



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

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

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

GFCW Practical Assignment

Author: Todd Chapman

Date: April 4, 2001

1. Security Architecture

Specification:

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

You must define access for:

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

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

Network Design:




Figure 1. Security Architecture

Figure 1 illustrates the security architecture designed to satisfy the specification. The philosophy behind the architecture is to keep the design simple, provide multiple defensive layers, and segment access to information on a need to know basis.

The first layer of defense is a Cisco router model 3640 running IOS 12.1. In addition to providing Internet connectivity, the router is configured with a simple set of packet filtering rules. These rules are designed to filter out common attacks such as IP address spoofing, blocking private addresses, controlling ICMP traffic, and preventing source routing. The packet filtering router is a good first line of defense but it is important

that it's rulebase does not impact the performance of it's primary job, routing. "Don't mirror the firewall rules." [Spitzner 2001] The router will be configured to send logging information to the Syslog server on the internal network.

The second layer of defense is two Checkpoint Firewall-1 2000 firewalls running on Sparc based Solaris 2.6 systems. The firewalls are configured for high availability using Stonesoft's, StoneBeat FullCluster software. This configuration allows the second firewall to take over the duties of the first if the first firewall fails due to hardware problems, denial of service attack, or is simply removed from the configuration for maintenance.

Unlink the perimeter router, the sole purpose of the firewall is to enforce the GIAC Enterprises security policy. While performance and connectivity are important, they must not come at the expense of security. For this reason the firewall rulebase denies all traffic by default, and then only traffic that is specifically permitted is allowed to pass. Firewall-1 uses stateful inspection to keep track of connection status and make decisions about when to forward network traffic or not. While stateful inspection is generally not as secure as an application gateway, it's higher performance makes a good choice for a perimeter firewall.

By using Checkpoint's VPN-1 Gateway on the perimeter firewall, secure access can be provided to partners and suppliers through encrypted VPN connections. In addition remote access can be safely provided to employees while traveling by using the VPN capabilities of VPN-1 along with Checkpoint's SecureClient. No dial-up access is allowed directly into the GIAC Enterprises network.

SecureClient was chosen over SecuRemote because of it's built in firewall capabilities. With SecuRemote, if the remote access client is compromised, the remote system may be used to gain access through the encrypted tunnel into GIAC Enterprises. SecureClient prevents connections from being accepted by the remote system while it is connected to the GIAC network through the tunnel.

Directly connected to the Firewall-1 systems is a service network. This network includes systems that need to be accessible from the Internet but do not need an additional layer of security. The DNS systems on the service network are part of the split DNS design. They only contain information that needs to be resolvable from the Internet. The DNS servers are OpenBSD systems running BIND. Although BIND has a history of vulnerabilities, it is run as a non-root user in a chroot jail to further protect the system in case a BIND security hole is exploited.

The mail system on the service network acts as an application proxy for incoming e-mail. This is a Solaris 2.6 system running Norton AntiVirus for Gateways version 2.1. This system is used to guard against denial of service attacks against the mail system and scans attachments for viruses.

The third layer of defense is a Symantec Raptor 6.5 firewall running on a Windows NT 4 workstation. This firewall provides a extra layer of defense for the web server network, database server network, and other servers on the internal and corporate networks. By using application proxies instead of stateful inspection and NT instead of Solaris, additional security is achieved because vulnerabilities in one technology or platform are unlikely to be found in the other.

Not only does the Raptor firewall protect various internal resources from the Internet, but it also protects those resources from each other. By placing the database server and web server each on their own network, access to these resources by VPN connections and employees on the internal network can be further restricted.

The Raptor firewall also enforces GIAC Enterprises' policy on Internet usage by restricting and monitoring the services and destinations that it's employees may use.

The public network, service network, and the network between the two firewall layers are monitored by OpenBSD systems running version 1.7 of the Snort intrusion detection system. Each switch on the network supports port mirroring. Time synchronization across all systems on the network is maintained using NTP so that logs and system events can be correlated during or after an attack.

Each firewall, server, and IDS host on the network is hardened to reduce the risk of being compromised. The minimum amount of software is installed and only necessary services are allowed to run. Tripwire is run on each system to detect intrusions.

Security Procedure:

In addition to the security architecture, which was represented in large part by Figure 1, we also employ a security procedure. The security procedure is comprised of other tools and processes used to maintain and monitor the security policy.

Before all other parts of the security procedure comes training. Security training is essential to ensure that security engineers have the appropriate understanding of the technologies they work with and the threats they are presented with. There are many sources of security training. Usually a mix of product training and vendor neutral training provide the best understanding of specific technologies and the big picture. Excellent vendor neutral training can be obtained through the SANS Institute (<http://www.sans.org/>).

The first part of the security strategy is the daily review of all logs. These include the firewall logs, IDS logs, and logs from all other systems as collected by the Syslog server. Perl will be used to write filters that help reduce the number of redundant and uninteresting log messages that are read each day. Unfiltered logs should be archived to tape or CDR media so that if a break in has been discovered the logs can be reviewed to see when the compromise might have occurred and what evidence of those events might have been missed.

The second part of the security procedure is to subscribe to security mailing lists for each vendor's product used in the security architecture. Vendor neutral mailing lists such as those found at www.securityportal.com should also be subscribed to. These lists should be reviewed on a weekly basis. Any applicable patches or service packs should be installed as soon as possible.

