Global Information Assurance Certification Paper

Don’t give a mouse a cookie....
Protecting Fortunes into the Future

Submission for GIAC Certified Firewall Analyst, version 3.0

Meg Layton
May 18, 2004

Abstract

GIAC Enterprises is a business which deals in the online sale of fortune cookie sayings. This is a tiered business, dealing almost entirely through remote/mobile workers, and is dependant on its secure infrastructure to keep itself competitive in the market. The industry is highly competitive, and GIAC Enterprises has been the leader for a number of years. It has remained the industry leader because it is willing to evaluate and implement new technologies, and it recognizes the value of security within its operations. Recently, the I.T. Director vacated the position within the company. The new I.T. Director could not find any documentation regarding the network, and undertook the task of documenting the current structure, investigating the reasons behind it, and understanding the current environment of the company. Proper documentation of a network is critical to being able to provide recommendations for future upgrades, meet scaling requirements, and adhere to the proper security posture. This paper will explain the current security architecture, and illustrate the security policy and component configuration. This will become a critical documentation of the current infrastructure. When appropriate, recommendations may be made for upgrades/issues that the company should look to address in the near future. An analysis of a competitor's infrastructure will be conducted in part 3, and the paper will end with a recommendation for Correlation Technology given the current infrastructure of GIAC Enterprises.

Part 1: Security Architecture

"If you give a mouse a cookie he will probably want a glass of milk…" -Joffe

Business Requirements

The main business requirement for GIAC is database uptime. When the web interface to the database or the database itself goes down, the company is virtually crippled, as it is in the business of offering the data contained in the database for sale. In addition, the security of the database and the data therein is critical. If the fortunes contained were leaked to a competitor, they would not be a marketable asset, and likewise customer information could be extremely valuable to a competitor. There is no hard physical product which is being sold, this company is selling data. The security of the data is therefore the most critical to the business continuing its operations. The company has been in business a long time, and has invested much in its infrastructure. When it was founded in early 1990, the world thought Windows was the best thing. As most businesses did, new computers were purchased with an expected depreciation of about six to ten years. Therefore, by the time new operating systems were available in the common marketplace, several of the enterprise operating procedures were already entrenched in the Microsoft genre, and redesigning or

replacing these procedures in favor of more secure ones has been slow. In some cases, the company was able to implement a more secure offering because when the business process was ready to change, so was the platform. Since the company's inception, it has taken a more pragmatic view of the computers, expecting them to depreciate at the three to five year mark. This has made newer technology available for integration at a more rapid rate over the past few years.

The company was founded by an independently wealthy man with an idea, and at the outset, business was difficult. Orders were received via fax, and GIAC needed to outsource the printing of fortunes and have actual paper fortunes delivered to the end user. Through the years, the business model has changed to acknowledge their strengths in the fortunes themselves and to take advantage of the changing technology. They are now a strictly e-enabled business, taking orders over the Internet and fulfilling orders through e-mailed data. This has reduced the costs associated with shipping and creation of hard-copy of fortunes, and they have reinvested this money to focus on their customer care. By ensuring their data is secure, the customer data is secure, and the customer gets the fortunes requested, they have built their enterprise into the leading fortune creator for cookies. When the company started dealing with the Internet regularly, they selected an ISP to host their web page as they did not have the infrastructure or the knowledge at the time to secure their public-facing site. Now, they continue to host their public-facing site with the ISP because they have chosen to focus their infrastructure on a customer-facing track. Their public-facing site does not change much, and due to their long-time relationship with the ISP, there is little cost to hosting the web page with the ISP. They are comfortable with the security of this static information being, essentially, out of their control.

