Global Information Assurance Certification Paper

Author: Patrik Sternudd

GCFW Practical Assignment for SANS Network Security 2000

1 Index

2 Introduction

This document is intended as th e practical assignment required for passing the GIAC GCFW certification provided by the SANS Institute. All information and solutions are based upon the by SANS given scenario. Of course, ideas and designs may work well enough in real cases. But personally I feel the chances to stumble upon a company that would actually earn $200 million per year on selling online fortune cookies are pretty insignificant. But who knows? Similar situations might arise.

I am not perfect myself and I know my knowledge within some aspects of IT Security is scant; I work more with firewalls and intrusion detection systems than penetration testing or policy writing. Since the field is so wide, I do not think many people are specialised in more than perhaps a couple of areas. If y ou are, well, congratulations, you probably got a well -paid job at least. For those who are not, the expectation I have is that you have fair knowledge of TCP/IP and firewalls (and various operating systems as well, but that should go without saying, I hope). You should also be familiar with common acronyms and port numbers for services as HTTP and DNS. Otherwise, you have to risk not knowing what I am talking about most of the time. That said, I hope you will find this document useful, I laboured long and hard (didn't we all?), but I also learned some new things and got a chance to think into greater detail; it is not every day you are hired to build a security infrastructure from scratch.

A difficulty I found was to keep the text within the right context. As a result I tend to change between them. In one instance, I write as I were the GIAC student discussing the problem, at others as a consultant or even a team of consultants hired by GIAC Enterprises.

A note on the IP addresses shown through this docume nt. Unfortunately, I had not the possibility to set up any testing environments using real IP addresses. Instead, I am using addresses from the private networks (RFC 1918) to symbolise different nets.
Below is a table showing which networks I am using, an d for what.

Network	Symbolise
10.0.0.0/8	The Internet
10.100.1.0/24	The network where the external IF on the firewall is
10.200.2.0/24	The network containing the evil hacker
192.168.1.0/24	Screened Network (DNS Server, etc)
192.168.2.0/24	Presentati on/Service network.
172.16.0.0/24	The network where the firewall console is
10.300.20.0/24	Our ISP's DNS resides on this network.

Any hostname resembling that of an existing company or organisation are coincidal; they are picked to make it easier for t he reader to understand who is doing what against whom.

3 First Assignment: Security Architecture

"Define a security architecture. The goal of your policy is to use filtering routers, firewalls, VPNs and internal firewalls to rapidly implement the VISA "Ten Commandments" to the extent possible at GIAC Enterprises, a new Internet Startup that expects to earn 200 million per year in sales of online fortune cookie sayings."

3.1 Preface

Some may object that my recommendations are merely overkill paranoia, way too expensive, not feasible, etc. It may be so, especially if a smaller company attempted this design. But in this case, I think this could prove a good solution. My reasons for this are found in the description of the company: " *a new Internet Startup that expects to earn 200 million per year in sales of online fortune cookie sayings"* . First, GIAC Enterprises expect to earn a lot of money. Therefore, they should afford to build a secure infrastructure. Second, since all revenues will come from online purchases, a security breach where an intruder might render the systems useless, corrupt data (for example, consider fortune cookie sayings containing harassment, whether it be sexual or religious), or theft of proprietary source code and/or information may prove even more costly than a sound security architecture in the first place. The whole business depends on the ability to deliver a qualitative service; failure to do so could lead to a decrease in sales or even worse, lawsuits (considering harassment or similar scenarios).

3.2 Architecture: A Brief Overview

There are five different logical areas. I will describe them only briefly before I into greater detail on each of them.

- The first area is where the business infrastructure resides.
- The second area contains workstations and servers for developers, administrators and executive personnel.
- In the third can be found workstations that are at a different geographical location due to a recent acquisition of a smaller company.
- The fourth area connects all hosts in the first and second areas to a separate network. Syslog messages will be sent to the Syslog server on this network. All other traffic will be dropped. The drawback is that all hosts are dual homed so it is important all routing are turned of between the interfaces. Neither of the last two networks will have connectivity to any other network, the Internet included!
- The last one, the fifth area exists at the same location as the first, second and fourth; it contains a network where the NIDS (Network Intrusion Detection System) sensors and monitors are located.
- One of the VISA Top10 is "Encrypt data sent across networks". However, this will not be done in this design. This is motivated by the use of numerous Intrusion Detection Systems. If all data were encrypted, the IDSes would all be rendered useless. It is my feeling one has to do this decision in every environment and try to decide which is the lesser evil. In this case I consider unencrypted data on the inside of the firewalls to be the lesser evil compared of not getting any advanced warning of Intrusions, or not being able to follow an occurring intrusion.
- Network Address Translation will be used, even though GIAC Enterprises has enough addresses to be able to do without. The advantage of this is that if the firewall would go

down, the servers would not be reachable since the firewall is the host translating. This eliminates the scenario where an intruder might attempt Denial Of Service attacks to bypass security.

3.3 A Commentary on the Network Diagrams

It is important to understand the diagrams do not necessarily show all hosts on the networks. For example, I have chosen not to draw the firewall management stations. In cases such may exist, they should be connected to a separate interface on the firewall. Additionally, there may be further services needed on the "Service Network" (see diagram later) that the company in question forgot to forward to the designer. Should such issues arise, it may be included in the design and physical network later. In fact, the company gave very little information on what they actually were going to use, so I have made a few additions on what I thinks may be necessary for the function of a company like GIAC Enterprises. In addition, as time goes, new services may (and most certainly will) be needed.
Another thing I would like to point out is that the diagrams do not show the count of hosts or firewalls. Where one HTTP-server can found on the diagram, there may be two, or even a cluster of ten servers in the physical world. This is for the company to decide, probably depending on traffic and load. The same applies to the firewalls. They may also be clustered for load balancing or redundancy. Excluded from this at the moment are the packet filtering firewalls behind the border router since they are faster and require less performance than an application layer firewall.

I also should mention that my choice of server platform is Sun Solaris running on Sparc. It is stable, easy to configure, quite scalable in hardware, and Sun Microsystems offers good service fast (if you have the money, that is). It can also be sufficiently stripped down (away goes X) and hardened, although it takes some effort. But it can still be done (compared to some other OSes I've made acquaintance with, although this is not supposed to be a "other vendor bashing session"). If you are more comfortable with Microsoft/Linux products, then use them. If I were to install an NT server, it would not be very secure since I work with Unix and not NT.

3.4 Border Routers

This design uses three Border Routers (of course, it is quite hard to have an Internet connection without having a router, but here those routers are deployed by the company and not the ISP). They implement the first layer of the defence, the first obstacle in the way of the wily hackers. Although very grand sounding, they do only a few tasks. Routers are not designed as firewalls and they should not be used as such. My design gives them the following functions:

- Drop all ICMP-types not needed for the operation of the networks (that is, administrators on the inside will be able to use ping and traceroute, but not those on the Internet. MTU discovery will be allowed as well as Unreach types coming from the outside). In the first environment (business front-end), only MTU discovery will be accepted.
- Drop any and all NetBIOS types.
- Perform spoofing counter-measures (no internal addresses from the outside and no external addresses from the inside, the later being both a courtesy to the community and protection against liability claims)

Note that the above requirements not should be considered as a full router access list, this topic has already been covered by previous classes, so I feel it is too much out of the scope.

3.5 Architecture: First, Second and Thi rd Areas

3.5.1 Network Diagram

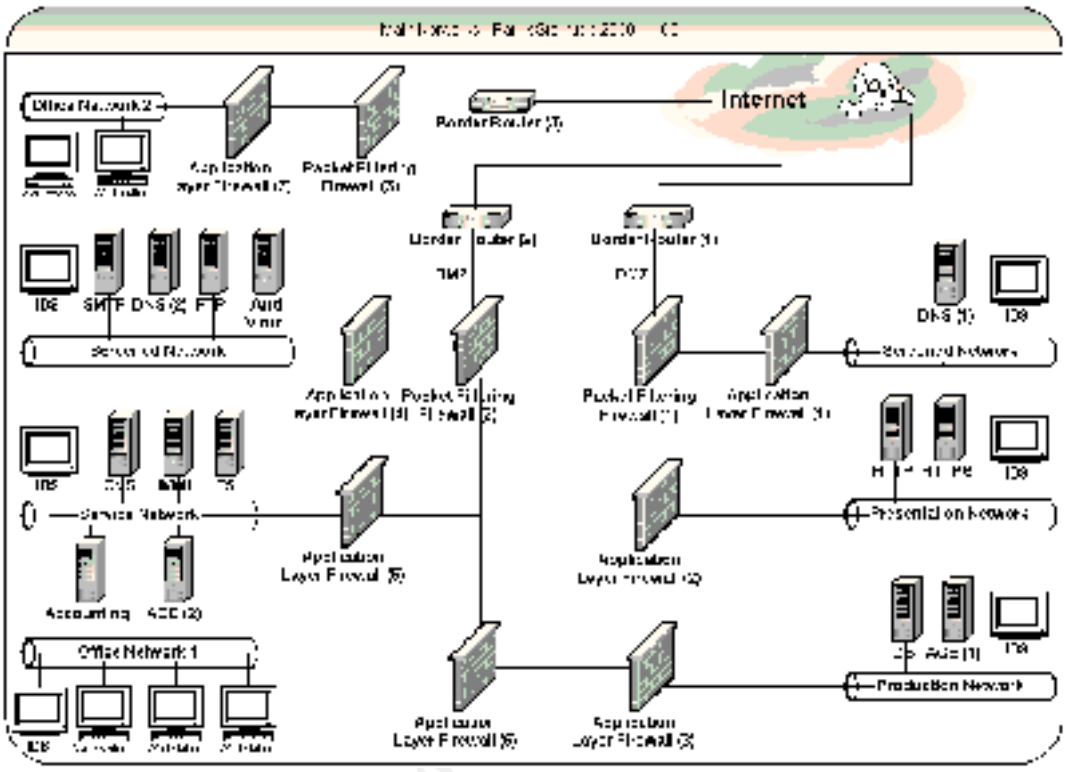

3.5.2 Description

I will begin with describing the process of a customer trying to access fortune cookies (hereafter named "cookies"). This traffic will initially pass the border router, which hopefully will not filter it out. If it does, the customer probably is not up to any good (or he should produce a good explanation why he has this sudden urge to get cookies over ICMP or NetBIOS). At this point, the protocols are either HTTP, HTTPS or DNS. If DNS is the case, for the typical customer it should be a DNS query within the UDP protocol. TCP connections (hey, if you want an answer for a hostname that is longer than 512 bytes, we do not want you here in any case) will only be accepted from the secondary DNS located at a trusted ISP. It is located there for redundancy considerations. To get permission to act as secondary DNS, the ISP had to guarantee it would take proper means to ensure the security of the underlying OS, as well as make sure it would not allow Zone transfers to or fro m any IP other than the primary DNS. Also, to limit the danger of DNS cache poisoning, none of the DNSes should allow recursion (this is set in the named.conf file, assuming you are running BIND version 8 from ISC). The only other traffic permitted through the packet filtering firewall is HTTP and HTTPS, provided they come on standard ports (80 and 443, respectively).

After the packet filtering firewall, the traffic will have to pass through additional application layer firewalls. These separate the inter nal networks from each other. The packet filtering firewall acts as a noise killer since it is faster than the application layer firewalls. Below is a schema describing how the networks may communicate with each other.

Source	Destination	Service
Internet	DNS Server	DNS
Internet	Web Servers	HTTP/HTTPS
Web Servers	Data Base Server	DB (Oracle)
Web Servers	ACE Server	RADIUS

The HTTP servers may talk to the backend databases (containing all cookies) but before allowing that connection to be made, the cu stomer must successfully authenticate via a SecurID token (RADIUS is used to communicate with the ACE Server – RADIUS is a standard, SecurID is proprietary so the application was developed for RADIUS).