In addition sites such as the Global Incident Analysis Center (<http://www.sans.org/giac.htm>) should be monitored for breaking information on security incidents and detects. The SANS Top Ten Threats (http://www.sans.org/infosecFAQ/threats/top_ten.htm) should also be periodically reviewed for applicability to the security architecture.

Next, the configuration of all components of the security architecture is guided by a change control process. All changes must be reviewed by the security team and documented. The use of sudo (<http://www.courtesan.com/sudo/>) and Tripwire (<http://www.tripwire.com/>) help track these changes.

All communication on GIAC Enterprises networks should be encrypted when possible. This means ssh and scp in place of telnet, rsh, rcp, and ftp. In some cases telnet must be used (Cisco router) but should only be used from within the GIAC network. Never from the Internet. This prevents the router password from crossing networks not under our control, as plain text. The default passwords must be changed on all systems. Guest accounts should be removed.

The login banners should be changed on all systems. The login banner should contain a warning that only authorized use is permitted, and the banner should not give an indication of the type or version of the system. Hiding the system type and version makes is more difficult for potential intruders to attack the system based on known vulnerabilities because they first have to expend effort identifying what they are attacking.

Lastly, periodic auditing and vulnerability testing of the security architecture should be conducted. Without these audits there is a great risk that the security policy as documented will drift from the implementation. Temporary changes to the security policy may become permanent or new vulnerabilities may be overlooked. The processes and tools for these audits are covered later in this document.

2. Security Policy

Specification:

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

- Border Router
- Primary Firewall
- VPN

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

Border Router Security Policy:

As stated previously, a simple set of rules are to be installed on the router to prevent simple attacks. The router configuration is listed and explained below. The configuration is divided into three sections. Basic configuration defines configuration parameters that apply to the router as a whole while each Access Control List (ACL) section defines the packet filtering rules for one of the router's interfaces.

Basic configuration:

line console 0 login password XXXXXXXX	Set a password for console access.
enable secret YYYYYYYY	Set a password for the IOS enable mode.
service password-encryption	Stores the password encrypted. The passwords are still sent to the router in plain text!
line vty 0 4 login password ZZZZZZZZ	Set a password for telnet access.
banner / WARNING: Authorized Access Only /	Create a warning banner.
logging 192.168.4.5 logging trap debug logging console emergencies	Configure logging parameters and have the router send logs to the IP address of our Syslog server.
no service udp-small-servers no service tcp-small-servers no service finger no ip http no ip bootp	Turn off some common services we don't want to pass through our router or server ourselves.
no cdp running no cdp enable	Turn off Cisco Discovery Protocol. What device and IOS version we are using is nobody's business but ours.
no ip source-route	Deny source routed packets. These packets might get routed back to a host that is spoofing the IP address of another host.
no ip directed-broadcast	Deny directed broadcast packets. These are sometimes used by attackers trying to amplify a DOS attack against another party.
no snmp	Turn off snmp service. If this service is required by GIAC Enterprises a hard to guess community name should be set.
no ip proxy-arp	Turn off proxy-arp service.

no ip unreachable	This rule prevents the router from accepting packets that may alter the routing tables of hosts on our network.
-------------------	-----------------------------------------------------------------------------------------------------------------

Serial Interface Access Control List:

This access list defines packet filtering rules for packets entering the router's serial interface from the Internet. By defining filters that are applied when the packets enter an interface instead of when a packet exits an interface, we save router resources by dropping packets that would otherwise be routed and then dropped as they try to exit the router. Sometimes input filters are not enough to get the job done, but for our simple configuration they work very well. (A backslash indicates the continuation of a line.)

interface serial 0	Identify the interface we are configuring.
ip address 223.2.246.1 255.255.255.252	Set the IP address and netmask of the interface.
ip access group 112 in	Apply access list 112 to the interface.
access-list 112 deny ip 10.0.0.0 0.255.255.255 access-list 112 deny ip 127.0.0.0 0.255.255.255 access-list 112 deny ip 172.16.0.0 0.15.255.255 access-list 112 deny ip 192.168.0.0 0.0.255.255	These four rules deny packets with IP addresses defined by RFC 1918 as being private addresses that should not be routed across the Internet.
access-list 112 deny ip 224.0.0.0 31.255.255.255	Do not allow multicast packets.
access-list 112 deny host 0.0.0.0	Do not allow packets with a host address of all zeros.
access-list 112 permit ip any any	Once an access list is created an implicit default deny rule is created. Because we want to pass all but a narrowly defined list of packets we have to add this rule that allows all other IP traffic to enter this interface.

Ethernet Interface Access Control List:

This access list defines packet filtering rules for packets entering the router's serial interface from the GIAC Enterprises public network.

interface ethernet 0	Identify the interface we are configuring.
ip address 223.2.245.1 255.255.255.240	Set the IP address and netmask.
ip access group 115 in	Apply access list 115 to packets coming in to this interface.
access-list 115 deny tcp any any 135 access-list 115 deny tcp any any 139 access-list 115 deny udp any any range 137 178 access-list 115 deny tcp any any 445 access-list 115 deny ucp any any 445	Block packets from common Microsoft services from leaving our network.
access-list 115 permit ip 223.2.245.0 \ 255.255.255.240	This rule allows packets to enter the router if they have a source IP address from the public part of our network.
access-list 115 deny ip any any	This rule denies all other traffic entering this interface.