In order to understand the security and network structure, it is necessary to understand how the company does business, and what the requirements are. The current security posture of the company is as follows:

- The company acknowledges three important factors influencing its security posture – the technology available today, the business requirements, and the people within the organization. It strives to unite these three factors for an overall security posture that is unmatched in the industry.
- Security is an issue that affects everyone, therefore every employee is security-aware and responsible for the security of the company's assets.
- I.T. performs routine audits and applies security components as needed to corporate systems accessed by any of the identified end-users.
- Security is not a one-time issue to be addressed. New vulnerabilities are always being discovered, and in order to remain secure and competitive, the current infrastructure must constantly be assessed. A single system, a single vendor, a single solution will not be found to address all of the needs. Network design must take into account security from all aspects.

- Tools are routinely assessed to ensure they are operating in the role in which they were intended. These tools include anti-virus, personal firewall, VPN client, token-authentication, intrusion detection, content filtering, and vulnerability management.
- Security does not end at the perimeter. Because of this, the network is divided into a "protected network" and an "internal network" – which is additionally protected from the outside threat. In addition to a perimeter firewall, client firewalls are implemented throughout the organization, integrated with anti-virus to ensure in-depth protection against blended threats.

> "Blended threats combine the characteristics of viruses, worms, Trojan Horses, and malicious code with server and Internet vulnerabilities to initiate, transmit, and spread an attack. By using multiple methods and techniques, blended threats can rapidly spread and cause widespread damage..... Effective protection from blended threats requires a comprehensive security solution that contains multiple layers of defense and response mechanisms."
> (Symantec AV Center)

Strict policies are in place and evaluated to control access points to the network. These policies include dictating hardware & software configurations of machines that have access to the internal network.

To address all of these requirements, security within the organization is addressed with a defense-in-depth approach that is designed with the following layers:

- Physical security: security of the offices and equipment
- IP Addressing scheme through the router
- Firewalls, VPNs, address translations
- DMZ hosts with Anti-virus, content filtering and additional services like http, FTP, LDAP. NIDS sensor.
- Internal firewall with address translation
- LAN: NIDS, HIDS, anti-virus on the host, client security
- Knowledge security: training and re-training of personnel and customers to ensure the infrastructure remains secure

These layers will be discussed more throughout this document.

Identified end-users

Customers

The general customer structure is as follows. Companies purchase fortunes bulk online. They do this through a secure ordering system. There are two class of customers:

1) cookie manufacturers which purchase large bulk orders of fortunes for their cookies. A certain guarantee goes into the large bulk orders – each saying may only be duplicated three times within a single large bulk order, and each customer will only receive the saying a maximum of six times

within a given 12 month span. These customers are the makers and marketers of mass-produced cookies, for restaurants and large markets. The common form of payment are purchase order numbers.

2) small specialty shop and family business orders. Because of the size and specialized clientele, the specialty shop market requires no duplicates within a given 12 month span. These customers are the makers and marketers of home-made or made-to-order cookies. These customers submit either purchase orders or company credit card information through these secure on-line system.

Requirements of the customers: Email, web connection, online sales. Protocols to be implemented: VPN, smtp, http, https.

Traveling Sales

By virtue of their environments, each market has its own space with its own restrictions. The specialty shop owners sometimes deal in specialized fortunes (for special occasions, theme fortunes, etc.) and require a specialized sales force targeted to their needs. This mobile sales force team reports to the Vice-President, Small Business sales. Each sales team member typically has about five accounts which he/she is responsible for, and a requirement to target three new accounts a year. The primary focus of this team is to build/maintain customer relations.

The Vice-President, Bulk sales operates a sales-force with a decidedly different focus. Because the focus of these customers is the quality of the data and assurance that the fortunes meet quality standards, this team has a larger account base for which it is responsible and spends a lot of time on data control. They are also responsible for targeting/acquiring new customers.

Requirements of the traveling sales: Email, web connection, access to online sales. Protocols to be implemented: VPN, smtp, http, https.