Data from established connections and UDP replies ar e allowed to pass in the reverse direction. All other traffic is rejected. As you can see, no connections may be initiated from the inside going out to the Internet, but data containing cookies may pass back to the customer via the TCP sessions already ini tiated.

To the left on the diagram, we have the corporate network where applications and cookies are developed to be later included in the services to the customers. This document do not focus on the good practices of doing such changes in testing environ ments, it is considered outside the scope.

The challenge here is to let the employees access the Internet to a certain degree, allow incoming emails, etc, and at the same time maintain security. To add more complexity, a second division of the company is located in another country (this used to be a smaller company before it got bought up and incorporated). Therefore, the use of a VPN is required. Also, the administrative staff must be able to pull out reports from the databases at need so a link to the production network is set up as well.
And of course, the company, being as big as it is, has several partners and providers out there in the dark cloud of the Internet. Those must be able to upload and download information at need, but only the information relevant to them. They must not access other confidential data and those transfers they do must be protected against eavesdropping to ensure integrity of both data and account information.

So how is this to be solved then?

Below are listed all service s allowed between the networks:

Source	Destination	Service
Internet	SMTP Server	SMTP
Internet	Workstations	ICMP (echo reply, unreachable)
Anti-virus Server	Internet	SMTP
Internal DNS Server	Internet	DNS Queries
Partners/Providers	FTP Server	SSH/VPN
Anti-virus Server	Mail Server	SMTP
Mail Server	SMTP Server	SMTP
Workstations	Internet	ICMP (echo request)
Workstations	Internet	HTTP/HTTPS/FTP
Workstation	Service Network	Mail/CVS/FTP/DNS

Most of this is pretty straightforward, but I think I should c larify the issue with incoming and outgoing email messages. There are three servers capable of handling SMTP; the primary MX

for GIAC Enterprises, the secondary MX, which is also acting as anti-virus host for incoming SMTP data, and finally the internal Mail Server where the users' mailboxes are stored. The firewall will allow incoming SMTP to the primary MX, but it will not allow it to send any messages out. This host will check that it, in fact, should relay the message. If so, it will forward it to the anti-virus server which first scans the message for viruses or other non-wanted contents. If the message pass the test, it will either send it to the Mail Server (in case the message is for an internal user) or out to the Internet (in case it is addressed to an external user). The secondary MX will not receive SMTP from any other host than the primary MX. The anti-virus server is secondary MX only because some mail servers on the Internet require the sending server to have a MX record. The virus definitions on the anti-virus server will be updated automatically on daily basis.

Partners and providers who want to use the FTP server will have to log on using a SecurID token. This will protect against stolen passwords and make tracking easier.

3.6 Architecture: Fourth Area

3.6.1 Network Diagram

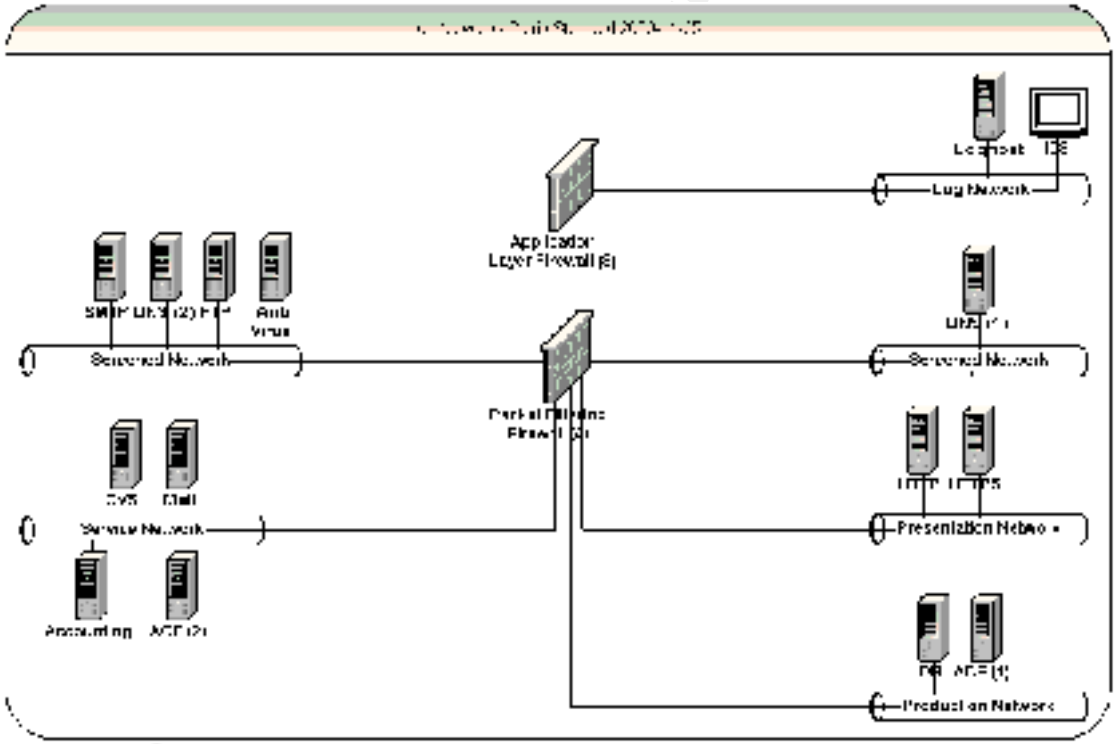

3.6.2 Description

As mentioned previously, all servers have two NICs, the first one being used for standard services, and the other, for sending system logs to the central log host. The hosts would be on different physical network s, using the same structure as the primary NIC. The networks will not be on a broadcast media, such as a hub, but rather a switch (it should probably be mentioned that a switch is not secure against sniffing attacks, but it is harder to accomplish and much more likely to be discovered). Those switches are connected to different legs to the packet filtering firewall, which is in place to ensure only valid traffic is let through. Since all syslog traffic is considered important, the standard syslog daemon is replaced with nsyslog by Darran Reed. It will give the network slightly more overhead since it uses TCP instead of UDP, but I do not consider this a big issue since we already are on a network dedicated for logging. The pro is of course the simple fact tha t TCP is not connectionless but rather makes sure the data is delivered. Another software will be deployed to make md5 checksums of incoming log events (this feature is of yet not incorporated in nsyslog). This is done to ensure higher reliability of data , as well as potential requirements as evidence in court.

After passing through the packet filtering firewall, the data will also pass through the application layer firewall where incoming data is inspected to make sure nothing else tries to trick the packet filter by using an accepted port in the firewall.

3.7 Architecture: Fifth Area

3.7.1 Network Diagram

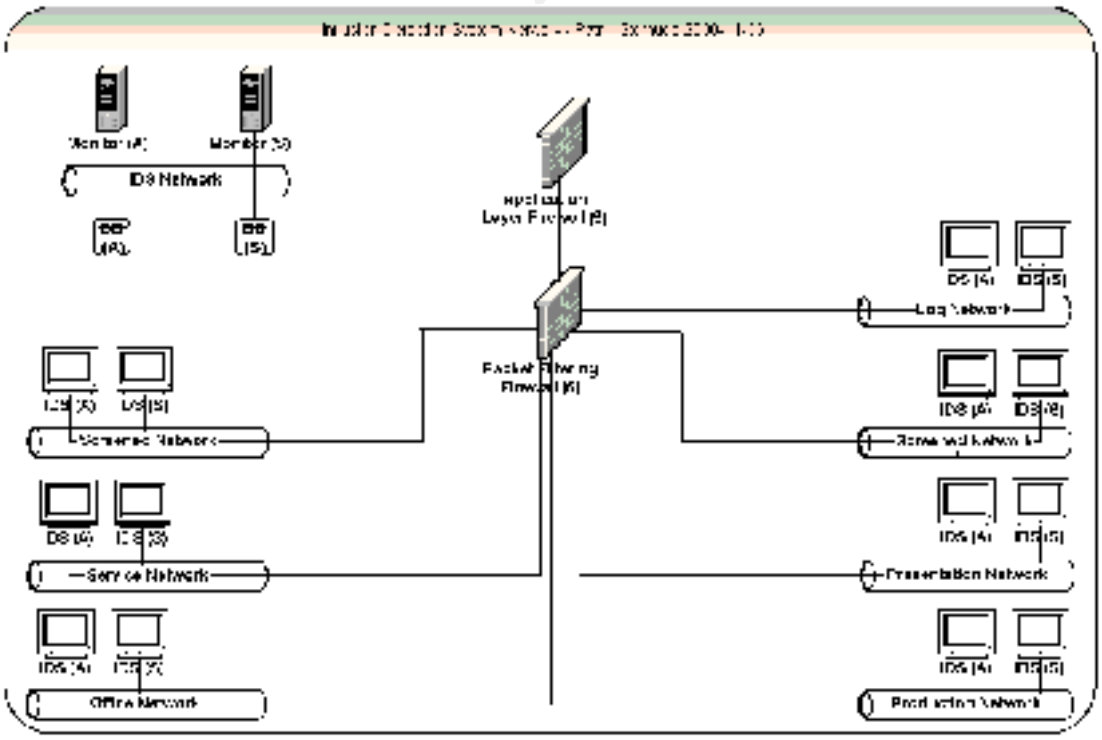

3.7.2 Introduction to IDS (if you are familiar with IDS, you may want to skip this)

NIDS sensors are placed on all networks. NIDS means "Network Intrusion Det ection System" which differ from Host Based Intrusion Detection in that the data is collected from the wire (today, I would guess the most common media is ethernet) instead of processes or audit modules (like Solaris BSM which has been shipped for some tim e now). The advantage is the ease with which you can discover attack patterns on different hosts at the same time. It is also a good early warning system that something is amiss. There are several vendors on the field of NIDS, Cisco probably being one of t he biggest with their Cisco Secure IDS, previously named NetRanger. There are open source alternatives as well, tcpdump based Snort and Shadow being among them. Intrusion Detection differs in another thing except whether it is network or host based. Intrus ions can either be detected using signatures of known attacks (drawback: you will probably not notice new, unknown attacks) or by anomaly detection. That is, everything that does not fit in is suspicious and triggers an alarm (the drawback here is the likelihood of false positives. Also, if the system is "learning" what the normal behaviour is, the dedicated attacker might slowly change his or hers habits, leaving the IDS to believe everything is normal). There are also various ways the intruder might attem pt to evade the IDS or even cause it to fail. But why then, should we use NIDS? The reason is, should an intruder target your network it is more than likely some (not all, but some) of his/her activities will get detected, causing an alarm to trigger.

3.7.3 Architecture Description

There will be two NIDS sensors on every production network at the main geographical site, the Log Network included. There are no sensors on the DMZ since this would yield too many alarms on various portscans or simply noise (misconf igured devices as routers, etc). There will be one signature based and one anomaly based system working together. What the first one might miss will most likely be noticed by the second. Special care must be taken to ensure all fragmented traffic is to cau se alarms since it most likely will be an attack, giving the fast connection to the Internet. Yes, fragments may arrive normally, but it would be worth to investigate nonetheless. The NIDS may not defragment correctly, but as long as it report that fragmented traffic was noticed, the security staff could investigate by hand.
All sensors will have the first NIC (Network Interface Card) in promiscuous (listening) mode without MAC- or IP-address. Additionally, it would not be able to transmit data, thus rende r it invisible on the network. This can be achieved in various ways, some commercial sensors has this by default. The other NIC will be attached to the IDS network which will be used for sending alerts or reconfigure the sensors. Each network with producti on hosts will have a corresponding network on which the IDS management interface will be connected. All of these networks are connected to a separate interface in a packet filtering firewall, with a rule base designed for only letting appropriate traffic f low (logs to the monitors, and configurations and updates to the sensors). This firewall will then pass through the data to the next firewall (application layer) before arriving to the monitors. The monitors will save all incoming data to their databases for future analysis or as evidence. From the monitors to the sensors, the traffic will pass in reverse order. As previously mentioned, the IDS network will have no connectivity to any other network, any upgrades in software or signature files will have to be applied via CD-ROM (after scanned for viruses).