Primary Firewall Security Policy:

The primary firewall in our security architecture are the two Firewall-1 (FW-1) firewalls that are configured for high availability using StoneBeat. The configuration of StoneBeat will not be covered in this document. Instead we will pretend that there is only one FW-1 firewall. In either case the configuration is

very similar, except that in the StoneBeat case there is a shared IP address for the pair of interfaces on each network.

In other words, for the public network the firewalls have IP addresses 223.2.245.2 and 223.2.245.3, respectively. In addition the IP address 223.2.245.4 is advertised as the IP address of the web server. If the primary firewall fails, the secondary firewall takes over this IP address and continues to server incoming requests to that IP address.

Before defining our Firewall-1 rulebase a discussion on how rules are matched and what makes up a rule would be beneficial. In general rules are matched top down. That is, the rules are numbered in the FW-1 policy editor and these rules are checked starting at the first rule and consecutively checking each rule until there is a match. Once a match is found the appropriate action is taken as indicated by the rule and no more rules are checked

This is complicated by the fact that FW-1 actually has two kinds of rules, explicit and implicit. Explicit rules and those that the administrator manually enters into the rulebase. Implicit rules are rules that FW-1 puts in the rulebase as a result of settings in the Properties Setup screen in FW-1. (Figure 2) These settings include IP options and IP spoofing settings. Implicit rules are not displayed by default in the rulebase. They can be revealed by turning on Implied Pseudo-Rules in the View menu of the Policy Editor. Rules are ordered in the rulebase in the following order:

Match Order	
1	IP Spoofing / IP Options
2	Security Policy "First" Rule
3	Administrator defined rules.
4	Security Policy "Before Last" Rule
5	Last administrator defined rule.
6	Security Policy "Last" Rule.
7	Implicit Drop Rule

An important property in the FW-1 Security Policy Properties Setup screen is the setting for which direction rulebase rules are applied in. The possible settings are inbound, outbound, and eitherbound. The most common interface direction setting is inbound and that is the setting that we will use.

When set to inbound, packets are matched against the rulebase as they enter the firewall, regardless of which interface they are entering. The advantage of this setting is that packets are inspected before they enter the firewall. The disadvantage is that packets originating from the firewall are not inspected. If an untrusted user has access to the firewall, that users connections will not be filtered by the rulebase.

The outbound setting is not used because it leaves the inbound interface vulnerable. The eitherbound setting is more secure because it checks packets going into and out of each interface. Checking each packet multiple times causes a performance degradation and is therefore not used because.

Figure 2 below shows the Firewall-1 security policy properties and how implied rules can be placed at the beginning of the rulebase, at the end of the rulebase, or before the last rule in the rulebase. Figure 2 also shows where the interface direction setting is set.

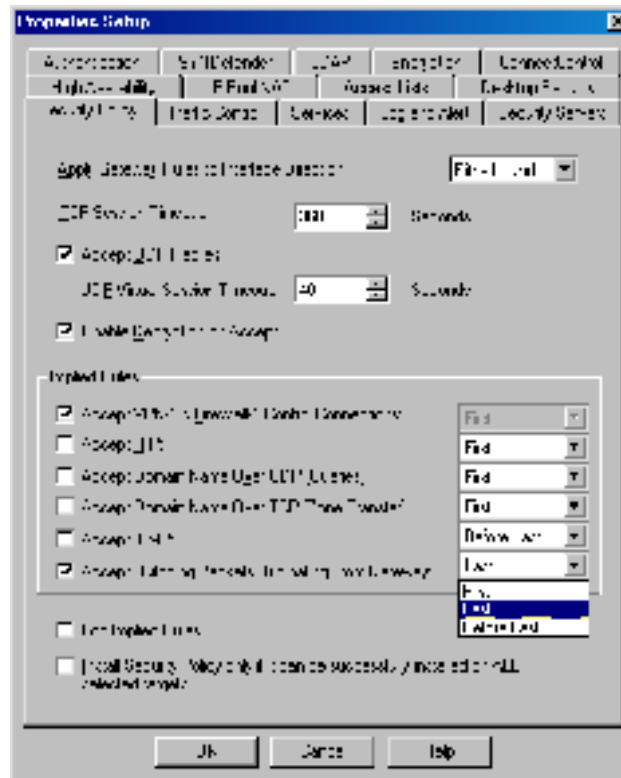


Figure 2. Security Policy Properties

Before describing the rulebase we will define management objects to use in our rulebase. These include workstation objects, network objects, and service objects. These objects help simplify the rulebase by replacing addresses and port numbers with more meaningfully named and color coded representations. The Service objects we use are pre-defined in Firewall-1.

Workstation Objects:

Workstation Objects represent individual machines including workstations and servers.

Object Name	Address
web.giac.com	192.168.3.2
smtp.giac.com	192.168.1.12
Secondary-DNS	160.4.3.2
Syslog-Server	192.168.4.5
fw1.giac.com	223.2.245.2, 192.168.1.2, 192.168.2.2
fw2.giac.com	223.2.245.3, 192.168.1.3, 192.168.2.3
dns1.giac.com	192.168.1.10
dns2.giac.com	192.168.1.11
Int-SMTP	192.168.4.4

Network Objects:

Network Objects represent entire networks of IP addresses.

Object Name	Network Number
Public-Net	223.2.246.0/28

Service-Net	192.168.1.0/24
Internal-Net	192.168.4.0/24
Partner-Net	Network number of applicable partner.

Group Objects:

Group Objects allow multiple Workstation and Network objects to be referenced under one name.