Contractors

All of the fortunes which the sales team market are supplied by individuals writing and researching their fortunes. The Vice-President of Supply ensures that these individuals are providing quality, non-repetitive fortunes for resale. Because of the nature of the business, there is a team of about 1000 individuals writing and researching the fortunes. Some of this team specializes in “quirky” fortunes, or fortunes targeted more at the Small Business unit, and even in “customizable” fortunes. When fortunes are submitted, they must go through consistency and duplicate checks before being accepted as a “marketable” fortune. In addition, fortunes are categorized to ensure that quick and easy references exist in the

database to ensure the proper fortunes are sold to the proper customer for maximum satisfaction.

Requirements: Email, File transfer, database access (protected, not internal). This indicates the need for a fixed IP to make use of VPN technology. Protocols that address requirements: VPN, encrypted smtp, file encryption.

Employees

GIAC retains two internal offices: GIAC North America and GIAC Europe. This ensures that they are the most visible fortune reseller in the world, with their fortunes being translated into multiple languages. Internal operations include Human Resources, Technical support, Database management, Financial Resources, Operations Management. Employees sitting in the office must have access to the web, access to the internal database, access to e-mail and there is the occasional requirement for additional software. (For example, Human Resources requires its own database and has been using Microsoft Access for many years.)

In order for in-office employees to communicate/share data with traveling employees/contractors, a file share/ftp server has been set up in the DMZ. Files are transferred back and forth, everything on this server is encrypted in order to preserve the company's confidential data. For this file encryption and for the transfer of fortunes through e-mail, PGP has been implemented on a company-wide basis.

Fortunes are received from the writers into the "received database" daily. These go through the duplication checks, categorization assurances, and translations and when they have been approved as marketable fortunes are placed in the SALE database. This coincides with the CUSTOMER database which tracks which customers have been sold which fortunes.

The finance office has chosen the financial software is SunSystems by SystemsUnion to support their general ledger and accounting needs. This financial data warehouse allows the two offices to share a single ledger structure. This data is resident on the internal network and is not accessible from the protected network.

Employees can, at any time, either be physically located on the corporate Local Area Network (LAN), or be traveling and require remote access to certain components of the infrastructure. The remote employees require a method of connection that is secure and fast, therefore a VPN solution is provided through which users access the network.

In a highly competitive field such as fortunes, the problem of the corrupt insider needs to be addresser. Security experts warn that the insider threat is a tricky

one to measure, but an increasing threat to companies. According to an article in ComputerWorld, "Overall, the number of computer-related offenses committed by insiders continues to rise rapidly each year." (Schweitzer) Therefore, careful hiring practices have been implemented and there is a special effort to keep employees satisfied at their jobs, give them ample career advancement, and an opportunity to take ownership in the company through their own individual contributions. When the company succeeds, everyone succeeds, and the employees are well-rewarded.

Requirements: Email, Web browsing, File transfer, database access (internal when on the LAN, protected when connecting remotely). This indicates the need for a fixed IP to make use of VPN technology. Protocols that address requirements: VPN, file encryption, HTTPS

General Public

The general public does not have access to the GIAC site, the internal web server is for approved partners, employees, or suppliers only. The general public only needs to access an informative site with information on how to contact, etc. Therefore, people requesting such general information are connected to the public web site. For this reason, the public web site is hosted by the ISP and does not reside on the protected networks of GIAC Enterprises. This also achieves some safety: whois has this name/webpage listed with the technical contacts of the ISP, since they were the ones that originally registered the domain name. This is the accurate information for attacks that may originate from within the ISP infrastructure and associated with this domain name. For additional outgoing, the domain name gets extended for the service utilizing it, i.e. mail is sent to MailGIACEnterprise.com. Additional functions (not mail functions) that use the information contained within the registrar database for this domain name are flagged as suspicious, or not even forwarded into the private domain – the preferred method of access is IP address. Anything listed in the registrar database contains no user names, but does contain an accurate e-mail that is hosted by the ISP which is regularly checked by IT staff.