3.8 Architecture: Host Based Protection

3.8.1 Servers

Except the good practice of keeping servers protected by firewalls, it is likewise important to keep the hosts secure by themselves. A good philosophy is to consider a host secure enough only if you can trust it to be safe were you to put it directly on the Internet without the protection of firewalls or filtering routers. A few recommendations in this topic are:

1. Install Tripwire on all servers. Keep the da tabases on read-only media such as CD-ROM. Compare files on the system with those in database at least twice a day.
2. Deploy host-based firewalls. For Unix servers, the choice would either be IPFilter by Darran Reed or SunScreen EFS Lite by Sun Microsystems (bundled with Solaris 8).
3. Log all syslog entries to a remote loghost (previously mentioned).
4. Allow no remote administration to important servers.
5. Use of secure passwords for all accounts.

3.8.2 Workstations

It is also important to have good protection on works tations. The most important thing coming to mind is a good anti-virus solution. Viruses may cause much damage and the company could be hurt from media attention concerning virus infections. Anti-virus software also tends to find many backdoors and trojans. The virus definitions should be scheduled to update every night (automatically check for new definitions and if they exist, download and install them).

4 Second Assignment: Security Policy

"For the purposes of this assignment, your security policy should be focused on implementation of requirement number 1 above "Install and maintain a working network firewall to protect data accessible via the Internet." For a baseline policy, use the filtering recommendations located at www.sans.org/topten.htm. You DO NOT need to repeat that information. Instead, focus on ADDITIONAL filtering you would recommend and why. Keep in mind you are an E-Business with customers, suppliers, and partners, you MAY NOT simply block everything! Your policy should implement your design above."

4.1 My Choice of Firewall And the Layout of This Assignment

I have chosen to use Firewall-1 & VPN-1 from Check Point Technologies for this assignment. There were several reasons for this. First, Firewall-1 is, if not the most, one of the most used firewalls on the market. So chances are most people reading this document will already be familiar with it. Second, it runs on most platforms so it was possible to get a test system set up within a month. Third, it has a GUI that is easy to understand and configure. And finally, since this is the firewall I am most used to, the choice was pretty easy. For the moment, the latest release is 4.1 with service pack 2, which is what I will use.

The underlying operative system in my examples are either Sun Solaris or Red Hat Linux. As I have already stated, I do not feel comfortable with Microsoft products and as far as my knowledge goes, Firewall-1 was originally developed on Sun Solaris. Additionally, I find it easier to load lots of NICs into an Enterprise Ultra server than any PC-based hardware (and yes, often many times more expensive, but it scales so much easier and run so much more stable).

For those of you who do not have experience with it, I have taken screen dumps from the GUI. I have not dumped every menu or configuration option; this would take too much time, and besides, Check Point already did this with their CCSA and CCSE course material. The screen dumps are all from the first firewall (labelled "Packet Filtering Firewall (1)" on the main diagram). The rulebases of the Application Layer Firewalls will not be covered at all since the time for the assignment is limited and the most interested events probably will occur at the packet filtering firewalls, them being closest to the Internet and all.

4.2 Firewall-1 Considerations

If you feel at ease with Check Point Firewall -1, you might want to skip this section completely and go directly to the next chapter (4.3).
This chapter contains information about some behaviours of Firewall -1 (for example, how the rule ordering works) as well as how the GUI looks.

4.2.1 Creating Workstation And Network Objects

To create rules, the administrator first need to set up workstation and network objects. These are used within the rules to grant or deny access to other networks or workstations. Please note that a workstation is any type of host (server, firewall, workstation or PC).

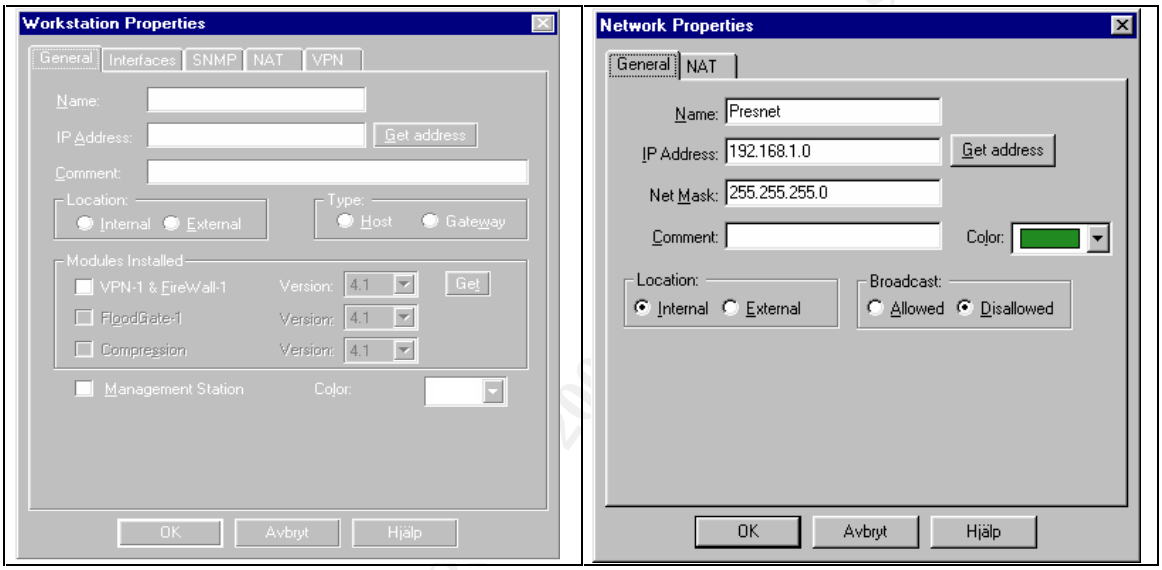

Above, to the left, a workstation object is defined. If it was a firewall, the "VPN -1 & FireWall-1" checkbox would be filled in the "Modules Installed" section. To the right, a network object is defined. It is possible to do automatic address translation from the objects, but it gives you fewer options as how to set it up. It also adds to the locations to check for problems when something is amiss. I prefer to keep all NAT in one place.

4.2.2 The Policy Properties

Many important configurations take place in the Policy Properties. The administrator must make sure only sensible values are set or the security of the firewall and network behind could be at stake.

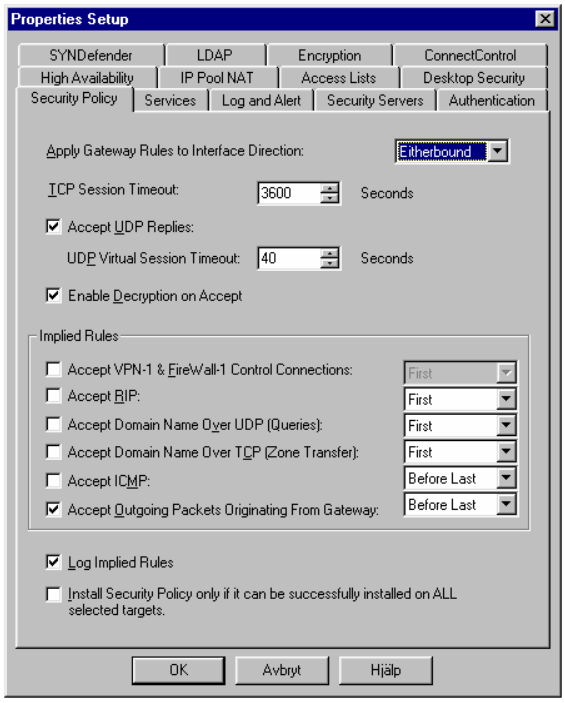

By default, many of these options are turned on. Worse, they are not even logged. This I consider a weakness in Firewall -1. A firewall, by default, should deny everything. Some things should be turned on, however. If the "Accept Ou tgoing Packets Originating From Gateway" is unchecked, the rulebase must take into account no packet will be allowed to leave any interface of the firewall unless explicitly permitted to do so. Likewise, without "Accept UDP Replies", returning UDP packets must be allowed regardless if they are replies or not, giving intruders additional possibility to get traffic through the firewall.

It is also important to remember the "Services" menu. If you know you will not need RPC or FTP, be sure to turn them off.

4.2.3 Rule Ordering

The rulebase of Firewall -1 is order dependent. The inspection engine will, for each packet, compare the contents with the policy, starting at the top and moving down. If a rule matches, the specified action is taken and the process starts all over again with the next packet waiting. Because of this, it is extremely import to consider the implications of placing a rule in the wrong place. In the worst case, a packet you intended to drop could instead be accepted. Rulebase ordering also effec t performance. Rules triggering on a lot of traffic should be placed as high up as possible.

Below is a screen dump from the Policy Editor:

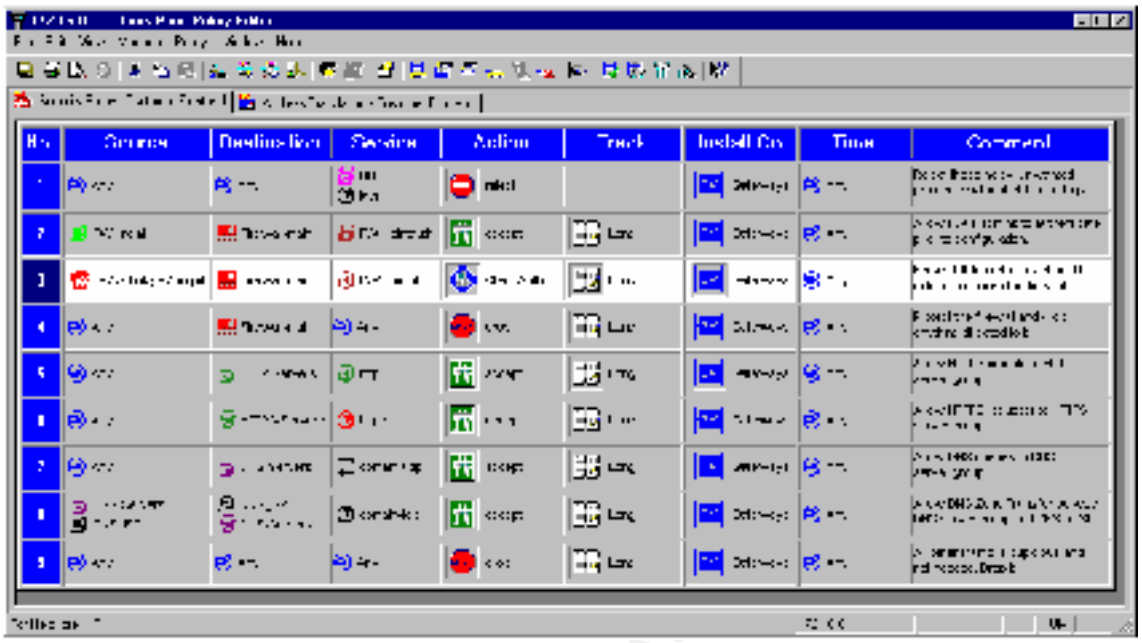

All settings in the Policy –> Properties (I will use this type of writing to indicate a sub -menu throughout this document) menu are considered as Rule 0 and therefore processed first.

