Global Information Assurance Certification Paper

SANS GIAC Firewall and Perimeter Protection Practical

SANS at Monterey 2000

Submitted by: Jeremy Brown
Date: November 22, 2000

I have broken this document down into three sub-sections, which relate to the three practical assignments. The first section addresses the security architecture of the network. This sections takes on the physical aspect of a secure network, which gives a birds-eye view on what devices go where, and what function they provide. The second section is one that deals with, not so much the physical aspects, but the logical aspects of the network. This being, filter router Access Lists, a firewall rule base and other bits of info, that restrict, access to critical data. The final portion of the paper is dedicated to testing and auditing the security implementation.

Part 1: Security Architecture

The basic layout of the network architecture is located below in **fig. 1.1**. The diagram shows there is one main gateway. This is a filter router, and it provides a simple redirection of traffic (probably a high-end Cisco 4000). If an external network or host is trying to access one of the external services, such as the mail server, or web server, it is passed to the insecure network (which I will address in detail later). If the incoming address is that of the partnership's IP (found out through the application of a standard ACL), it will be forwarded to the firewall, where IPSEC is applied. If the incoming traffic is that of an allowed service, such as a few specified TCP and UDP protocols it is allowed to pass through, to the firewall. However, certain extended ACLs have been added to ensure a bit more comfort (denying such services as LDAP and RIP); and other items that have no reason to be entering our network.

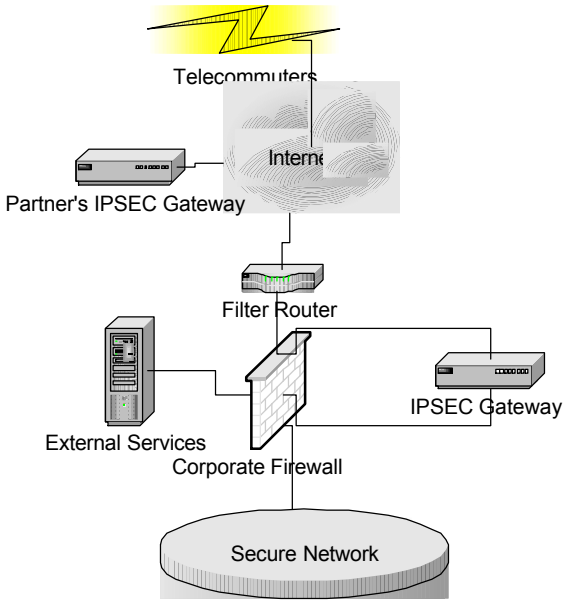

fig. 1.1

As we took note of earlier, the external services network (or insecure network), provides services for the outside. (**fig. 1.2**) The requested services are forwarded through the filtering router and to the firewall. The firewall goes over its basic rule base to ensure that no tampering goes on with the HTTP, mail and FTP servers. In addition to the logging that occurs at the firewall level, there is additional logging provided at each of the servers (i.e. all connections made to the FTP server are logged). The stored data that resides on each of these servers, is encrypted using GnuPG. http://www.gnupg.org.

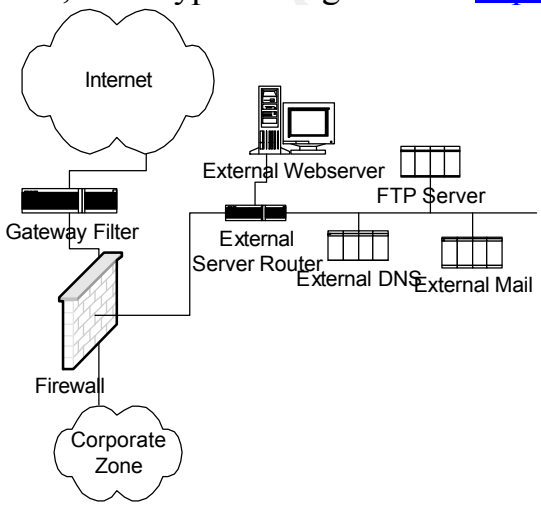

fig. 1.2

This encryption is just to ensure that, if indeed root is compromised, credit card

numbers, names, fortunes etc, are undecipherable.

The transactions that occur over the web, are highly susceptible to attacks. Information such as phone numbers, names, and credit card numbers are encrypted using SSH. SSH is implemented on the HTTP server (external web server).