Additional requirements/restrictions

Fortune lifecycle

Based on the requirements of the different factions listed above, GIAC Enterprises depends on a distributed database architecture. The database which the fortunes are received is a dynamic database, with many read/writes a day. This database exists in the protected network with the Web Servers, and is accessible through the VPN. The fortunes are gathered four times a day, go

through checks as mentioned above, and are transferred in batches to the internal database, which resides on the more secure internal network. This is the database which contains customer information and the "marketable" fortunes. Once the data is transferred, those fortunes are no longer available from the protected network.

Orders are received through our web interface. These orders are placed in the order database and are transferred four times a day to the internal database for fulfillment. Such a separation ensures the protection of data from SQL injection (an attempt to inject a SQL query/command as an input) and CGI scripting vulnerabilities. In addition, orders at either location are transferred both to their local internal database AND to the second office's internal database through an asynchronous replication method. (Although replication is not an accurate term, because in this usage everything is one way. No data comes back to the external order database or external fortune database to be stored.)

After the transfer of the data to the internal database, the data is immediately deleted from the protected network. This ensures that in the event of perimeter compromise, under six hours of data is available on the database. Attackers that succeeded in cracking the perimeter would not have write access to the protected internal database. Indeed, if an attack is reported as successful on the fortune database in the protected network, fortunes that were resident on this database are discarded, ensuring complete customer satisfaction that fortunes from GIAC are beyond compare in this market.

Orders are filled through a batch fulfillment process and then e-mailed to the end user in an encrypted file which the customer would unencrypt with their private key. Based on this process, mail security is also a concern.

Sales people and customer have designated "customer handlers" that work within the offices of GIAC. These handlers are the ones who can answer questions on order status, deliver information on specific fortune requests, etc. They frequently transfer files to the servers in the DMZ to share information with those on the road, including sales status reports, progressive customer reports, and target customer information. Transferring all of this data in a secure way is critical to maintaining a competitive edge.

Vendor Alliances

Aside from the early dependency on Microsoft, two vendors repeatedly appear in the network configuration: Symantec and Dell. Symantec has been providing anti-virus support since the Norton days, and as they expanded into other security technologies were a natural choice for the administrators, who already had a level of comfort with the support and the products. Dell has also been in place for a long time as the hardware vendor of choice, due to the experience

this company has had with the hardware support and best choice for low price. As long-time customers of both vendors, GIAC experiences special price considerations and has appropriate service levels from both vendors to meet their needs.

Basic Implementation Aspects: Hardware/Software/Placement

The GIAC security posture dictates defense-in-depth. With this in mind, a variety of services and vendors has been implemented – this methodology helps to offer complete protection so vulnerabilities on certain products/vendors don't propagate over the entire network.

A word about file encryption

Where encryption is required, each employee, each contractor, and each customer uses a PGP Key. This offers confidentiality and integrity to data, as well as authenticity and non-repudiation for the e-mail communications. Mail clients are equipped with PGPNotes, a plugin that allows Lotus Notes users to encrypt attachments from within their Notes client. Keys are stored on the company's public key ring. Customers add to their own keyring the office staff and their personal "handlers" and sales point of contact, while the office key ring is more complicated, needing to track all of the contractors and all of the customers as well as the employees. Contractors, on the other hand, only require the office staff on their key ring. New customers are provided with the key rings they require when they place their first order. PGP has been in use since the company first started transferring data, and is updated to the latest version – because it has been in use for such a long time, it is a trusted and integral part of securing the company's data. Great care is taken to remove keys for employees that no longer work at the company, and to ensure the keyrings are always accurate.