Group Object Name	Grouped Objects
fw.giac.com	fw1.giac.com, fw2.giac.com
Ext-DNS	dns1.giac.com, dns2.giac.com
Internal-Networks	Internal-Net plus any other applicable internal networks.
SecurityGroup	Workstation objects representing members of the security group.
RAservers	Workstation objects representing resources that should be made available to remote users.

Rulebase:

Each rule in the FW-1 rulebase has 9 components. They are:

1. Rule Number, assigned automatically by FW-1.
2. Source address.
3. Destination address.
4. Service.
5. Action (Accept, Drop, Reject, etc.)
6. Track (Short Log, Long Log, Alert, etc.)
7. Install On (which firewall to install the rule on)
8. Time (when the rule is in effect)
9. Comment (Descriptive text to remind the admin. what the rule was for.)

For this document components 7 and 8 will be ignored since they are the same for every rule. Each rule will be installed on both of our FW-1 firewalls and the rules will have no Time component, thus being in effect at all times. Also, since each rule will be described in detail the Comment field will also not be listed.

No.	Source	Destination	Service	Action	Track
1	Any	fw.giac.com	Any	Drop	Long
2	Partner-Net, Internal-Networks	Partner-Network, Internal-Networks	Any	Encrypt	Long
3	Any	web.giac.com	HTTP/S	Accept	Long
4	Any	smtp.giac.com	SMTP	Accept	Long
5	Any	Ext-DNS	domain-udp	Accept	Long
6	Secondary-DNS	Ext-DNS	domain-tcp	Accept	Long
7	Public-Net, Service-Net	Syslog-Server	Syslog	Accept	Long
8	smtp.giac.com	Int-SMTP	SMTP	Accept	Long
9	Internal-Networks	Any	HTTP/S	Accept	Long
10	SecurityGroup	Cisco	Telnet	Accept	Long
11	SecuirtyGroup	Public-Net, Service-Net	SSH	Accept	Long
12	RemoteUsers	RAservers	Any	ClientEncrypt	Long
13	Any	Any	Any	Drop	Long

Rule number 1 is commonly known as the “Stealth Rule.” Any connections directly to the firewall will be dropped. This helps prevent attacks on the interfaces of the firewall itself. The firewall can still be managed remotely and accept VPN connections because of the implicit rule that FW-1 creates when the security policy property “Accept VPN-1 & Firewall-1 Control Connections” is enabled.

Rule number 2 allows VPN connections between our internal networks and those of a specified partner.

Rule number 3 allows connections to our web server. Because the web server is protected by an additional firewall we don't have to worry about who can connect to which IP address on the firewall. This will allow us to restrict customers, suppliers, and partners to different virtual servers. As GIAC Enterprises grows these virtual servers can be split into physically separate servers.

Rule number 4 allows mail connection to our external SMTP server.

Rule number 5 allows DNS lookups from the Internet on our external DNS servers. These are restricted to UDP DNS queries as we don't expect a DNS query result to be too large for a DNS reply. This helps reduce the number of possible attacks on the BIND daemon.

Rule number 6 allows our secondary DNS server on the Internet to do zone transfers.

Rule number 7 allows servers on our public and private nets to send syslog messages to our syslog server.

Rule number 8 allows our external mail server to forward messages to our internal mail server.

Rule number 9 lets users on our corporate network browse the Internet.

Rule number 10 allows members of the security group telnet access to the router to make configuration changes.

Rule number 11 allows members of the security group use ssh to connect to servers on the public and service networks.

Rule number 12 is the rule that allows remote users use SecureClient to securely access internal resources. Only a select number of servers may be accessed through these connections.

Finally, rule number 13 is commonly known as the “Cleanup Rule”. This rule drops and logs any packets which have not already been matched. While FW-1 automatically creates an implicit drop rule at the end of the rulebase, that rule doesn't do any logging, making the cleanup rule necessary.

Figure 3 shows the Firewall-1 policy editor with part of our rulebase displayed.

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

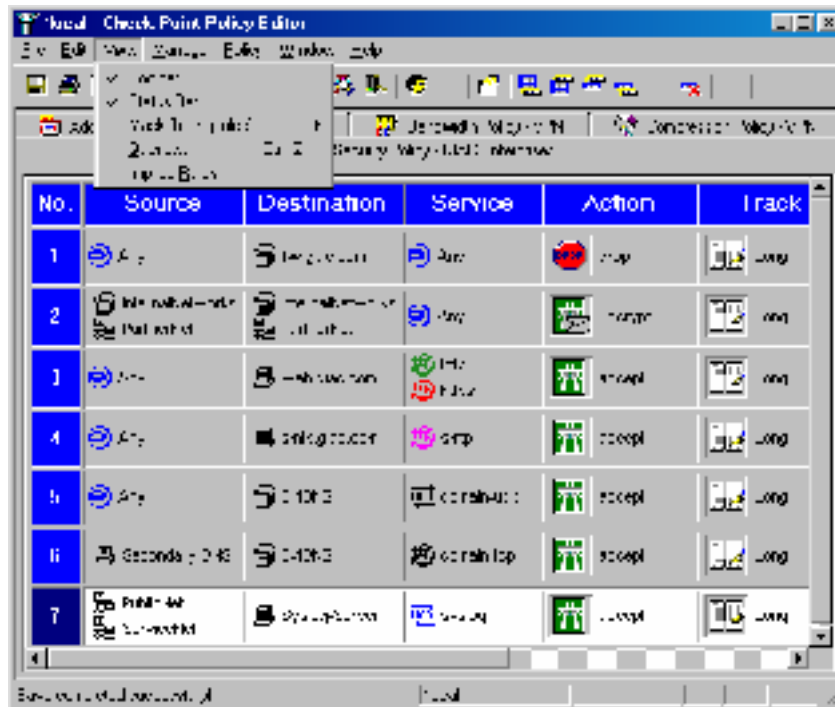


Figure 3. Policy Editor

Network Address Translation:

NAT (Network Address Translation) is used in our security architecture to accomplish two tasks:

1. Hide our web server and other servers on our service network behind IP address aliases on the firewall.
2. Hide outgoing HTTP and HTTPS connections from our corporate network behind one IP address alias on the firewall.