This next diagram shows the VPN / Partnership connections to the internal network. The two IPSEC Gateways are displayed also. The gateways provide encryption/decryption for the data that passes between the two networks. The traffic, after being allowed to pass through the filter router, is forwarded by the firewall to the proper IPSEC node (note: see IPSEC server in figure 1.3). The IPSEC server just carries the info to establish an active SA for the two security gateways. **There are two entries in the security gateway's SAD. One entry is for employees, and one for the partnership's connection. This is where the address is translated and forwarded to that specific user's

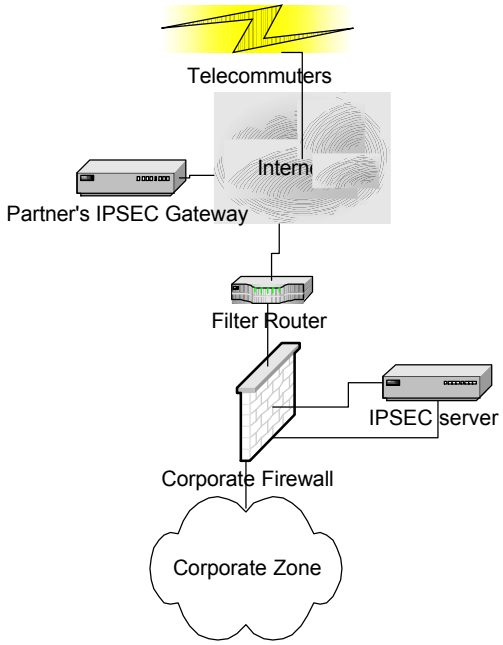

fig. 1.3

In order for the Corporate firewall to allow any traffic to flow inside the internal network,
it must have an active SA established. In addition to this, the firewall will only allow services like: FTP, DNS and Telnet, to be allowed through, if, they have already been initiated by a host on the internal network.

Close-up of corporate topology:

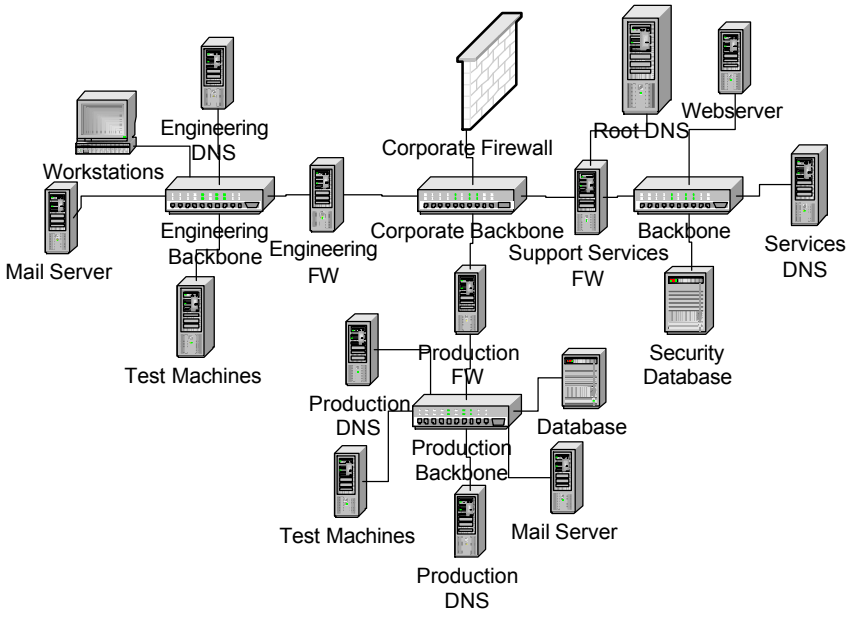

Fig 1.4

In this diagram, most of the applications that are used by all users, are located at the access layer of this network. There are also firewall's located at each BU's gateways.

This restricts access and unwanted traffic from flowing in and out of the network segment (provides services at a need-to-know basis). The databases located in the Production's network segment are all encrypted, as with the security database on the Support Service's network. Support service network also provides auditing capabilities, and checks patch updates for various WIN X and Linux servers. Support also does virus updates and check CERT warnings daily.

Part 2: Security Policy

This network security policy's goal, is to ensure that vital data be protected from any prying eyes, and any services deemed to be a threat to critical data denied. In addition to this the DMZ, and outmost areas of the network should be monitored very closely for warning signs. By warning signs I mean, taking precautions before security has been compromised, and data stolen. There are four basic security layers, which will be discussed in this order. The first layer of security is the filter router. This is the main gateway to the outside, and contains a couple egress ACLs and a bunch of extended ACLs.