Physical security

While physical security is not usually a device that can be called out in a network diagram, it is an integral part of the security infrastructure of any network. In the case of GIAC Enterprises, when they established their U.S. presence a few years ago, physical security was at the forefront of their minds. The competitors are always looking for a new way of finding out their secrets, and therefore every employee is taught the basics to ensure the security of their data. The following precautions were taken with both offices:

- Card keys on an alarmed system: employees are required to swipe their cards for access beyond the reception area. Access to the server rooms is carefully restricted to authorized server personnel. Employee access

and alarms are reported to an alarm log server, which is manually copied to a file and “sneakernetted” across the room and entered on the syslog server. (Integration to the network ought to be considered, with a reminder that the alarm keypads and keyreaders will then be another entry point into the network and additional security measures may be required.)

- Doors have been installed to resist forcible entry
- Walls have required fire rating
- Cleaning crews are bonded, ensured and the same two people have been cleaning the offices since it opened.
- Server rooms are equipped with tiled anti-static flooring, above-head cabling racks, and exclusive power drops
- Equipment in server rooms is rack mounted
- Paper shredder is in use to shred critical documents
- Dedicated temperature control/AC units for the server room
- All equipment is connected to dedicated an APC UPS (appropriate sizing is considered) so that power fluctuations/brown outs do not affect the equipment
- Appropriate fire precautions have been taken. Since there are two offices and Europe/America have slightly different requirements/restrictions, the power and fire layouts are beyond the scope of this document.
- Server rooms are on the bottom floor, at the end of the hallway from the kitchen/rest rooms (as illustrated). This placement is intentional for two reasons: a) the ceiling had been known to leak in the event of heavy rains. Therefore, placement of the rooms on the second floor was not an option. b) to eliminate possible hazards as a result of piping issues, the goal was to separate it as much as possible from the main pipes in the building.

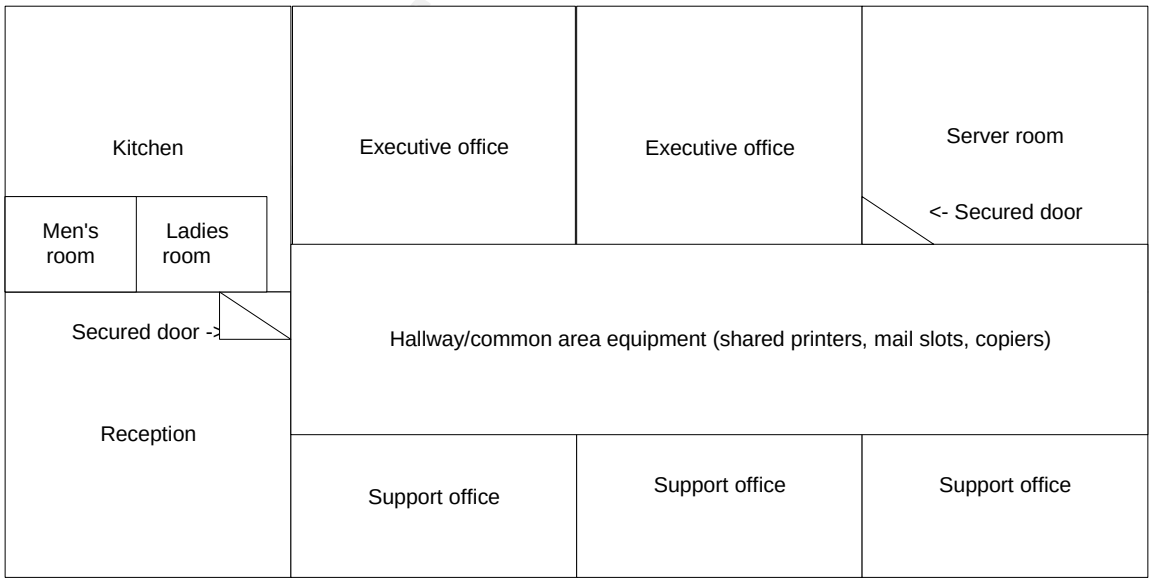

Office layout on first floor (in the interest of space additional floors/offices not depicted here.)

