While not living up to all the expectations of the reviewers (one comment was "...it's a case of the best of the worst"), Nessus was able to identify 90% of the vulnerabilities correctly. For network-level scanning, nmap will work fine (although a second tool might help in backing up our claims).

Continued on next page

Assignment #3 – Audit Your Security Architecture, Continued

Assessment Plan, continued

The plan for assessing the vulnerabilities of the network and its systems continues from the previous page.

We should plan to carry out the scans at two different intervals on two different days, one without any users (during the “midnight” hours on a Friday night, after the backups have been run) and one with as small a group of users as possible (say around 7-8am on a Monday morning); these two times should provide a good range of opportunity for scanning without disrupting the business or its systems.

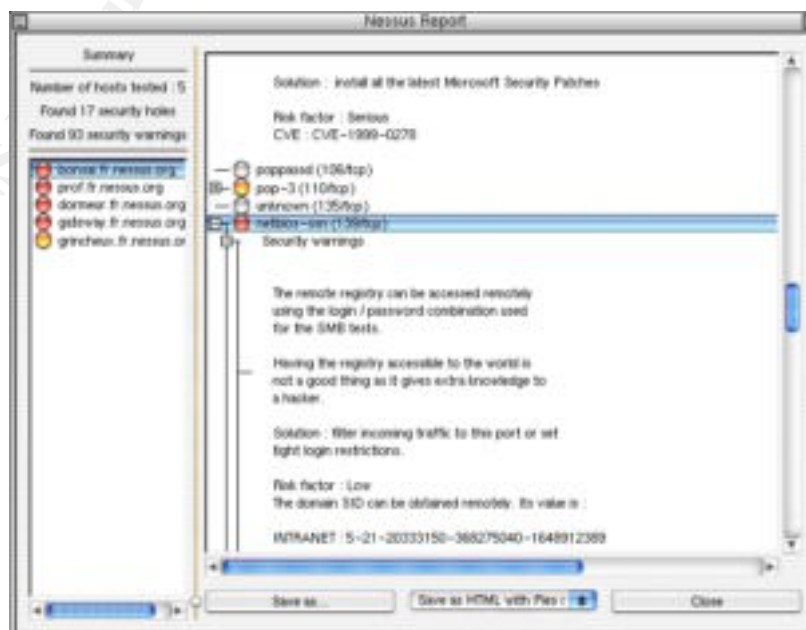
Our plan will have to be provided ahead of time for review by system and network administrators (to ferret out any dependencies that would be impacted by the scanning and the interaction with the hosts). Everyone will be notified as to the unavailability of network and system resources during the scanning process.

The final step will be to coalesce the collected data into a final report with findings, suggested courses of action to remove immediate vulnerabilities, and recommendations for tools, scans, and frequency intervals that administrators may use to keep the systems “in the pink”.

Implementation of the Plan

Once the plan has been reviewed and accepted by the appropriate parties, it then needs to be scheduled for implementation.

The testers first install the Nessus clients to the appropriate machines, then setup and run the scan. The output of the scan will resemble the following picture:



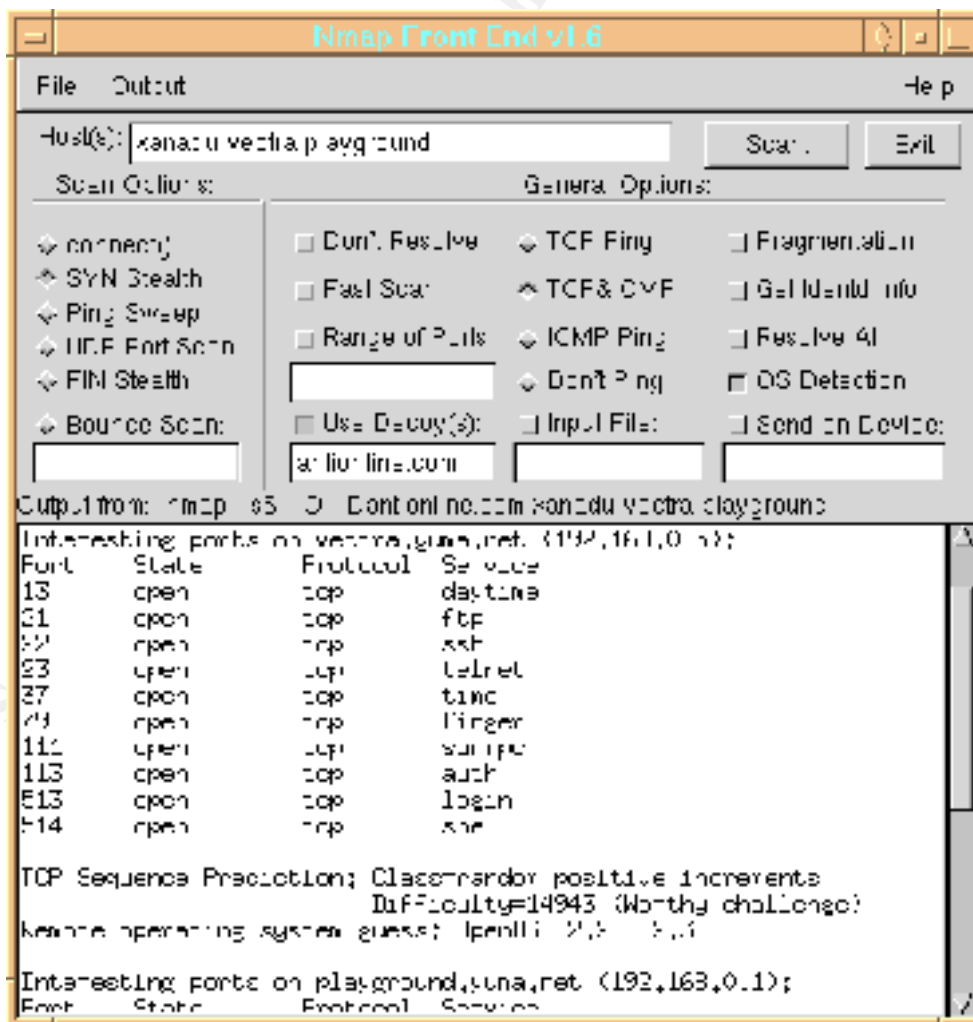
Continued on next page

Assignment #3 – Audit Your Security Architecture, Continued

Implementation of the Plan, continued

This sample output (from <http://www.nessus.org/demo/third.html>, where sample report formats with the complete report are also available) was run against only five hosts, but as you can see, it identified a serious risk at the top of the window that suggested the latest Microsoft Security Patches were not installed on that particular host. In addition, a low-risk vulnerability is identified, with a suggested solution for filtering incoming traffic to the netbios-ssn (139/tcp) port or setting tight login restrictions. Particulars of the specific test that generated these results can be located at <http://www.nessus.org/demo/first.html> and <http://www.nessus.org/demo/second.html>.

Next, the testers will need to run tests with nmap from outside the border routers, from inside the corporate network, and from each of the SSNs. A sample of the nmap output, shown in the following picture, indicates the results of TCP SYN packets sent in stealth mode to three specific hosts:



Continued on next page

Assignment #3 – Audit Your Security Architecture, Continued

Implementation of the Plan, continued

The implementation of the plan continues from the previous page.

The output from the nmap scan is a list of the open ports (closed or filtered ports are not called out by number), the TCP sequence prediction calculation, and a guess of the remote operating system. More options are available, with more information presented.

And the final piece of the implementation would be to cull all the information, including the research done by the researcher, into a series of reports, graphs (possibly), and recommendations – both short-term and long-term. The final presentation of the findings should also include an executive summary for senior level IS/IT management, as well as intimate details for the security, network, and system administrators.

Perimeter Assessment

The team's assessment of the perimeter identified several key points, listed as follows:

- Two different types of stateful inspection firewalls might be more work than necessary, if the vulnerability scans are continued, and the patches/upgrades are made in a timely fashion. Implementation of the corporate security policy is fine on both systems, but idiosyncrasies between the different appliances make some unnecessary differences between the rulebases. In addition, recovery may be easier with an update, working copy of the exact configuration of the appliance.
 - Caution with programmers, especially the CGI and web-based application ones. Make sure to run HTTP and CGI vulnerability scanners as well (due to the fact that our revenue is based primarily on our selling of the fortunes via an e-commerce web-based application). And stay away from Windows NT on the SSN, if possible, with significantly more exploits available and more black-hats looking for new ones that have not been patched yet.
 - VPN devices and placement could also be improved. There may be less need for the application/proxy firewalls in the design except in front of the Data subnet if the VPN devices were placed in front of the firewall.
 - No anti-virus scanning appeared to be present except on the SMTP mail gateways on the SSNs. Web surfing, FTP downloads, and other normal user activities are also vulnerable to viruses. The placement of an anti-virus filter should most likely be directly behind the firewalls.
 - No ICMP at all may be too strict on the border routers. Some VPN implementations will not work correctly without some communication via this channel. If rules are relaxed on the border routers, then more careful monitoring and logging will need to occur, so that GIAC Enterprises does not fall victim to a smurf attack..
-

Assignment #4 – Design Under Fire

Task

The purpose of this exercise is to help you think about threats to your network and therefore develop a more robust design. Keep in mind that the next certification group will be attacking your architecture!

Select a network design from any previously posted GCFW practical (<http://www.sans.org/giactc/gcfw.htm>) and paste the graphic into your submission. Be certain to list the URL of the practical you are using. Design the following three attacks against the architecture:

1. An attack against the firewall itself. Research vulnerabilities that have been found for the type of firewall chosen for the design. Choose an attack and explain the results of running that attack against the firewall.
2. A denial of service attack. Subject the design to a theoretical attack from 50 compromised cable modem/DSL systems using TCP SYN, UDP, or ICMP floods. Describe the countermeasures that can be put into place to mitigate the attack that you chose.
3. An attack plan to compromise an internal system through the perimeter system. Select a target, explain your reasons for choosing that target, and describe the process to compromise the target.

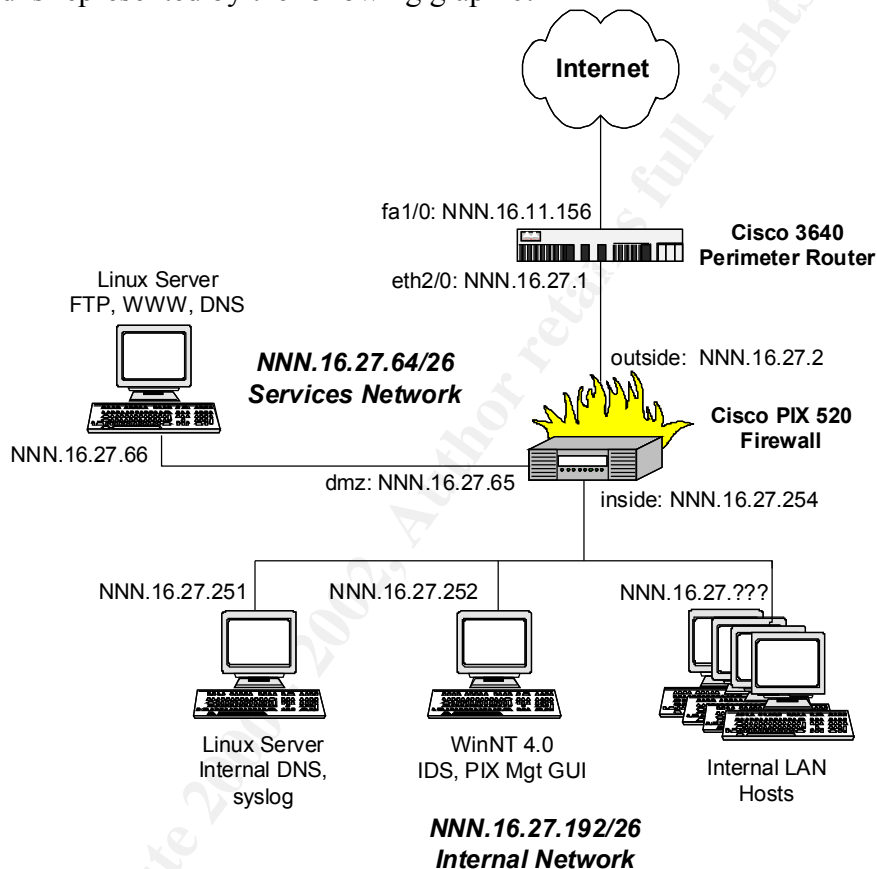
Note: this is the second time this assignment has been used. The first time, a number of students came up with magical "hand-waving" attacks. You must supply documentation (preferably a URL) for any vulnerability you use in your attack, and the exploit code that you use to accomplish the attack. The purpose of this exercise is for the student to clearly demonstrate they understand that firewall and perimeter systems are not magic "silver bullets" immune to all attacks.