4.2.4 Actions on Matched Rules

Firewall-1 has eight different actions to be applied on a packet matching a rule. Of these, "Accept" (letting the packet pass) and "Drop" (discard the packet without notification) are the most common ones in a rulebase. However, sometime you will want to use "Reject" which will send a TCP RESET back to the host attempting the connection. "Client Auth" is used when the initiating host must log on and authenticate before being allowed to pass (this is based on IP number). "Encrypt" is used when deploying VPNs. "Client Encrypt" is used together with SecuRemote (another Check Point product). "User Auth" requires user authentication every time a session is initiated (works only with standard services as telnet, smtp and ftp). Finally, "Session Auth" resembles "Client Auth" but requires separate software while "Client Auth" can be done either with telnet or HTTP.

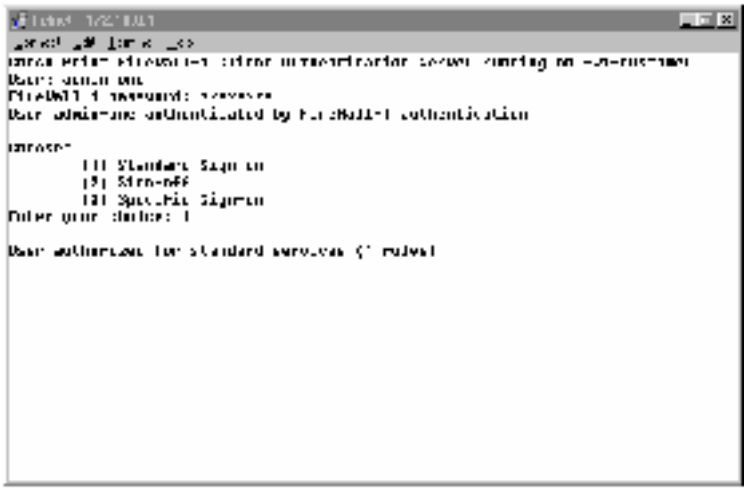

Above is an example of a Client Auth connection . A telnet to port 259 (default port) was made to the firewall and the administrator successfully gained permission to start the Policy Editor at the management server!

4.2.5 Firewall-1 Logging

When it comes to logging, Check Point gives you the following opti ons: "Short", "Long", "Account", "Alert", "Mail", "SnmpTrap", and "User Defined". By not applying any of these options, you chose not to log at all. "Short" and "Long" only differs on how much data will be logged. "Long" is the option I use most because if I wish to log something I want all data available to be able to understand better what I am seeing. Alert pops up an alert box on the GUI and Account records the data in a format made for keeping accounting records. "Mail", "SnmpTrap" and "User Defined" c an all be configured in the "Policy -> Properties -> Log and Alert" screen (Lance Spitzner uses the "User Defined" to turn the logging facility of Firewall -1 into an Intrusion Detection System, more information about this can be found at "http://enteract.c om/~lspitz").

Below is an excerpt from the GUI's log viewer (the colouring and icons makes the reading easier when you are monitoring in real time):

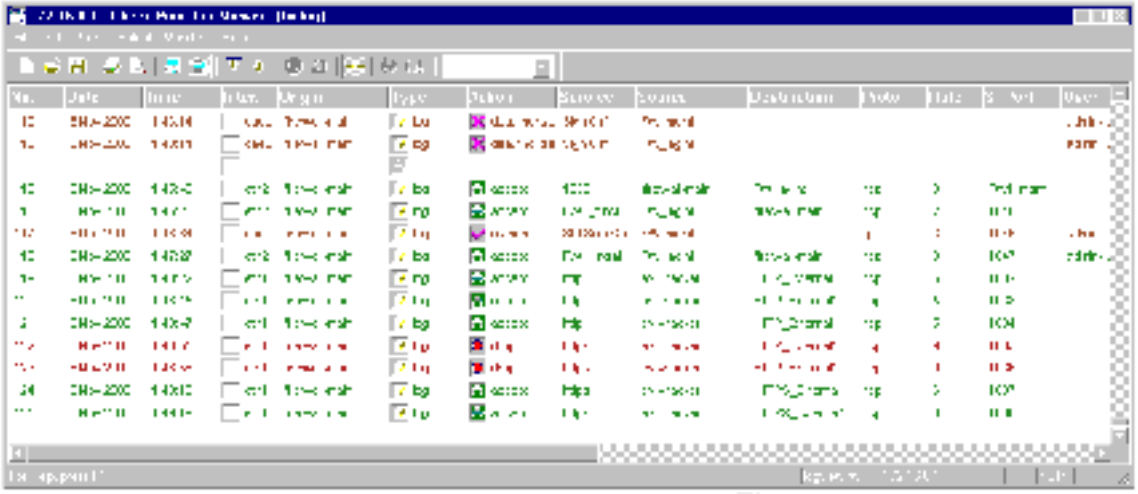

As you can see, there are a lot of useful data here (for example, the Interface information is of great worth when you are trying to figure out why your packet is dropped when you actually did permit it in your rulebase).

If you would rather read the information in standard ASCII, use the "fw log –ft" on the management server (-ft means it will begin from last position in the log file and display new, incoming events as they occur).

You can also chose to view only open (active) connections:

4.2.6 Network Address Translation

NAT in Firewall-1 (Network Address Translation) require two things to be done in the OS to ensure proper functionality (in fact, any functionality at all). First, you need to add a static route to the "virtual" address, using the real IP as gateway.

Adding routes in Sun Solaris and Red Hat Linux
[root@solaris_box]# route add [tra nslated IP] [real IP] [root@linux_box]# route add [translated IP] gw [real IP]

Second, you need a static ARP for the translated IP. The ARP -address must be set to the MAC of the external interface of the firewall. This technique is called "Proxy ARP" (if you

want to avoid trouble, and need NAT, you probably should not use Windows NT since it seems to have problems with Proxy ARP).

Creating Proxy ARPs in Sun Solaris or Red Hat Linux (don't forget the 'pub')
[root@giac_fw]# arp –s [translated IP] [MAC of external interface] pub

Below is a dump from Check Point's Address Translation window.

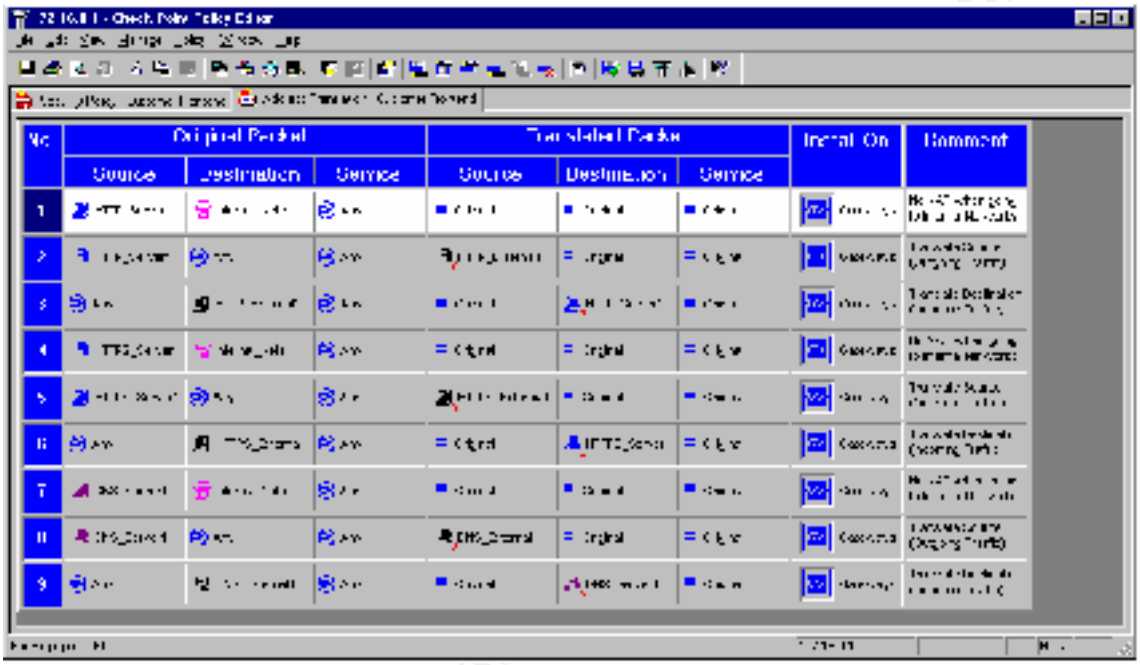

4.2.7 Spoofing Protection

Spoofing is bad. We do not want spoofed packets to traverse our networks. First, it means an intruder could possibly bypass our security devi ces by feigning source IP. Second, it could be used for Denial of Service attacks against other sites on the Internet. It is of equal importance to protect against spoofing both inbound and outbound. This is easily achieved in Firewall -1.

First you open your workstation object for the firewall and change to the "Interfaces" screen. All four Interfaces are listed (if they do not show up, click on "Get").

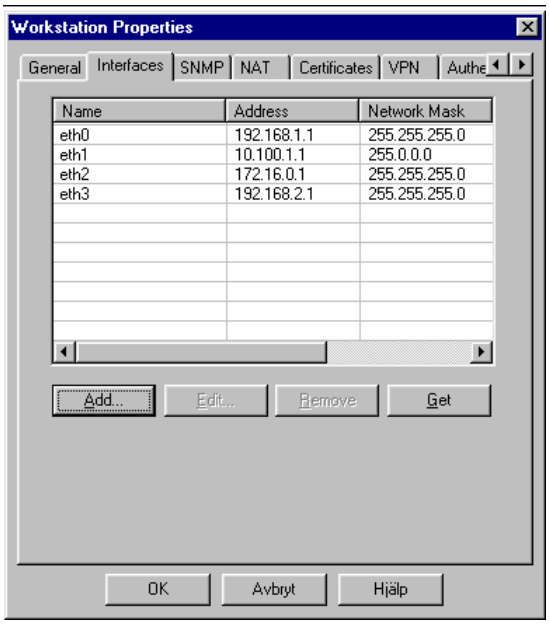

Then you need to modify the Security Property on all interfaces. Be sure to remember to change tracking from "None" to "Log", otherwise you will have a bad time trying to figure out why packets are rejected (and it is very good to know if someone tries a spoofing attack on you). Shown below is the Properties of the Presentation Network:

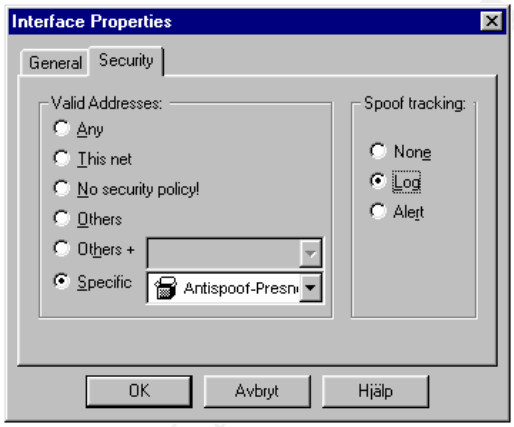

"Specific" is chosen as "Valid Addresses", meaning the administrator assigns a object already defined. This option must be used when you have Network Address Translation in place. You first create a group where you will add the network object for the particular netw ork (or for the truly paranoid, all separate hosts, making it harder for someone to add new hosts on that net) on that interface, together with the object of the translated address. If you do not, all packets addressed to the virtual address will be rejec ted. If you do not have NAT, you can make it easier for yourself and specify "This net" as valid addresses. Below can be found the Group Properties of the group previously used for antispoofing.

A similar group must be created for the Screened Network.