Security Training

The best security policy and implementation in the world will not help a company that does not train its employees on what to look for, how to react to potential security vulnerabilities within the system. The following measures have been incorporated into a security-awareness program that is given to every employee upon hire, and reinforced with practical training, e-mail notifications, and "Security Aware – Security Star" awards that are given to employees every quarter.

- ✓ Know your system. If you see something suspicious: STOP what you are doing. RECORD what you see on the screen and NOTIFY I.T. personnel. Something suspicious may be that your username is not the one that comes up when you get back from lunch, it may be an error message you have not seen, or it may be new icons in your taskbar. An ineffective recording is "something strange happened." An effective recording is "I got this error message right after I logged in."
- ✓ Take note of the people around you.
- ✓ Every employee is required to maintain a telephone log of who called and the business nature of the conversation, even if it is just "Joe from I.T. changing your password"
- ✓ I.T. has a rotating passphrase that will be listed in the opening line of authenticated e-mail that they issue. By authenticating the passphrase employees will be able to tell if the instructions contained within are valid – this is in addition to the key rotation which would further validate the e-mail and its contents
- ✓ I.T. releases concise instructions, and keeps an up-to-date FAQ that helps employees with the basics
- ✓ Organizational policies are enforced and re-educated often, and policy on web postings (newsgroups, mailing lists, etc) is reviewed and enforced.

Employees are rewarded: if they report a security event not covered in a FAQ, the event and the person is noted as a Security Star, and every employee reporting a valid event goes into a pool for drawing gift certificates each quarter. The theory is that valid security events either lead to updated FAQs if it is something worth noting, or lead to enhanced security/early incident response if it indeed is indicative of a compromise. No employee should be afraid of reporting something strange, or be so caught up in their tasks of the moment to miss the sign of a potential attack.

Basic server Best Practices

In all of the server components, the following best practices are applied to harden the systems: sample files have been removed, unneeded services have been

turned off, and servers have been inventoried with a digital hash of every file where appropriate.

The backup strategy for critical machines is as follows:

- Full backup each night (ArcServe Backup Release 11 for Windows is the primary software in place. It has been fully patched to the latest release. Where appropriate, releases for the additional O/S's have been implemented.)
- Friday backups stored off site (safety deposit box)
- Monthly rotations of set of base tapes to ensure no degradation of tapes
- Quarterly restore of data tests to ensure that data can be restored in timely and efficient manner
- Where applicable, lockdown tools have been applied.

Log files are kept separate from the servers (in order to maintain logging in the event of the compromise of a separate server). Log files are sent to a syslog server and reviewed weekly looking for possible attack indicators. (Correlation technology, which is addressed in part 4 of this paper, could help with a faster assessment of the network and would span the multiple vendors implemented here.)

Banners identifying O/S and software types have been disabled, and sign-on warning banners have been posted at all logon points where it was technically feasible. These banners warn against unauthorized use and informs the users on the possibility of monitoring of their transactions.

Where technically feasible, RAID arrays level 5 have been implemented. RAID (Redundant Array of Independent Disks) makes use of multiple drives for fault tolerance and performance improvements. Level 5 ensures data and parity information is spread across multiple disks (this process is known as striping). "Fault tolerance is maintained by ensuring that the parity information for any given block of data is placed on a drive separate from those used to store the data itself." (PC Guide)

Additional redundancy considerations

Wherever possible, the network has been enabled to be as redundant as possible. For instance, should one router interface fail, the router to the other office should pick up the traffic and route externally through this interface. All the firewalls have been implemented with a spare interface for replacements, and you will see in the IP address table reserved addresses for future redundant implementations or replacements as needed. Each office has in stock a spare hub that allows for instant replacement of the hardware, and in addition the

database is configured in such a way that data loss is minimal should a server fail.