The next layer of security is the corporate firewall. Depending on how the firewall deals with the traffic, the layer three security may or may not be implemented. The third layer of security, which is applied to all remote employee connections, and all partnership connections, is the IPSEC service. Depending on which two types of traffic, the packets could be forwarded two one of two IPSEC nodes. Each node contains a SAD for establishing two SA's. Finally, the last layer of defense is the internal firewalls. These firewalls are inside the "secure" perimeter. They divide the business into 3 BU's and restrict information on a need-to-know basis.

As I mentioned earlier, the Boundary router (a Cisco 4000) is responsible for distinguishing traffic and forwarding it appropriately. The router has a main gateway connection (denoted as Serial 1/1), with an IP address of 192.21.71.1. The data arriving to this interface is immediately put through a series of ACLs.

Boundary Router Configuration:
No ip direct-broadcasts
No ip source-route
No service Finger
No ip http
No ip bootp
Banner //// Warning Unauthorized Usage of Fortune Maker's Network, Will Result in
Arrest and Possible Conviction /////
Access-list extended egress permit ip 192.71.21.0 255.255.255.0
Access-list extended egress deny ip any any logging

Interface serial 1/0
Ip address 192.21.71.1 255.255.255.0
access-group 100 in
access-list 100 deny udp any any eq 137
access-list 100 deny udp any any eq 138
access-list 100 deny tcp any any eq 139
access-list 100 deny tcp any any eq 109
access-list 100 deny tcp any any eq 110
access-list 100 deny tcp any any eq 111
access-list 100 deny udp any any eq 111
access-list 100 deny tcp any any eq 143
access-list 100 deny udp any any eq 520

access-list 100 deny tcp any any eq 389
access-list 100 deny udp any any eq 389
access-list 100 permit tcp any any
access-list 100 permit udp any any

If the arriving packets do not correspond with any of the deny parameters, it is forwarded to the address of the firewall, which is 192.71.21.1. Remember, this filter router has only a few major functions. First, to deny specific services, that which should never be allowed to enter the network. Second, to provide some protection from packet spoofing (egress filtering), and third to by default allow everything else to be dealt with by the firewall. The router starts at the top the rulebase and works itself downward. Each filter can be applied by using simple command line entries (as you see in the example above).

I should probably go over a few of the configuration lines so we get the idea.

Access-list extended egress permit ip 192.71.21.0 255.255.255.0
Access-list extended egress deny ip any any logging

These lines just verify that the outgoing packets are from our network's firewall, and not some sort of spoofed packets.

access-list 100 deny udp any any eq 137
access-list 100 deny udp any any eq 138
access-list 100 deny tcp any any eq 139

These entries to the access list deny netbios services to our network. This prevents the 911 Worm and exploitation of windows systems.

The final rules on the rule base:

access-list 100 permit tcp any any
access-list 100 permit udp any any

are applied last. These are applied if the traffic does not fit any of the above specified requirements.

The corporate firewall is the main area of defense for our perimeter. The firewall we have to set up for Fortune Maker's, is a PIX firewall. This firewall will be running NAT. Name Address Translation, will allow all our services to appear as they are coming from this NAT box, when in fact they are sourced from a ten net address. This provides a higher level of security, by not giving attackers any additional information concerning our network. The only external addresses attackers could see, would be in the external DNS which, by use of split DNS, only lists necessary addresses (Mail, HTTP and DNS server) . In addition to this, the PIX box will be retrieving information about its IPSEC SAD.

The firewall should by default, deny all traffic, and allow services from very specific ports. So, we know our final rule on any rulebase should be to drop all packets, also known as the lockdown rule.

First off, our PIX has one incoming interface and one outgoing interface. The incoming interface is for all traffic flowing into the network. The outgoing interfaces are where we apply our rules. We can define three levels of security, which relate to the different networks we have.