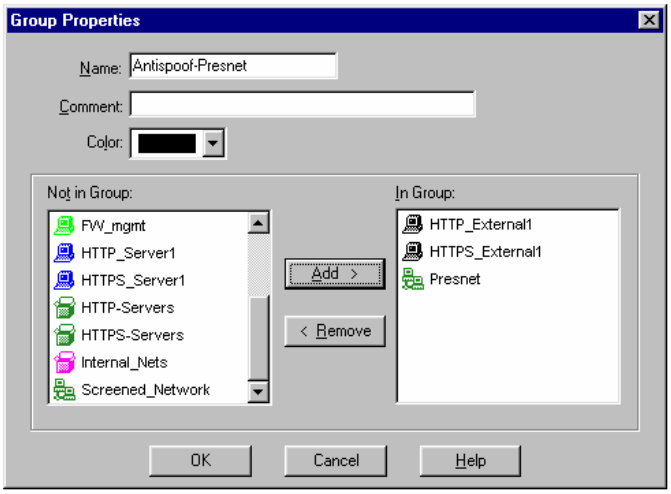

But how should we define all addresses on the Internet? This is for the configuration of the external interface. Luckily, Check Point provide us with the "Other" option which will match every address not defined on any of the other interfaces. Again, rem ember to check the "Log" radio button in the tracking properties.

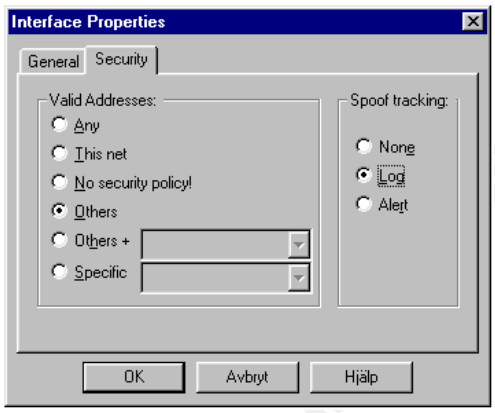

4.2.8 The SYNDefender

Another thing we do not want into our networks is SYN flooding attacks. It makes our servers feel bad so we should do something about it. Firewall -1 comes with something named SYN Defender. Below the properties can be seen.

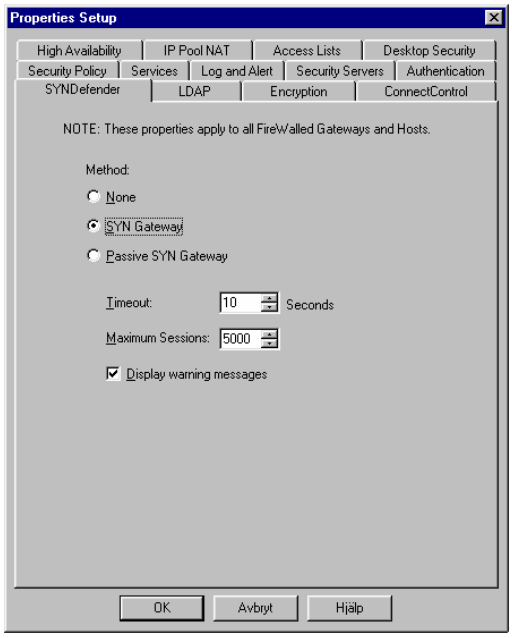

It is up to the administrator to decide which (if any) option is to be used. The SYN Gateway intercepts incoming SYN packets, sends back SYN -ACK and wait for the final ACK. If there is none, Firewall-1 drops the whole connection and the server behind the firewall will never notice the exchange. This could lead to false positives for potential intruders who use SYN scans to map the network because they will get an answer no matter if the host on the other side is up or not.

The Passive SYN Gateway listens for incoming SYN packets. When the server replies with the SYN-ACK, the timer is started. If no ACK returns within the defined interval (default is 10 seconds), the firewall sends a RST to the s erver, spoofing the source as the initiator of the connection. Whichever solution is deployed, the "Display warning messages" should be checked.

4.3 The First Firewall: Protecting the Production Environment

No	Source	Destination	Service	Action	Track
1	Any	Any	NBT ident	reject	
	Explanation				
	This rule is processed first; it is an effective noise killer for all those stray NetBIOS types coming in. It also takes care of ident requests. Note that the action is reject rather than drop, meaning mail services asking for ident not will have to wait for timeout. Also, since this is considered noise I don't want it to clutter the logs so the Track type is omitted. If this was not first in the base, we could still get a lot of log entries from rule number 4 (the firewa ll lockdown rule, which drops and logs all matches).				

No	Source	Destination	Service	Action	Track
2	FW_mgmt	Firewall-main	FW1_clntauth	accept	Long
	Explanation				
	The firewall needs to be administrated over the network since the firewall host itself has no X-windows and it takes too much effort to compile rule bases without the GUI. But even though we only allow the MGMT box to administrate (see next rule), we want to make sure the user requesting the GUI console indeed has permission to do this. To be ab le to do this, we need to allow the MGMT to access the client authentication services.				

No	Source	Destination	Service	Action	Track
3	Fw-Adm@FW_mgmt	Firewall-main	FW1_mgmt	Client Auth	Long
	Explanation				
	Here we are; members of group Fw -Adm can connect to the management server if successfully authenticated. Had we placed this rule before the last one (number 2), we would not have been able to administrate the firewall at all since the firewall would not let us reach the authentication service				

No	Source	Destination	Service	Action	Track
4	Any	Firewall-main	Any	drop	Long
	Explanation				
	So now all that should be able to connect to the firewall are allowed to do so, this is defined in the two pervious rules. Any other traffic going to the firewall will be dro pped.				

No	Source	Destination	Service	Action	Track
5	Any	HTTP-Servers	http	accept	Long
	Explanation				
	This is where we start with the actual services; HTTP traffic should be let through to all members of the HTTP-Servers group (all physical HTTP Servers on the Presentation Network together with their translated IP addresses). This rule is first in the "Service part" of the rulebase because I expect the largest portion of the incoming traffic to trigger on this one, no need to put it further down in the ba se.				

No	Source	Destination	Service	Action	Track
6	Any	HTTPS-Servers	https	accept	Long
	Explanation				
	Since all customer data (either customer logon data or fortune cookies) must be transmitted securely, we need to give them access to port 443 (HTTP over SSL) on all HTTPS-Servers (and their corresponding translated IP-addresses).				

No	Source	Destination	Service	Action	Track
7	Any	DNS-Servers	domain-udp	accept	Long
	Explanation				
	This rule may well trigger more often than both HTTP and HTTPS but the amou nt of data will be far less. It may also be less traffic due to the caching behaviour of the DNS servers. In any case, what this rule does is to give everyone the possibility to send DNS queries (over UDP) to the DNS Servers group.				

No	Source	Destination	Service	Action	Track
8	DNS-Servers DNS_ISP	DNS_ISP DNS-Servers	domain-tcp	accept	Long
	Explanation				
	For the functionality of DNS, the primary and secondary server must be able to do zone transfers. This rule does just that; note that Firewall -1 supports multiple hosts or groups of hosts in the Source, Destination and Service fields (as you know, DNS zone transfers takes place over TCP, so domain -tcp signifies port 53 of the TCP protocol).				

No	Source	Destination	Service	Action	Track
9	Any	Any	Any	drop	Long
	Explanation				
	Cleanup time! Everything we did not previously allow, we want dropped and logged. Here is for example where stray telnet sessions will go (together with a lot of other peculiarities that often may be seen on the Internet), as well as intrusion attempts and port scans.				

4.4 The Second Firewall: Protecting the Corporate Network

The first rules are the same as the in the first firewall but I include them anyway so that they are not forgotten. Notice that if the company where to use Micros oft servers for file sharing/printing and so forth, we would have to allow NetBIOS from the Office Network located in the other country. Since this particular company in fact uses Unix servers (and no, RPC is not used between the firewalls), this is not an issue and all NetBIOS traffic can be safely discarded without further notice.

No	Source	Destination	Service	Action	Track
1	Any	Any	NBT ident	reject	
	Explanation				
	This rule is processed first; it is an effective noise killer for all those stray NetBIOS types coming in. It also takes care of ident requests. Note that the action is reject rather than drop, meaning mail services asking for ident not will have to wait for timeout. Also, since this is considered noise I don't want it to clutter the logs so t he Track type is omitted. If this was not first in the base, we could still get a lot of log entries from rule number 4 (the firewall lockdown rule, which drops and logs all matches).				

No	Source	Destination	Service	Action	Track
2	FW_mgmt	Firewall-main	FW1_clntauth	accept	Long
	Explanation				
	The firewall needs to be administrated over the network since the firewall host itself has no X -windows and it takes too much effort to compile rule bases without the GUI. But even though we only allow the MGMT box t o administrate (see next rule), we want to make sure the user requesting the GUI console indeed has permission to do this. To be able to do this, we need to allow the MGMT to access the client authentication services.				

No	Source	Destination	Service	Action	Track
3	Fw-Adm@FW_mgmt	Firewall-main	FW1_mgmt	Client Auth	Long
	Explanation				
	Here we are; members of group Fw -Adm can connect to the management server if successfully authenticated. Had we placed this rule before the last one (number 2), we would not ha ve been able to administrate the firewall at all since the firewall would not let us reach the authentication service				

No	Source	Destination	Service	Action	Track
4	Firewall-main FW_Office2	FW_Office2 Firewall-main	FW1 IPSEC	accept	Long
	Explanation				
	As we are going to deploy a VPN to the other site, the firewalls need to exchange keys (ISAKMP) and allow incoming ESP (Encapsulated Security Payload) as well as AH (Authentication Header). *If this rule were placed after the next rule, the VPN could not work ; the firewall would even drop the SA negotiation.*				

No	Source	Destination	Service	Action	Track
5	Firewall-main FW_Provider1	FW_Provider1 Firewall-main	FW1 IPSEC	accept	Long
	Explanation				
	Same as above, but for a provider instead of another division wit hin the company.				

No	Source	Destination	Service	Action	Track
6	Any	Firewall-main	Any	drop	Long
	Explanation				
	So now all that should be able to connect to the firewall are allowed to do so, this is defined in the two pervious rules. Any other traffic goi ng to the firewall will be dropped.				

No	Source	Destination	Service	Action	Track
7	Office_Net	! Internal_Nets	echo_request	accept	Long
	Explanation				
	Allow employees to ping hosts on the Internet				

No	Source	Destination	Service	Action	Track
8	! Internal_ Nets	Office_Net	echo-reply unreach mtu_discovery	accept	Long
	Explanation				
	Allow replies from pings together with ICMP unreachable messages and MTU discovery.				

No	Source	Destination	Service	Action	Track
9	! Internal_Networks Internal_Mail_Server	SMTP-Server	smtp	accept	Long
	Explanation				
	We want to receive email from the Internet; however, we do not want our internal users or hosts to connect to our main SMTP Server. The exclamation mark indicates that the Internal_Networks is negated, that is, everyon e not coming from our internal nets (that would be the Internet) will be allowed. There is one exception to this, though. We want our Internal Mail Server to send outgoing mails through our SMTP Server. If the firewall administrator has objections against using negated hosts or networks, this could be done as three rules instead (this is why I feel negate is good, it just saved me two addition rules). The first rule would accept traffic from the Internal Mail Server, the second would deny the Internal_Netw orks, and the third would give everyone (Any) access to it.				

No	Source	Destination	Service	Action	Track
10	AntiVirus_Server	! Internal_Networks Internal_Mail_Server	smtp	accept	Long
	Explanation				
	Again I use negate to allow the outbound smtp server (th e Anti-virus Server, why this is was covered in the previous assignment) to send emails out into the world. Of course, I do not want the SMTP Server to start talking to internal hosts (this is where the "negate" comes into play). Still, there is one except ion: the Anti-virus Server needs to talk with the Internal Mail server to deliver incoming messages (this is the second row in the Destination field)				

No	Source	Destination	Service	Action	Track
11	DNS_Server	! Internal_Nets	domain_udp	accept	Long
	Explanation				
	Our Internal DNS Server must be able to ask other servers on the Internet, but only for UDP queries. It will not be accessible from the Internet, since it contains no external zones. It should of course not talk with Internal hosts either, so ne gate is used.				