The system is also implemented in such a way as to allow for future growth. The IP addressing table allows for additional security devices such as additional firewalls or routers, should that become necessary. The IP addresses show that the networks and subnets have been implemented with liberal growth opportunities to allow for additional servers or databases, perhaps the eventual movement of the static web site in-house.

Address ranges in use

1. 198.100.1.0/25 and external routable IP addresses: This block of addresses was allocated to GIAC for its public presence. Because GIAC only has a few external addresses, below 198.100.1.10 is America, and above is Europe. This allows plenty of room for growth should their needs change.
2. 192.168.113.0/25 and 192.168.114.0/25 internal side of network between two firewalls (DMZ) Uses class C private addresses based on RFC 1918. Purposely uses oddly numbered addresses to allow for merging with other entities that may use more commonly numbered addresses.
3. 10.193.111.x internal network and 10.193.112.x. Uses the class A 10/8 prefix private addresses based on RFC 1918. Purposely uses oddly numbered addresses to allow for merging with other entities that may use more commonly numbered addresses.

198.100.1.0 Network

.1 ISP (America)
.2 internal router face
.3 Backup router
.7 Broadcast
.8 Network
.9 Firewall

.12 ISP (Europe)

.13 Backup router
.17 Broadcast
.18 Network
.19 Firewall

The firewall does redirects for the following:
.20 SMTP Gateway
.22 Internal DNS

.24 FTP server

Since the web server is not carrying anything for public consumption, only the database front-end for connection via VPN, it does not require a routable Internet address.

DMZ

192.168.113.0/25 and 192.168.114.0/25
192.168.11x.7 Broadcast
192.168.11x.8 Network
192.168.11x.1 DMZ side of external firewall
.2 DMZ side internal firewall
.3 & .4 router addresses for office-to-office T1
.5 router loopback
.6 router
192.168.11x.10 and above servers/hosts as follows
.10 SMTP
.11 Syslog
.12 HTTPS
.13 DNS
.15 FTP
.16 NIDS
.17 Database

Internal Network

10.193.111.x and 10.193.112.x
10.193.11x.2 Internal side of internal firewall
10.193.11x.7 Broadcast
10.193.11x.8 Network
10.193.11x.10-20 servers
.17 Database
10.193.11x.25+ are hosts to 10.193.11x.254 max

Connectivity

Since this is business depends on connectivity, great pains were taken to ensure that their business could grow to meet their changing needs. There are two T1s connected to the router: the first connects to the ISP, and the second is a dedicated link between the offices. While both managed bandwidth and big bandwidth tactics were considered when installing internet connectivity, big bandwidth won out in large part because WAN link pricing dropped at a critical

implementation time, and because the Service Level is guaranteed. The point-to-point connections have proven to be reliable. In order to maintain appropriate levels of service, traffic flow is reviewed monthly. The T1s offer sufficient bandwidth to allow several traveling salespeople and customers to connect at one time, given that the majority remain on about 56kb dial-up in the U.S. and 64kb in Europe. With a T1 providing bandwidth of the commonly accepted 1.544 Mbps, (with overhead for framing at 8Kbps), then the usable bandwidth would indicate around 24 people connected simultaneously. This is seldom the case, and the T1 connections have proven to be quite adequate for the system implementation.

The basic connectivity diagram follows. The DSU/CSU representations included here are actually WAN cards within the border routers, and will be defined with the router. Voice connectivity at both sites is currently handled over a separate link through a PBX. The voice is called out in this paper for two reasons: a) it is an integral part of how the company does business, and b) the continuing convergence of voice and data networks presents new considerations and challenges when evaluating the security of a network. This company should be looking at video integration over their leased line between offices and voice integration into the data network is moving into the main stream. In addition, while seldom considered, voice systems, such as PBXs, offer their own security holes and opportunities for reconnaissance within organizations, and therefore must be considered when considering the security posture of a company.