In both cases the computers on our networks have private addresses. NAT allows these system to communicate with systems on the Internet by having their private addresses translated to legal public addresses by the firewall.

For each publicly accessible server on our private network, we have to configure NAT using the NAT tab of that workstation object. Packets leaving the GIAC network for the Internet are translated using Static Source Mode. Packets returning to the GIAC network from the Internet are translated back using Static Destination Mode. This is accomplished with the following steps:

1. Bring up the Workstation Properties dialog box for the appropriate server and click on the NAT tab.
2. Select the "Add Automatic Address Translation Rules" option.
3. Choose "Static" as the Translation Method.
4. For Valid IP Address, specify a legal, publicly available address to use and click OK. For example: 223.2.245.4
5. Configure the Solaris based firewall to correctly accept and route packets for the "Valid IP Address" using the *arp* and *route* commands.

Figure 4 demonstrates NAT setting for the web server.



Figure 4. Setting static mode NAT for the web server.

To hide an entire internal network behind one public IP address, Firewall-1's Hide Mode is used. The steps for configuring Hide Mode are:

1. Bring up the Network Properties dialog box for the appropriate internal network and click on the NAT tab.
2. Select the "Add Automatic Address Translation Rules" option.
3. Choose "Hide" as the Translation Method.
4. For Valid IP Address, specify a legal, publicly available address to use and click OK. For example: 223.2.245.5
5. Configure the Solaris based firewall to correctly accept and route packets for the "Valid IP Address" using the *arp* and *route* commands.

Configuring the VPN Connection:

For GIAC Enterprises VPN connections, the IKE (Internet Key Exchange) encryption scheme is chosen when supported by the remote network. IKE was chosen because it is standards based and supports key exchange. For the security protocol, Authentication Header (AH) was chosen. AH was primarily chosen because many of GIAC Enterprises' partners are likely to be located in foreign countries where using DES, which is mandated by ESP (Encapsulating Security Payload) may not be legal. For the encryption protocol 3DES will be used when legal and exportable DES otherwise. MD5 will be used for data integrity. The steps for configuring IKE are:

On the local firewall:

1. Open the Workstation Properties dialog for the local firewall and select the VPN tab.
2. For the encryption Domain select the network object represent the partners network.
3. Select IKE as the Encryption scheme and click Edit.
4. Select the Pre-Shared Secret option and click Edit Secrets.
5. Enter a Pre-Shared Secret for the authentication method.
6. Select OK twice to exit the dialogs.

On the remote firewall:

7. Open the Workstation Properties dialog for the remote firewall and select the VPN tab.
8. For the encryption Domain select the network object represent the GIAC Enterprises network.
9. Select IKE as the Encryption scheme and click Edit.

10. Select the Pre-Shared Secret option and click Edit Secrets.
11. Enter a Pre-Shared Secret for the authentication method.
12. Select OK twice to exit the dialogs.

On both local and remote firewalls:

13. Add a rule to the top of the rulebase. Set:
 - Source = Partner-Net and Internal-Networks
 - Destination = Partner-Net and Internal-Networks
 - Service = Any
 - Action = Encrypt
 - Track = Long
14. Right click on the Encrypt icon in the Action column to set the encryption properties for the rule.
15. Click the Edit button and set the encryption properties defined earlier.
16. Click OK twice to exit the dialogs.
17. Verify, install, and test the policy.

Figure 4 shows the VPN dialog for the firewall object.

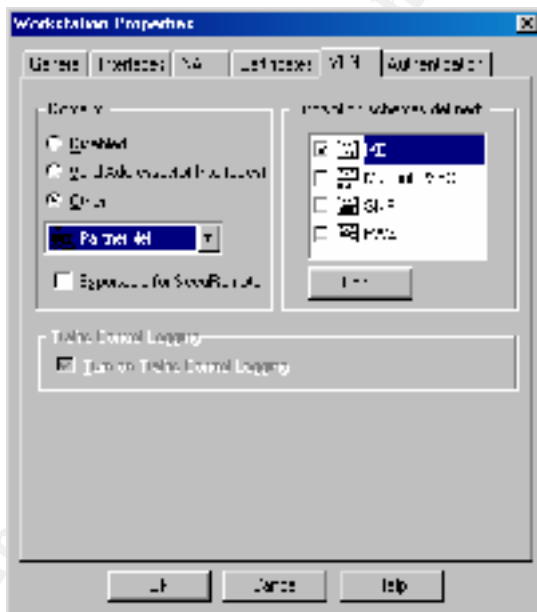


Figure 4. Configuring the firewall for a VPN.