Pix configuration:

hostname pixfirewall
interface ethernet0 auto
interface ethernet1 auto
ip address outside 192.71.21.2 255.255.255.0
ip address inside 10.1.1.1 255.255.255.0
ip address services 192.71.28.1 255.255.255.0
ip address ipsec 10.2.1.1 255.255.255.0 See External Services Note
static(outside, services) 192.71.21.2 192.71.28.1 netmask 255.255.255.0
syslog level 5
conduit permit tcp host 192.71.21.2 eq www any
conduit permit tcp host 192.71.21.2 eq 1080 any
conduit permit tcp host 192.71.21.2 eq ftp any
conduit permit tcp host 192.71.21.2 eq ssl any
conduit permit tcp host 192.71.21.2 eq dns any
conduit permit udp host 192.71.21.2 eq dns any
static(outside, inside) 192.71.21.2 10.1.1.1 netmask 255.255.255.0
syslog level 7 See inside Note
conduit permit tcp
conduit permit tcp host 192.71.21.2 eq tcp established-port
conduit permit udp host 192.71.21.2 eq udp establish-port

route outside 10.0.0.0 255.255.255.240 172.17.11.3 1
static(outside,ipsec) 192.71.21.2 10.2.1.1 netmask 255.255.255.0 See IPSEC Note
crypto isakmp policy 2

authentication pre-share
hash md5
crypto isakmp key inside address 10.2.1.1
crypto isakmp key outside address 192..0.20

no rip outside passive See Firewall Services Note
no rip outside default
no rip inside passive
no rip inside default
no rip services passive
no rip services default
no snmp-server location
no snmp-server contact

By default, the PIX box will drop any packets that are not a part of the rule base. However, to better address security issues concerning commonly probed services, such as port 137 probes, we allow high level logging on some of the interfaces.

External Services
This section of the rule base defines what is allowed to go into the external services network. The protocols I have specified to be allowed to traverse the firewall are:
HTTP (port 80 and 1080), FTP, SSL and both ports for DNS. These are the only services that will be allowed to pass onto the 192.71.28.1 segment. There is a pretty good amount of logging done on this interface. The log information will tell basic login information and a log of services denied and applied.

Inside
The portion of the config responsible for controlling the incoming traffic to the internal network is pretty simple. It provides a minimal amount of incoming services. The only services it allows through are established UDP and TCP connections. That is, services such as FTP and HTTP that do not have a SYN flag enabled. This means that the firewall will check to make sure that the traffic was initiated by an internal host, and the packets are not just any old traffic wandering inside.

IPSEC

The IPSEC gateway must utilized a couple protocols to ensure that authentication and encapsulation occurs for all VPN traffic. The command lines concerning ISAKMP, are the required protocols for this to take place. (Introduction to IPSec--http://wwwin.cisco.com/cmc/cc/so/neso/sqso/eqso/dplip_in.htm) These are the only protocols that have been enabled, since traffic must be encapsulated with IPSEC to be allowed through.

Firewall Services

These couple lines in the configuration (such as no service snmp), are to deny certain services the PIX should not be providing to any hosts. These services include RIP routing and SNMP.

There is yet another layer of protection which comes after IPSEC and the Corporate Firewall. This the BU's firewalls. There are three FW-1 implementations at each network node. A close-up of the network architecture is below (**fig. 2.1**), to give you a better visual idea of what we are talking about.

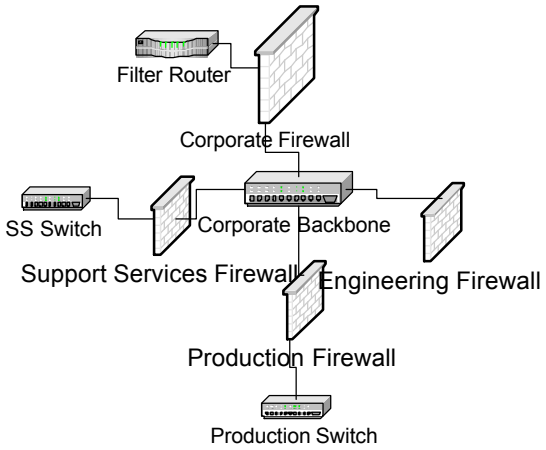

fig. 2.1

You see, the Corporate Backbone, is divided into three distinct VLANs. These VLANs in turn, have FW-1's behind them. Below, I have defined the rule base that will be implemented on layer 4 of the security architecture.