Continued on next page

© SANS Institute 2002, Author retains full rights.

Assignment #4 – Design Under Fire, Continued

Network Design The network design that I have chosen to attack was submitted by Adam Payne in August 2000 [9], and is represented by the following graphic:



Firewall Attack Adam has used a Cisco PIX 520 firewall in his design. Without knowing the exact version of PIX that he has implemented in his design, I have deduced from his References list that it is very likely version 5.1.

After searching on www.securityfocus.com for known issues with version 5.1, I decided on a very simple attack which seems to have carried over into version 5.2, but has yet to be identified as a vulnerability in version 5.3.x, the “SMTP Content Filtering Evasion Vulnerability”. The SecurityFocus discussion from <http://www.securityfocus.com/vdb/bottom.html?section=discussion&vid=1698> is reproduced below for further clarification:

During communication with an smtp server, if the "data" command is sent before the more important information is sent, such as "rcpt to", the smtp server will return error 503, saying that rcpt was required. The firewall, however, thinks everything is alright and will let everything through until receiving "<CR><LF><CR><LF>.<CR><LF>". It is then possible for the attacker to do whatever he wishes on the email server.

Assignment #4 – Design Under Fire, Continued

Firewall Attack, continued

The description and discussion of the Firewall Attack continues from the previous page.

An example of the exploit code used for this type of attack follows:

```
helo <mail server name>
mail from: nobody@nowhere.com
data          (from here on, pix disables fixup)
expn guest    (now I could enumerate user and have access to all commands)
vrfy oracle
help
<whatever command I want>
quit
```

The result of running this particular exploit is, in the case of my exploit code, nothing very malicious but rather most informative. Remember that any commands passed through this exploit will be executed on the PIX firewall, not the SMTP server. Nevertheless, a potentially dangerous proposition for such an easy exploit.

DoS Attack

I choose to attempt a DDoS attack with a UDP flood, called a “fraggle” attack (lesser known cousin to the very popular ICMP flood or “smurf” attack).

In order for my attack to work, I will need to use my 50 compromised DSL/cable modems as my amplifying network. Sending a stream of spoofed UDP packets to the broadcast addresses of the amplifying network on port 7 (echo) will cause a flood of responses to be sent to the victim’s address. Basically, my assumption is that Adam’s network (which does not specify the size of his link) is connected to the Internet .by something about the size of a T1. The 50 ~500Kbps DSL/cable modems combined are significantly larger than that. Therefore, multiple reply packets (ICMP unreachables or UDP echo responses) will flood the victim’s pipeline, choking off any other traffic – in effect, denying all other services.

The simplest countermeasures that can be put in place to mitigate becoming a victim of this type of attack are to drop all ICMP packets from entering your network at the border router, and to supply ample amounts of bandwidth to your network. Much more effort should be placed on preventing becoming part of an amplifying network for this type of attack. `No ip directed-broadcasts, and no icmp any any (smurf only)` directives should be added to the access-list for the external interface of the border router.

Continued on next page

Assignment #4 – Design Under Fire, Continued

Attack Against Internal System

Let's assume that I want to deface Adam's web site. He has a Linux server in his SSN that is running HTTP, FTP, and DNS services. Taking into consideration that BIND is one of the largest security headaches for a security administrator, I would attempt to attack BIND, with the ultimate goal of owning the box and posting my own page in the stead of Adam's home page.

The basis for my attack against BIND will be the information from CERT Advisory CA-2001-02 "Multiple Vulnerabilities in BIND" (posted at <http://www.cert.org/advisories/CA-2001-02.html>). The Advisory lists 4 known vulnerabilities; depending on the version of BIND that Adam is running in the SSN, I may be able to use either 1 BIND 8 or 2 BIND 4 exploits to cause the execution of code and take over the box.