3. Audit

Specification:

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

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

Audit Design:

The purpose of auditing an organization's security is two fold. The first purpose is to ensure that the organization is in compliance with it's security policy. The second is to "verify the existing level of protection and make recommendations to improve it." [Wendt 2001]

A complete security audit has two parts: 1. Use software and hardware tools to test the security architecture for vulnerabilities and compliance with the security policy. 2. Review records and interview employees to make sure that the procedural components of the security policy are being followed.

A basic vulnerability test should be conducted on a monthly basis. With some up-front work this test can be mostly automated, minimizing the time and expense to help ensure that the test is performed as scheduled. Every 6 months a full audit should be conducted. Both of these can be performed by GIAC Enterprises employees and should take less than a day to complete. If GIAC does not employ a full time auditor it is important that a full audit be conducted by a qualified manager who is not involved with daily security practices. Having someone involved with the audit who is not too close to the process helps to keep the audit objective and reduces the risk that shortcuts will be taken.

On an annual basis an audit should be performed by a third party who specializes in security practices and vulnerability assessments. This type of audit will probably take several days. References from satisfied clients should be obtained before hiring a security consulting firm. This company will have the resources to keep up on new vulnerabilities and intrusion techniques and will be prepared to perform scans from within your network and from the Internet.

For all audits and vulnerability scans written permission of management should be obtained.

Vulnerability scans should take place during non-business hours. In an information business on the Internet all hours can be considered business hours, therefore scans should be conducted during low traffic periods. System administrators should be available during all scans to react to any potential problems that may be caused by the tests. Review of security records and employee interviews can be conducted during normal business hours.

Vulnerability Analysis:

There are many tools available to assist in vulnerability testing. Our preference will be tools that can be automated so that with a little planning and time invested up front, doing a periodic vulnerability analysis

can be accomplished with minimal effort. For this vulnerability assessment we will use the ping command, telnet command, the Nessus security scanner, and the tcpdump command.

Ping Tests:

The ping command can be used to test the responsiveness of each firewall's interfaces. This test can be conducted from each IDS system to the local firewall interfaces in our security architecture. For example figure 5 demonstrates a ping test from the IDS system on our public network. This confirms that rule number 1 on our firewall is being enforced. ICMP echo requests directly to the firewall interface are being silently dropped.




Figure 5. Sample ping test.

Telnet Tests:

Many people do not realize that telnet can be used to connect to any TCP port on a system. This is useful for debugging web servers, mail servers, and many other application that communicate with text based messages. From the IDS server on our public network we can confirm that we have HTTP access to web.giac.com. Figure 6 demonstrates how to use the telnet command to test access to a web server.

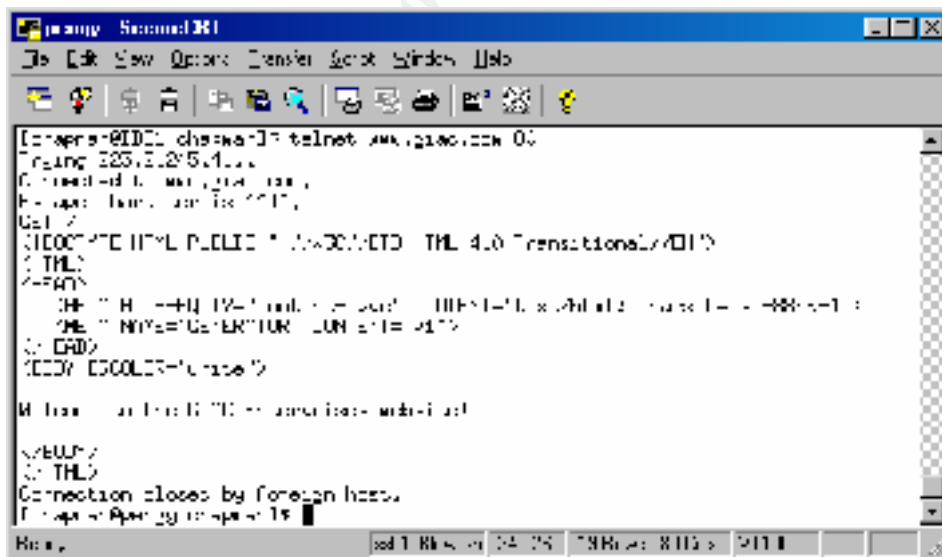


Figure 6. Using telnet to test a TCP service.

From our service network we can confirm that our dns1 server does not have access to our internal SMTP server. The SMTP port is 25. Figure 7 illustrates this test.




Figure 7. Failed telnet access to a TCP service.

Nessus:

Nessus (<http://www.nessus.org/>) is an open source tool capable of vulnerability testing and port scanning, and can produce nicely formatted reports to help communicate it's findings to management. Nessus uses a client/server architecture. There are several Nessus clients available including a command line client that can be automated through the Unix cron facility. While some might question the quality of an open source security scanner, a recent Network Computing review (<http://www.nwc.com/1201/1201flb1.html>) found Nessus to be the most complete security scanner in it's software category. In fact it is not unusual for open source tools to be updated for new vulnerability tests more quickly than their commercial competitors.

A computer with Nessus installed should be temporarily connected to the switch on each GIAC Enterprises network. The entire network should be scanned for active IP addresses, open ports, and vulnerabilities of available services.

Figure 8 illustrates Nessus configuration in action. Nessus has a plug-in architecture that allows new tests to easily be developed and added. Here we see that all plug-ins have been enabled except for ones that might cause a server to crash. The first window displays the plug-ins and the second window displays options for the currently selected plug-in.