Rule No.	Source	Destination	Service	Action	Track
1	Firewall-admin	Firewall	Firewall1	Accept	Long
2	Internal	External	DNS	Accept	Short
3	Internal	External	HTTP	Accept	Short
4	Internal	External	SSH	Accept	Long
5	External	Internal	All	Deny	Short

The different segments do have varying rule bases that disallow traffic coming

in from any other BU's. The above rulebase is a generic rule base that will be applied to all layer 4 firewalls.
The three network segments: Fortune Engineering, Production and the Services network, all add a higher degree of security due to a second layer of firewalls. In addition to more security, this solution also limits information on a need-to-know basis. It performs filters that are not specified by the filter router, and the corporate firewall.

Part 3: Auditing Your Security Architecture

Perhaps the most important step to implementing a security architecture is, verifying the validity and functionality of your security system. To do this, we must first plan the audit, otherwise, things might get crazy.

Our first step of planning is deciding on a time and logical location for the audit to take place. I believe, it would be wisest, to do the audit on a weekend. This is primarily for the fact, that port scanning and password cracking etc. can be very disturbing for network performance. What I mean by logical location for the audit, is "where in the architecture will I set my terminal?" To be most thorough, we will start the major audit with the NMAP scanner just outside the firewall, to test its rulebase. Then, we will do scans from beyond the filter router, to test it's ACLs, and how they work in accordance with the PIX rulebase. After the initial audit takes place, the security admins from the Services BU, do bi-weekly audits of the architecture.

After the first step of planning has been completed, we can now think about more realistic attacks. These should take place during work hours, as to simulate an actual attack with real conditions. The logical location of this attack should definitely be outside the perimeter. The process should test the IPSEC gateway as well as normal network traffic.

Now, how do we actually go about to doing an audit?

The initial audit should consist of several elements. First a port scanning test, to show which vital ports have been left open. A NMAP scan on the Corporate firewall will give us some information that we already new:

PORT	STATE	SERVICE
21/TCP	Open	ftp
25/TCP	Open	Smtp

53/UDP	Open	DNS
53/TCP	Open	DNS
80/TCP	Open	http
1080/TCP	Open	http
443/TCP	Open	Ssl

These are the services provided by our external network(ie. SSL, Web services, mail and ftp). When doing this process, a complete port scan of the entire perimeter should be included. Since NMAP is freely available at: http://www.insecure.org/nmap, cost of the initial audit should not be too bad. However, this is assuming you can find a lab administrator not to busy to manually scan the network in his/ her spare time.
At the time of the initial audit, a couple meetings should take place, establishing precedents for security. Logs should be checked every day on the filter router, and all four firewalls. These logs should contain all the information necessary for deciphering whether or not our E-business is getting attacked. Also, a security administrator should be appointed to work in the Services BU, checking for new patches from:
sunsolve.sun.com In addition to monitoring Suns patch updates, Windows updates will be watched for patches also (www.microsoft.com/technet/security). CERT advisories will be viewed also.
From the time of the initial audit and at every bi-weekly audit after, a password cracker such as L0PHTCRACK should be used on the password databases. Too often have I seen passwords compromised because of overlooking weak passwords. (http://www.l0pht.com/l0phtcrack/register.html) The fact is that buying a license for this is cheap enough (in the ballpark of 100-800 dollars) that it serves its purpose.

There are other various and sundry ways to check the network for critical information. One of these, is a utility called Snort (available www.snort.com). It is a packet sniffer, like tcpdump, that will allow our auditors to view traffic going between specific network nodes. An example of where we might want to use this would be between the Production Firewall and the External webserver. That is, this would allow us to ensure that no information such as credit card numbers could be seen in plain sight to anyone monitoring the backbone switch.

Another perspective we can take on the security audit is that of a "black hat." We can approach from outside the perimeter and first try a couple brute force attacks on the external servers. This means trying to authenticate oneself one the http server, ftp server and ssh server. If that does not heed to our

expectations, we can test the IPSEC gateway.

The small bi-weekly audits hit the network from a couple angles. As I mentioned earlier, a password cracker such as L0PHTCRACKER, will be used during these audits. An automated script, with full logs will go through the process of doing external scans of the network. The script will also check the validity of the Fingerprints generated by MD5 from the IPSEC gateway. A new shared, secret will be implemented for the partnership connection coming into our network. This is to try and prevent the IPSEC gateway from being compromised.

There are definitely some obvious weaknesses with the security architecture. No network is totally secure. In this scenario, the external or, services network is fairly susceptible to attacks. That is, the external mail server is vulnerable.

Also, the VPN traffic that has to be allowed through, is only as secure as the encryption and authentication from the IPSEC server. Any IPSEC node, could attempt to authenticate itself on our server.