No	Source	Destination	Service	Action	Track
12	Provider1	FTP_Server	FTP	Encrypt	Long
	Explanation				
	Provider1 may FTP to the FTP Server but only if the traffic comes encrypted. This is the second part of the VPN configuration. Unencrypted tr affic will not match this rule; nor will traffic that cannot be successfully decrypted (it will be matched, but not accepted).				

No	Source	Destination	Service	Action	Track
13	Provider_Hosts Partners_Hosts	FTP_Server	SSH	accept	Long
	Explanation				
	Some prov iders and partners may not be able to use VNPs themselves, but we still want them to do their work. Furthermore, we want it encrypted! *We solve this by allowing SSH connections to the FTP Server (which will be configured to allow sftp, but they will not b e able to log on into the box with a normal SSH session).*				

No	Source	Destination	Service	Action	Track
14	FTP_Server	ACE_Server	RADIUS	accept	Long
	Explanation				
	When someone attempts to log in to the FTP Server, he or she must authenticate themselves via SecurID. The FTP Server will issue a RADIUS query to the ACE Server. Needless to say, if the query never arrives, the ACE Server will have problems to respond so this rule is quite necessary.				

No	Source	Destination	Service	Action	Track
15	Office_Network2	Internal_DNS	domain_udp	accept	Long
	Explanation				
	The subdivision's hosts should be able to query the Internal DNS, but of course, no Zone Transfers.				

No	Sources	Destination	Service	Action	Track
16	Office_Network	Internal_DNS	domain_udp	accept	Long
	Explanation				
	Internal hosts should also be able to get DNS functionality (and again, no Zone Transfers).				

No	Source	Destination	Service	Action	Track
17	Office_Network	! Internal_Nets	http https ftp	accept	Long
	Explanation				
	This is what will be allowing out to the Internet. Note the negated Internal_Nets. We don't want our workstations to reach anything else (except what is defined above).				

No	Source	Destination	Service	Action	Track
18	Office_Network2	CVS_Server	CVS	Encrypt	Long
	Explanation				
	The remote Office network must be allowed to communicate with the CVS server, provided that the traffic is encrypted.				

No	Source	Destination	Service	Action	Track
19	Office_Network2	FTP_Server	FTP	Encrypt	Long
	Explanation				
	The remote Office network must be allowed to communicate with the FTP server, provided that the traffic is encrypted.				

No	Source	Destination	Service	Action	Track
20	Office_Network2	Mail_Server	MUA	Encrypt	Long
	Explanation				
	The remote Office network must be allowed to communicate with the internal Mail Server on the port the MUA (Mail User Agent, for example Lotus Notes) use.				

No	Source	Destination	Service	Action	Track
21	Any	Any	Any	drop	Long
	Explanation				
	Cleanup time! Everything we did not previously allow, we want dropped and logged. Here is for example where stray telnet sessions will go (together with a lot of other peculiarities that often may be seen on the Internet), as well as intrusion attempts and port scans.				

4.5 Third Firewall: Protecting the division at another location

No	Source	Destination	Service	Action	Track
1	Any	Any	NBT ident	reject	
	Explanation				
	This rule is processed first; it is an effective noise killer for all those stray NetBIOS types coming in. It also takes care of ident requests. Note that the action is reject rather than drop, meaning mail services asking for ident not will have to wait for timeout. Also, since this is considered noise I don't want it to clutter the logs so the Track type is omitted. If this was not first in the base, we could still get a lot of log entries from rule number 4 (the firewall lockdown rule, which drops and logs all matches).				

No	Source	Destination	Service	Action	Track
2	FW_mgmt	Firewall-local	FW1_clntauth	accept	Long
	Explanation				
	The firewall needs to be administrated over the network since the firewall host itself has no X-windows and it takes too much effort to compile rule bases without the GUI. But even though we only allow the MGMT box to administrate (see next rule), we want to make sure the user requesting the GUI console indeed has permission to do this. To be able to do this, we need to allow the MGMT to access the client authentication services.				

No	Source	Destination	Service	Action	Track
3	Fw-Adm@FW_mgmt	Firewall-local	FW1_mgmt	Client Auth	Long
	Explanation				
	Here we are; members of group Fw-Adm can connect to the management server if successfully authenticated. Had we placed this rule before the last one (number 2), we would not have been able to administrate the firewall at all since the firewall would not let us reach the authentication service.				

No	Source	Destination	Service	Action	Track
4	Firewall-main Firewall-local	Firewall-main Firewall-local	FW1 IPSEC	accept	Long
	Explanation				
	As we are going to deploy a VPN to the main site, the firewalls need to exchange keys (ISAKMP) and allow incoming ESP (Encapsulated Security Payload) as well as AH (Authentication Header). If this rule was placed after the next rule, the VPN would not work; the firewall would even drop the SA negotiation.				

No	Source	Destination	Service	Action	Track
5	Any	Firewall-main	Any	drop	Long
	Explanation				
	So now all that should be able to connect to the firewall are allowed to do so, this is defined in the two pervious rules. Any other traffic going to the firewall will be dropped.				

No	Source	Destination	Service	Action	Track
6.	Local_Net	! Internal_Nets	echo_request	accept	Long
	Explanation				
	Allow employees to ping hosts on the Internet				

No	Source	Destination	Service	Action	Track
7.	! Internal_Nets	Local_Net	echo-reply unreach mtu_discovery	accept	Long
	Explanation				
	Allow replies from pings together with ICMP unreachable messages and MTU discovery.				

No	Source	Destination	Service	Action	Track
8.	Local_Net	! Internal_Nets	http https ftp	accept	Long
	Explanation				
	This is what will be al lowing out to the Internet. Note the negated Internal_Nets. We don't want our workstations to reach anything else (except what is defined above).				

No	Source	Destination	Service	Action	Track
9.	Local_Net	Internal_Nets	CVS	Encrypt	Long
	Explanation				
	This traffic should pass through the VPN to the main site. The traffic must be originating at the local net, we do not want to allow strange packets with other IPs.				

No	Source	Destination	Service	Action	Track
10.	Local_Net	Internal_DNS	domain_udp	Encrypt	Long
	Explanation				
	This traffic should pass through the VPN to the main site. The traffic must be originating at the local net, we do not want to allow strange packets with other IPs.				

No	Source	Destination	Service	Action	Track
11.	Local_Net	FTP_Server	FTP	Encrypt	Long
	Explanation				
	This traffic should pass through the VPN to the main site. The traffic must be originating at the local net, we do not want to allow strange packets with other IPs.				

No	Source	Destination	Service	Action	Track
12.	Local_Net	Mail_Server	MUA	Encrypt	Long
	Explanation				
	This traffic should pass through the VPN to the main site. The traffic must be originating at the local net, we do not want to allow strange packets with other IPs.				

No	Source	Destination	Service	Action	Track
13.	Any	Any	Any	drop	Long
	Explanation				
	Cleanup time! Everything we did not previously allow, we want dropped and logged. Here is for example where stray telnet sessions will go (together with a lot of other peculiarities that often may be seen on the Internet), as well as intrusion attempts and port scans.				

5 Third Assignment : Audit Your Security Architecture

"For the purposes of this assignment please assume that you have been assigned to provide technical support for a comprehensive information systems audit for GIAC Enterprises, a electronic commerce facility that is the largest supplier of electronic fortune cookie sayings in the world. The firewall analyst has set their firewall up according to their base + recommended enhancements security policy that happens to mirror your assignment 1 security policy exactly. "

5.1 Preparing the Assessment

Since this is an assessment rather than a penetration testing, we will not have to figure out what the network design looks like. We have been presented with diagrams and appropriate documentation so we have a good idea about what hosts there are on the networks. But of course we will need to scan through the networks to see if any new hosts exist. And we most certainly will have the co-operation of the IT department of GIAC Enterprises. Therefore, the assessment team would invite one of the IT department to be part of the team during the work. This should assure GIAC Enterprises we know our stuff, and at the same time give the technician a chance to learn how to do his own assessments. Another good thing is that he or she would provide a witness of our actions. This will lesser the suspicions that we might be doing something we really should not. Naturally, every step taken must be carefully documented and even logged using the Sun Solaris "script" command (a program that logs all input and output to a file).

The assessment should be scheduled to take place in morning, but still within working hours. Suggested start is at 9am, with a deadline at 3pm. Since some of the scans may take long time, one day may not be enough, so the assessment team might be working for a few days in a row. The reasons to keep it within office hours are the following:

1. Cost efficient. It costs a lot to have technicians working overtime. Even though GIAC Enterprises may have 24x7 personnel, the company doing the assessment may not.

2. If, somehow, something goes wrong (it should not, but who knows), skilled staff who knows their system will be at the location and ready to help.

3. It gives management a chance to drop by and see what is being done to their system, or have a educational but friendly chat.

The reasons for keeping it between 9am and 3pm:

1. People will have arrived and had time to get into the day's routine.

2. A couple of hours after the assessment finished are available in case something caused instability not immediately noticed.

3. Since GIAC Enterprises is "the largest supplier of electronic fortune cookie sayings in the world", any disruption in the service are likely to affect many people over world, so because of time zones, no part of the day is more suited than any other. The most

important thing in case of a disruption is to get the system back in order as quick as possible, which can be done most conveniently during normal working hours .

5.2 Agreements With GIAC Enterprises, to be Made Before the Assessment

There are some matters that must be decided and agreed upon before the assessment can start. Those will be between the company offering the assessment and GIAC Enterprises:

1. Liability. It is possible that one or more system may stop working during the assessment. It must be an agreed that the Consultant Company has no liability in such cases.

2. Confidential Information. It is likely that the assessment team will gain access to sensitive information. Each member of the team must sign a NDA before beginning the work. This agreement could give the Consultant Company the possibility to keep data from the assessment for internal use.

3. Pricing and costs. Will a certain amount of time be paid fo r and nothing more (maybe six hours on three following days, at a total of 18 hours), or will the assessment team be paid for the actual hours they work?

4. Competence and number of members of the assessment team. The minimum number of members should be two. This has two reasons: first, a single person may forgot to test certain things, and second, if they are two, they will have an easier time to report what actually happened (there are legal aspects as well, if there are two persons, the first can witness what the other did, and since both are skilled in assessment business, both will probably give good account of events that occurs).

5.3 Estimate of Costs and Effort

I calculate the assessment to take between 24 and 30 hours of effective time. Trying to be f air with both parts, the deal should probably state the minimum time they have to pay for. In this case it should be 24 times two (remember, we are not working alone here). Any exceeding time up to the 30 hours per person can be used as well, but then a fe e per hour will be charged. Any work over 30 hours will be charged at a higher rate. The assessment team will bring their own laptops for testing; no software will have to be installed in production environments. The personnel doing the test will be compet ent in their area, but not among the very best either (Author's note: I have no clue how much this would cost in any other country than Sweden and here those guys would go for about 2000 SEK/hour, which is $200. However, I have a distinct feeling those guy s would cost more in the US so I suggest you calculate the overall consultant fee yourself, using the formula 30*2*(whatever a competent security consultant cost at your place). Then you will get it in your own currency as well).

There are some risks invo lved, of course. No one knows for sure how every computer will react if it is hit by a SYN scan. Should one host fail, others that are depending on it may fail as well. So I would like to stress the need for system engineers, being at place during the assessment. This adds to the overall cost as well, but it is hard to give any numbers.