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

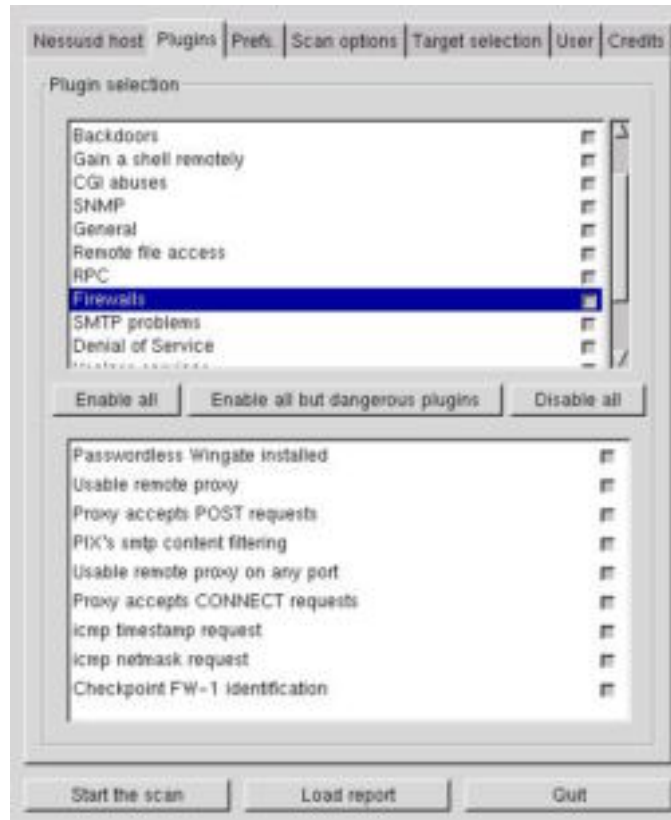


Figure 8. Nessus Plug-in Configuration.

Figure 9 shows the results windows for a scan against one host. This host is not on the GIAC Enterprises network. Security holes are show in red, warnings are shown in orange, and notes are shown in green.

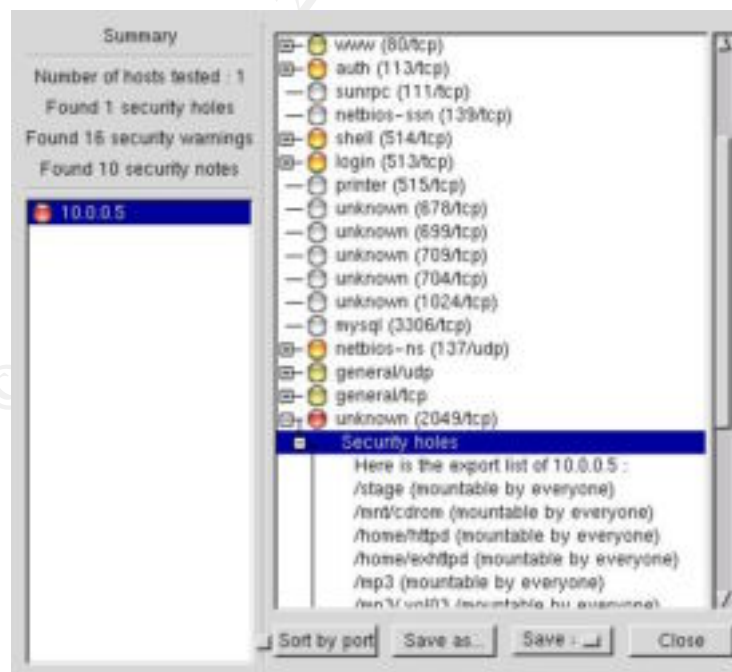


Figure 9. Nessus scan results.

The results of the Nessus scan should be compared against the security policy. Each vulnerability found should be addressed by the security team and fixed as soon as possible. Security notes should be reviewed and potential risks addressed. If a service is found that is not part of the security policy, then the service should be disabled or the security policy should be updated.

Tcpdump Test:

Using the tcpdump command from a workstation temporarily connected to the switch on the GIAC Enterprises public network we can inspect the contents of IP packets routed in and out of the GIAC network. Packets associated with a VPN connection should not carry any unencrypted data. (Note: The switch must support port mirroring and be configured correctly to see the traffic on the network analyzer's port.) The syntax of this command would be:

```
# tcpdump -x -s 1500 -w dump.out host <GIAC firewall addr> and host <partner  
firewall addr>
```

The `-x` option prints the packet in hexadecimal and the `-s` option indicated how many bytes of the packet to print. The `-w` option specifies a file to save the output to in raw format. A program such as Ethereal (<http://www.ethereal.com/>) can be used to read the output file and do ASCII conversions on the hexadecimal data. The data should be un-intelligible due to the encryption. Ethereal can also be used to do the network sniffing directly.

Interview and Reviews of Procedure:

The bi-annual audit should include employee interviews. The purpose of the interview is to determine if the employee fully understands and follows the security policy and procedures. Failure of an employee to correctly follow the proper procedure should be addressed with additional training.

Also, the configuration of all security systems (firewalls, routers, servers) should be compared against the security policy document. Any differences in the specification of the security policy and the implementation should be addressed. This will help prevent temporary changes in configuration of a firewall or other equipment from becoming permanent.

© SANS Institute 2000 - 2002
Authentic Training Solutions
Author retains full rights.