Once the box is owned, I can then cover my tracks and install my web page(s) with the same name(s) as Adam's.

Unfortunately, without DNS, the Internet would not be the giant communications vehicle that it is today. The best way for Adam to guard against this attack would be to (1) update his version of BIND to 4.9.8 or 8.2.3 or 9.x, and (2) divorce the DNS from the web server and put it on its own server. The second recommendation would not prevent an attack using the vulnerabilities from the CERT advisory, but it would add another level of depth to the defense (the vulnerabilities of one server are not necessarily the same vulnerabilities of a second, especially if the OSes are different). Each box could be locked down and armored according to the resident OS and software and the service(s) it provides.

© SANS Institute 2000 - 2002

End Notes

- [1] Spitzner, pp. 204-242.
 - [2] The SANS Institute, "SANS Network Security Roadmap 2001".
 - [3] Eldridge, pp. 2.
 - [4] Spitzner, p. 58.
 - [5] Cisco, "Improving Security on Cisco Routers", pp. 2-16.
 - [6] Eldridge, pp. 2-11.
 - [7] Keeney, pp. 1-2.
 - [8] Brenton, pp. 129-152
 - [9] Payne, p. 4.
-

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

References

- Author Unknown. "SANS Resources – How to Eliminate The Ten Most Critical Internet Security Threats". Available at <http://www.sans.org/topten.html>.
- Brenton, Chris. *VPNs and Remote Access*. Presented at SANS New Orleans 2001 on 31 January 2001.
- Cisco Systems, Inc. "Building a Perimeter Security Solution with the Cisco Secure Integrated Software". Available at http://www.cisco.com/warp/public/cc/pd/iosw/ioft/iofwft/tech/firew_wp.htm. Posted on 17 September 2000.
- Ibid. "Improving Security on Cisco Routers". Available at <http://www.cisco.com/warp/public/707/21.html>.
- Eldridge, Brett. "Building Bastion Routers Using Cisco IOS", *Phrack Magazine*, 9(55). Available at <http://www.routergod.com/bastion/bastion.html>.
- Higgins, Kelly Jackson. "Norfolk Southern Gets E-Business on Track". *Network Computing*, 11(24):96-8.
- Ibid. "Tunneling Through VPN Security". *Network Computing*, 12(7):87-9.
- Ibid. "Voice Systems Overhaul Nets Wireless, IP and Analog". *Network Computing*, 12(6):75-7.
- Keeney, Frank. "Untitled". Available at <http://pasadena.net/cisco/secure.html>. Posted on 30 December 1998.
- Kelly, Brian M. "GIAC Firewall And Perimeter Protection Curriculum Practical Assignment". Available at <http://www.sans.org/>.
- Langley, Richard. "Securing Your Internet Access Router". Available at <http://www.sans.org/infosecFAQ/firewall/router.htm>. Posted on 23 January 2001.
- McClure, Stuart, and Joel Scambray. *Hacking Exposed: Network Security Secrets & Solutions*. Berkeley, CA: Osborne/McGraw-Hill, 1999.
- Moskowitz, Robert. "Holding a Defensive Line". *Network Computing*, 11(24):43.
- Payne, Adam. "SANS GIAC Firewall and Perimeter Protection Practical Assignment". Available at http://www.sans.org/y2k/practical/Adam_Payne.doc. Posted in August 2000.

Continued on next page

References, Continued

SANS Institute, The. "SANS Network Security Roadmap 2001", 4th ed. Winter 2001.
Created by Michele D. Guel.

Spitzner, Lance. *Advanced Perimeter Protection and Defense In-Depth*. Presented at SANS New Orleans 2001 on 30 January 2001.

Spitzner, Lance. "Building Your Firewall Rulebase". Available at
<http://www.enteract.com/~lspitz/rules.html>. Last modified on 26 January 2000.

Stelzner, Jeff. "SANS GIAC Level 2: GCFW – Firewalls, Perimeter Protection, and VPNs Practical Assignment". Available at
http://www.sans.org/y2k/practical/Jeff_Stelzner_GCFW.doc. Posted on 19 May 2001.

Stevens, W. Richard. *TCP/IP Illustrated, Volume 1: The Protocols*. Boston, MA:
Addison-Wesley, 1994.

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