5.4 Scanning for Open/Unprotected Services

5.4.1 Portscan Methodologies (Introduction to Portscanning)

As you all know, "portscanning" is the term for testing what ports responds on a remote system. There are several different ways this may be done. You can either use the TCP connect() scan meaning an attempt will be made to connect to a remote port using the full three way handshake. Or, you could send only an SYN, FIN, URG or any other TCP flags. Although you of course will not get any connections this way, the remote system will still respond. Depending on this response you will know whether the port is listening or not. You even have the X-Mas Tree flag, meaning you will send a packet containing all TCP flags at the same time. All types have their advantages and disadvantages. It is very unlikely the connect() scan will cause any harm to some older TCP stacks while X-Mas Trees may do just that. On the other hand, a firewall mig h t block incoming SYNs (please see note about SYNDefender in Assignment 2) but perhaps not an ACK, FIN or RST (if it fails to check if there actually is a connection before applying the filter).

5.4.2 The Tool of Choice

The tool of choice is of course nmap by Fyodor (http://www.insecure.org). The version I will be using for this is nmap 2.53 (below is a screen dump from the graphical front-end).

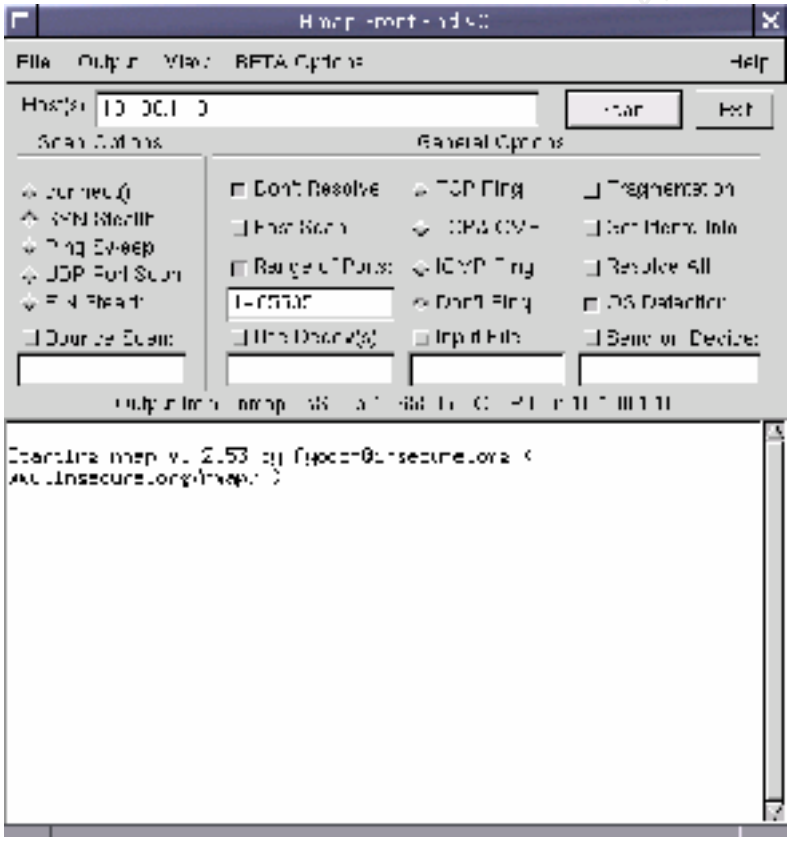

If there is a better scanner out there, I have not heard of it, and I am quite sure it is not free.

5.4.3 Portscan Methodology During the Assessment

When scanning for open ports I will first do a SYN scan and, should it not yield any result, try the connect() scan should the first one fail. The reason of my choice is simply that I do not wish to cause any Denials of Service at this stage. One must be prepared this will take some time since firewalls (or routers with ACLs) have this tendency to drop unwanted packets instead of sending RSTs or ICMP messages. For each such port that is blocked, we will have to wait for a timeout.

5.4.4 Scanning the Border Router

The Border Router is the first perimeter so it is important it is configured correctly. It would be pretty bad if anyone on the Internet could tftp new configuration files to it. The router has the IP address of 10.100 .1.254. Our first scan is from the Internet:

A Nmap scan directed at the Border Router:

```
[root@evil_hacker]# nmap –v –sS –P0 –T3 –p 1-65535 10.100.1.254
Starting nmap V2.53 by fyodor@insecure.org ( www.insecure.org/nmap/ )
Initiating SYN half-open stealth scan against (10.100.1.254)
Adding TCP port 23 (state open)
The SYN scan took 9834 seconds to scan 65535 ports
Interesting ports on (10.100.1.254):
(The 65533 ports scanned but now shown below are in state: filtered)
Port        State       Service
23/tcp      closed      telnet
Nmap run completed -- 1 IP address (1 host up) scanned in 9834 seconds
```

Here is a problem. Telnet seems to be allowed from the Internet. Otherwise, it looks good. A scan from the inside shows the same thing. We need to check it out more closely, i t could be something else. Telnet is most likely an appropriate tool.

Check on a potential security hole (1)

```
[root@evil_hacker]# telnet 10.100.1.254
Trying 10.100.1.254…
Connected to 10.100.1.254.
Escape character is '^]'.

User Access Verification

Password:
```

This looks like a Cisco router to me, so some common cisco default passwords (cisco, admin, router and blank) should be tested.

Check on a potential security hole (2)

```
Password:
Password:
Password:
% Bad passwords
Connection closed by foreign h ost.
```

None of the common passwords helped to gain access, so the password seems to be sensible enough. At least something!

Recommendation: Close down and block the telnet service, a border router should not be managed remotely, and especially not with t elnet.

5.4.5 Scanning the Firewall

This scan must be done from the DMZ; otherwise, the border router may block certain traffic.

A Nmap scan directed at the firewall

```
[root@evil_hacker]# nmap  –v –sS –P0 –T3 –p 1-65535 10.100.1.1
Starting  nmap V2.53 by fyodor@ insecure.org ( www.insecure.org/nmap/ )
Initiating SYN half-open stealth scan against  (10.100.1.1)
Interesting ports on  (10.100.1.1):
(The 65533 ports scanned but now shown below are in state: filtered)
Port        State       Service
113/tcp     closed      auth
139/tcp     closed      netbios-ssn

Nmap run completed  -- 1 IP address (1 host up) scanned in 10103 seconds
```

This looks good, nothing serious at all. The same scan would have to be applied from the inside to make sure it is safe from both directions.

Recommendation: Since the border router block NetBIOS, there is really no need for the firewall to reject them. Drop them, and turn the firewall into a black hole instead.

5.4.6 Testing the Rulebase

Now that we know our firewalls to be secure, we can begin testing if they do their job protecting the internal hosts. Again, we will be sitting on the DMZ, using nmap to see what traffic we may get through. As we will not know by looking at the nmap result whether the packet was dropped by the packet filtering firewall or by the interna l host, it is important to watch the firewall log as well.

For this test, we will nmap the DNS, HTTP and HTTPS Servers. A change in our tactics concerning nmap is that we will use TCP connect scan instead of SYN scan. This is because we do not want any in terference by the SYNDefender (We would not know whether the firewall or the host on the other side answered).

Nmap against the three internal hosts

```
[root@evil_hacker]# nmap  –v –sS –P0 –T3 –p 1-65535 10.100.1.10
Starting  nmap V2.53 by fyodor@insecure.o rg ( www.insecure.org/nmap/ )
Initiating SYN half-open stealth scan against  (10.100.1.10)
Adding  TCP port 80 (state open)
Interesting ports on  (10.100.1.10):
(The 65532 ports scanned but now shown below are in state: filtered)
Port        State       Service
80/tcp      open        http
113/tcp     closed      auth
139/tcp     closed      netbios-ssn

Nmap run completed  -- 1 IP address (1 host up) scanned in 10142 seconds

[root@evil_hacker]# nmap  –v –sS –P0 –T3 –p 1-65535 10.100.1.11
Starting  nmap V2.53 by fyodor@insecure.org ( www.insecure.org  /nmap/ )
Initiating SYN half-open stealth scan against  (10.100.1.11)
Interesting ports on  (10.100.1.11):
(The 65532 ports scanned but now shown below are in state: filtered)
Port        State       Service
113/tcp     closed      auth
139/tcp     closed      netbios-ssn
443/tcp     open        https

Nmap run completed  -- 1 IP address (1 host up) scanned in 10130 seconds

[root@evil_hacker]# nmap  –v –sS –P0 –T3 –p 1-65535 10.100.1.12
Starting  nmap V2.53 by fyodor@insecure.org ( www.insecure.org/nmap/ )
Initiating SYN half-open stealth scan again st  (10.100.1.12)
Interesting ports on  (10.100.1.12):
(The 65533 ports scanned but now shown below are in state: filtered)
Port        State       Service
113/tcp     closed      auth
139/tcp     closed      netbios-ssn

Nmap run completed  -- 1 IP address (1 host up) scanned in 10160 s  econds
```

Not very surprising. Looking from the DMZ, the firewall sure did its magic. On the first server we found http open (the web server), on the second we found https (the ssl server) and on the third we found nothing (we do not have the IP of the exte rnal DNS so we were not permitted to do any Zone Transfer. On all servers, ident and NetBIOS were rejected, just as it should.

So what about the firewall log?

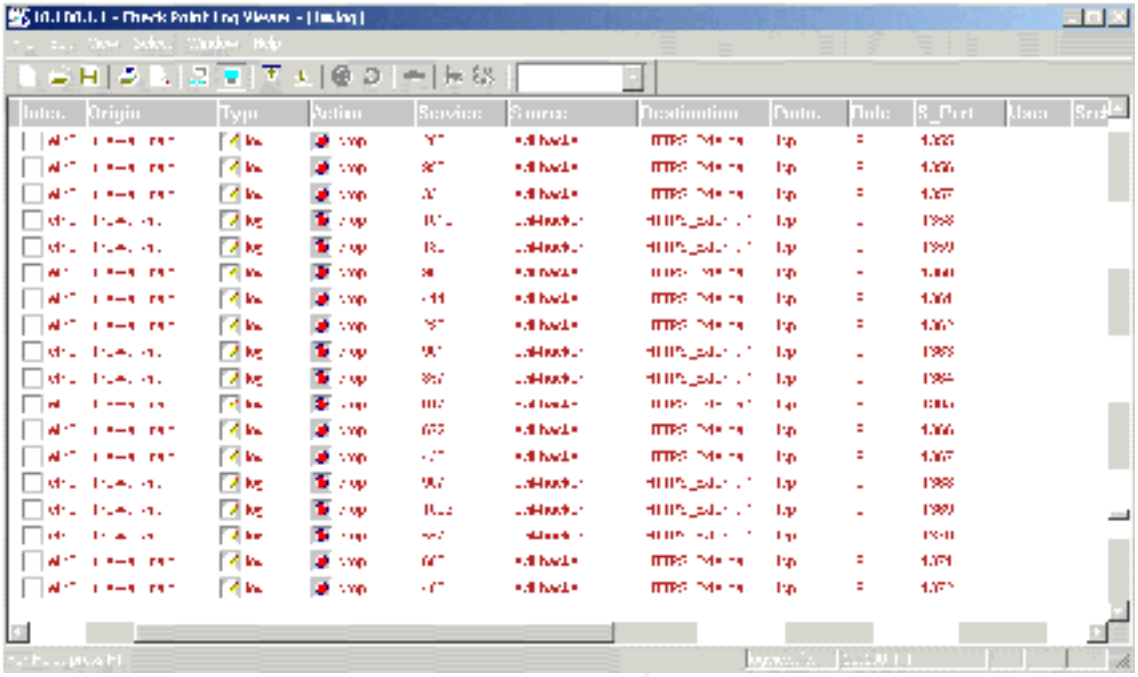

The above dump appears to be in the middle of the port scanning. A lot of drops can be seen in the logs, connections originating from the evil -hacker and targeting the translated address of the HTTP Server.