4. Design Under Fire

Specification:

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

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

Attack Target:

The target for this attack is shown in the figure 10 below. This security architecture is from the GCFW practical located at http://www.sans.org/y2k/practical/James_McMahon_gcfw.doc . The reason this security architecture was chosen is the lack of any firewalls behind the primary firewall. If the primary firewall can be compromised many systems will be exposed for further attack.




Figure 10. Attack target.

Attack 1:

The firewall in the security architecture is a Firewall-1 4.0 system. There are 13 security alerts listed for this firewall at <http://www.securityfocus.com/> . Of those 13 vulnerabilities none of them allow compromise of the firewall itself with the installed rulebase. The best chance for attacking the firewall is a denial of service attack. Perhaps if GIAC Enterprises loses critical business functionality due to the attack, changes will be made to the rulebase which might expose other vulnerabilities.

Firewall-1 can be made to use 100% of available processor time logging illegally fragmented packets. A program written in C capable of generating this attack can be obtained at <http://www.securityfocus.com/data/vulnerabilities/exploits/jolt2.c>. The firewall must be upgraded to Firewall-1 4.1 service pack 2 to eliminate this vulnerability. A temporary workaround is to disable console logging.

Attack 2:

The second attack is a denial of service attack against the web server and possibly the firewall itself. By using systems on the Internet to flood the web server with connection requests we can cause the web servers IP stack to consume all of its resources and prevent it from servicing valid requests. This is done with the tool Neptune which can be downloaded from <http://www.niksula.cs.hut.fi/~dforsber/synflood/programs/SYNpacket.tgz>.

Neptune send SYN packets (TCP connection requests) to the web server with a forged source address that does not exist but whose network address is routable. The web server will respond with a SYN/ACK packet but since the forged address does not exist the TCP 3-way handshake will never be completed. The queue in the web server's network stack will fill up with half open requests until it can no longer handle additional requests. Eventually each request will timeout and be cleared from the queue, but with a steady flood the denial of service can be maintained.

Firewall-1 has a feature called SYNDefender which is designed to prevent this type of attack, but it is turned off by default. When SYNDefender is turned on it can run in two modes; SYN Gateway and Passive SYN Gateway.

SYN Gateway is the first of the two options. In this mode Firewall-1 accepts the SYN packet in its own queue and immediately sends a SYN/ACK to the forged source. Eventually the connection will timeout and Firewall-1 will clear the connection from its queue having never sent a packet to the web server. The problem with this mode is that Firewall-1 can then become the target for the denial of service attack as its resources can be consumed with half open connections.

With Passive SYN Gateway Firewall-1 forwards the SYN packets to the web server but tracks the open connections. If the SYN/ACK isn't seen by the firewall within the configured timeout, or the number of half open connection exceeds the configured limit, Firewall-1 will send a RST packet to the web server, terminating the connection. The SYNDefender properties can be seen in the figure 11 below.



Attack 3:

The existing firewall rulebase offers little opportunity to attack internal systems. The best chances for compromising an internal system are the external SMTP and external DNS servers on the service network.

The SMTP server will accept TCP connections from any host on the Internet. By telnetting to it's service port (25) we can probably determine which SMTP server it is running. We can then use know buffer overflows or other exploits to spawn a root shell on the machine on the SMTP port. If the firewall isn't running any special SMTP filters it will not know that future connections are to a shell because it's based on stateful inspection technology as opposed to proxy or application gateway technology.

If attacks against the SMTP service don't work then DNS may possibly be vulnerable. TCP based DNS connections are only accepted from secondary DNS server. Using nslookup we can determine which secondary DNS servers are located on the Internet. One of these servers might not be protected by a firewall. Nessus could then be used to identify vulnerabilities in the secondary DNS server and attempt a compromise. If successful the GIAC Enterprises internal DNS server can be attacked from this host. It might be the case that recent BIND vulnerabilities may not have been patched because the DNS server is behind a firewall and considered safe.

Supplemental Attack:

Perhaps the simplest denial of service attack against an enterprise requires no attack against the security architecture itself. Some domain name registrars such as Network Solutions (<http://www.networksolutions.com/>) allow a domain name registrant to choose from multiple security methods to prevent unauthorized changes to their domain settings. Many registrants choose the simplest method which is an e-mail form with the requested changes.

With one forged e-mail an attacker could change the DNS entries for the target domain. The attacker could set these values to those of other compromised systems and by changing some or all of the domain's DNS information, direct connections to another site. These systems may have trojan programs on them that request usernames and passwords. The attacker then has authentication information which could be used to break into the target domain.

Sometimes the simplest attack may be the most effective.

© SANS Institute 2000 - 2002

References and Resources

- [Brenton 2001] Brenton, Chris, "VPNs and Remote Access"
January 2001, SANS New Orleans conference.
- [Northcutt 2001] Northcutt, Stephen, "TCP/IP for Firewalls and Intrusion Detection"
January 2001, SANS New Orleans conference.
- [Spitzner 2001] Spitzner, Lance, "Advanced Perimeter Detection and Defense In-Depth"
January 2001, SANS New Orleans conference.
- [Stevens 1994] Stevens, Richard, "TCP/IP Illustrated, Volume 1"
1994, Addison-Wesley
- [Wendt 2001] Wendt, Carla, "Network, Perimeter, and System Audit Review"
February 2001, SANS New Orleans conference.
- [Zwicky 2000] Zwicky, Cooper, Chapman, "Building Internet Firewalls, 2nd ed."
June 2000, O'Reilly and Associates

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