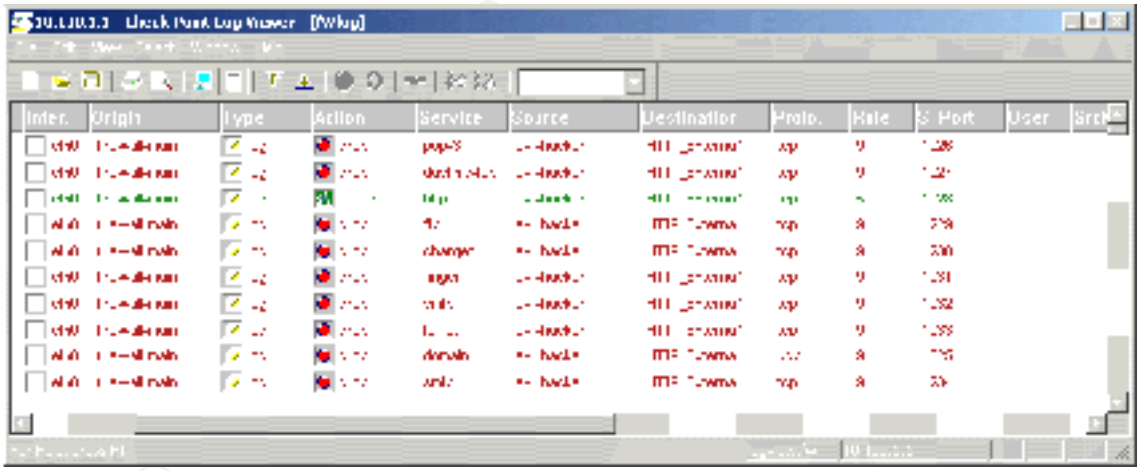

Above we have several drops and then one accept. The hacker found out that the HTTP server was listening on port 80. Except that one por t, everything looks red.

5.5 Confirming That ICMP Traffic Behaves as Expected

This is a quick test to make sure people from the Office_Networks can ping the Internet but not internal servers. And that Internet may not ping anything internal.

Pinging an Internet host from the Office Network

```
[admin@ws12]$ ping www.sans.org
www.sans.org is alive
```

Pinging an Internal Server from the Office Network

```
[admin@ws12]$ ping mail.giac -enterprises.com
no answer from mail.giac -enterprises.com
```

Pinging an Internal Se rver from the Internet

```
[admin@ws12]$ ping mail.giac -enterprises.com
no answer from mail.giac -enterprises.com
```

ICMP seems to work just fine. We will leave it at that (and yes, it was verified that the mail server was up and running).

5.6 Scanning for New Hos ts

Routine check: No new hosts or interfaces should be found other than those correctly documented. Below I view nmap in use at a whole network. Since no irregularities were found in any network, I will only show the first one.

Searching the Screened Net work

```
[root@evil_hacker]# nmap 192.168.1.0/24  –p 1-1024

Starting  nmap V2.53 by fyodor@insecure.org ( www.insecure.org/nmap/ )
Interesting ports on  (192.168.1.1):
(The 1022 ports scanned but not shown below are in state: filtered)
Port        State       Service
113/tcp     closed      auth
139/tcp     closed      netbios-ssn

Interesting ports on  (192.168.1.10):
(The 1023 ports scanned but not shown below are in state: closed)
Port        State       Service
80/tcp      open        http

Interesting ports on  (192.168.1.11):
(The 1023 ports scanned but not s hown below are in state: closed)
Port        State       Service
443/tcp     open        https

Interesting ports on  (192.168.1.12):
(The 1023 ports scanned but not shown below are in state: closed)
Port        State       Service
53/tcp      open        domain

Nmap run completed  – 256 IP addresses (4  hosts up) scanned in 523 seconds
```

5.7 Checking DNS

The DNS is a very important part of the system. Should it fall victim to an attacker, domain names could be changed to point to other sites, tricking customers to give away passwords or credit card numbers . We will check a couple of things just to make sure everything is in order.

5.7.1 Zone Transfers

We want to know whether the DNS Servers permits zone transfers. If it is possible, the data could be used for deciding which hosts on the network that would be th e easiest to attack. The tool I use for this is "dig", which comes with the ISC BIND installation. For this check to be efficient, the laptop must be placed on the same network segment as the Name Server (remember, the firewalls are blocking port 53 TCP fr om everyone except the secondary DNS). 192.168.1.12 is the internal IP address of the DNS.

Testing for Zone Transfers

```
[root@evil_hacker]# dig @192.168.1.12 axfr giac -enterprises.com

; <<>> DiG 8.2 <<>> @192.168.1.12 axfr giac -enterprises.com
; (1 server found)
;; Received 0 answers (0 records).
;; FROM: evil_hacker to SERVER: 192.168.1.12
;; WHEN: Tue Nov 15 20:01:10 2000
```

Since we got zero answers and zero records back, the DNS Server seems to have strict ideas about whom it is going to exchange zones with. But remember, there are another DNS out there that contain our zones: the secondary one, at our ISP.

Zone Transfer Check no. 2 – ISP's DNS (our secondary)

```
[root@evil_hacker]# dig @10.300.20.22 axfr giac -enterprises.com

; <<>> DiG 8.2 <<>> @10.30 0.20.22 axfr giac -enterprises.com
; (1 server found)
;; Received 0 answers (0 records).
;; FROM: evil_hacker to SERVER: 10.300.20.22
;; WHEN: Tue Nov 15 20:05:14 2000
```

No problem there either. Very nice.

5.7.2 Version of DNS Server

We need to figure out what version of DNS server is running. Again the tool is from the freely available ISC BIND source code. But now we are going to use "nslookup" instead of "dig".

Getting version of DNS Server

```
> set q=txt
> set class=chaos
> version.bind
Server:  ns.giac-enterprises.com
Address:  10.100.1.12

VERSION.BIND    text = "8.2.2 -P5"
```

The DNS Server claims it is running ISC BIND version 8.2.2 -P5 (patch level 5). This will have to be checked with the admin of the site since it could be a spoof, but the probability of this is not very high.

Recommendation: Upgrade to patch level seven (8.2.2 -P7), which contain a couple of security fixes. Also, change the text entry to something else, giving potential intruders a harder time to figure out what system it is. It does no t have to be a real name.

5.7.3 Recursion

There is one more thing we want to check; whether recursion is allowed or not. If it is, it could make life easier for anyone wanting to attempt a DNS cache poisoning, and furthermore, there is really no use in having the DNS work for other people. Once again, we are going to use "nslookup"

Checking for allowed recursion

```
>set q=a
>server ns.giac-enterprises.com
Server: ns.giac-enterprises.com
Address: 10.100.1.12

> www.sans.org
Server:  ns.giac-enterprises.com
Address:  10.100.1.12

Name:    www.sans.org
Served by:
- M.ROOT-SERVERS.NET
        202.12.27.33

- I.ROOT-SERVERS.NET
        192.36.148.17

- E.ROOT-SERVERS.NET
        192.203.230.10

- D.ROOT-SERVERS.NET
```

```
        128.8.10.90

- A.ROOT-SERVERS.NET
        198.41.0.4

- H.ROOT-SERVERS.NET
        128.63.2.53

- C.ROOT-SERVERS.NET
        192.33.4.12

- G.ROOT-SERVERS.NET
        192.112.36.4

- F.ROOT-SERVERS.NET
        192.5.5.241

- B.ROOT-SERVERS.NET
        128.9.0.107
```

Since we got back "Served by" and a list of root servers, the DNS server obviously do not allow recursion. Well and good.

5.8 Mail Relay Check

One thing that could prove embarrassing is if the mail sy stem accepts and relays mail it should not. Sooner or later, someone will discover it and use it for UCE (Unsolicited Commercial Email, aka spam) or harassment. So we will just to a quick check and make sure the Mail server is configured properly. Since SM TP require no special control characters as the telnet protocol does, we can use netcat instead. Netcat (nc) is cleaner, smoother and faster. It is the Swiss Army Knife of networking.

Checking for Open Mail Relay

```
[root@evil_hacker]# nc mail.giac -enterprises.com 25
220 mail.giac-enterprises.com ESMTP Postfix Mail Daemon Running
helo evilspammer.net
250 mail.giac-enterprises.com
mail from: spammer@evilspammer.net
250 Ok
rcpt to: security@consultant.com
554 < security@consultant.com >: Recipient address reje cted: Relay access denied
```

The system seems to be configured as it should. However, to make the life harder for those banner grabbing types, something should be done about the banner.

Recommendation: Change the string "Postfix Mail Daemon Running" to so mething less descriptive. Perhaps "Some Kind Of MTA" would do the job.

5.9 Anti-virus Testing

Nothing complicated but nevertheless a check that needs to be done. We want to verify the virus protection. An email will be sent to the contact person at GIAC Ent erprises, containing the "virus" EICAR.COM. Actually, it is not a virus in the real sense of the word, but rather a DOS executable that prints out " EICAR -STANDARD-ANTIVIRUS-TEST-FILE". Most (if not all) vendors of anti-virus software detects EICAR (more i nformation about EICAR can be found at http://www.eicar.org). If the file is allowed through the Anti -virus Server, something is amiss. Otherwise, the system is working. Naturally, the assessment team could offer to send a new "in-the-wild" virus to check whether the definitions are updated correctly.

5.10 Perimeter Analysis

5.10.1 Recommended Actions (summary)

1. Upgrade the DNS (ISC BIND) to patch level seven (8.2.2 -P7). The current patch level contain a few bugs which could be used for causing Denial of Service con dition which contain a couple of security fixes. At the same time, change the bind.version text to something less descriptive.

2. Change the Postfix banner to contain other information.

3. Modify the firewall rules to drop NetBIOS instead of rejecting it.

5.10.2 Suggested Architecture Improvements

Since we both had the network diagram and had a chance to chat with the primary security technician, some suggested architecture improvements should be presented.

1. Remove all access from the Office networks to the Internet except the SMTP that still must function for the company to operate. This is proposed because there are trojans that communicate with their servers over ICMP or HTTP. Still, Internet connection may be handy so the solution may be to deploy another network where a couple of workstation will be placed, protected by a firewall that will allow everything outbound but nothing inbound. If agreeable by management, this environment could be used by employees wanting to surf the web as well as technicians seeking i n knowledge bases. A slightly less dramatic solution could be to require client authentication prior to accessing the Internet.

2. Add "PCs" to the IDS network, using dual interfaces and a monitor plus database in the back end. Those should be running both A ntiSniff (a program that checks for Network Interface Cards running in promiscuous mode). AntiSniff will not be able to detect commercial IDSes running without MAC or IP but may well find other hosts that suddenly takes an unseemly interest in the network traffic. They should also be running Arpwatch, just for keeping an eye out for arp floods or similar things. More information about AntiSniff can be found at "http://www.l0pht.com/antisniff/" whereas Arpwatch comes with almost any Linux Distribution.

3. Consider implementing Tivoli or OpenView to get a centralised management platform.

4. Implement a new server for NTP (Network Time Protocol), log files do not make much sense if the all hosts have different times. Preferably, the master NTP server should get the time from three or more GPS receivers.

6 Closing Down

This is the end of this work. Time do not allow for more; after all we all have normal work to take care of during office hours (and sometimes evenings too). Thanks for taking your time and reading it through (or did you start at bottom line and began reading upwards?). I sincerely hope you think it were worth your time.