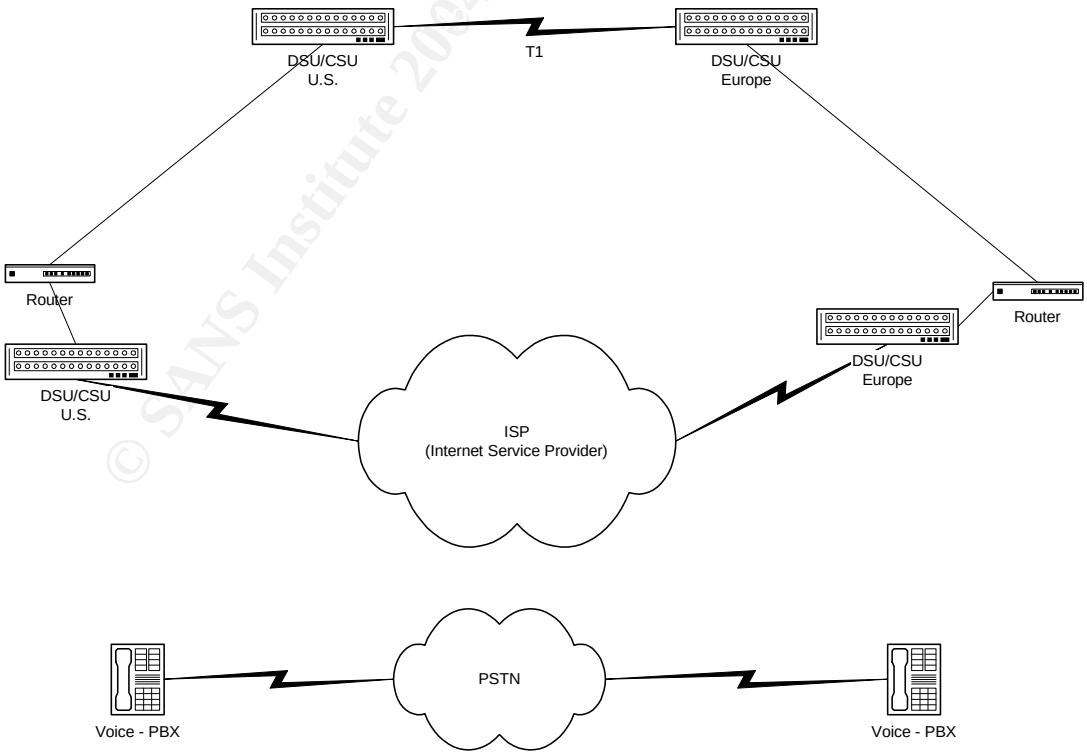

Separating Voice from Data

The PBX is a Telrad Digital 128 with Imagen Voice Mail. This system has an upgrade path to the Digital 400 through the dealer should future business models require it. All of the office employees as well as the traveling/remote office workers and the fortune writers have established voice mail accounts that allows for easy communication between the workers.

The following security measures have been taken on the voice network:

- system passwords have been changed – ensuring someone with knowledge of the system can't log in and perform reconfigurations
- "default" voicemailbox password has been changed so that new accounts are not assigned the out-of-the-box system default password - ensuring someone with knowledge of the system can't log in and perform reconfigurations
- system is audited to ensure initially assigned password is changed within 10 days
- analog ports (required for modems) are configured only to dial out – no incoming calls are permitted on analog ports
- analog ports are renumbered monthly, which ensures that even if someone bypassed the security features already implemented, they would need to repeat the exercise to "find" the port again every month
- modems are only permitted in the finance officer's office, and are used for banking transactions. These workstations are equipped with smart-card readers required by the bank to further secure the transactions.
- the maintenance port cable to the PBX is manually connected on an as-needed basis and the port is password-protected – ensuring remote reconfiguration is performed only as approved
- voice mail and phone extensions are regularly audited to ensure inactive employees are removed
- the default system messages have all been re-recorded on the voice mail to eliminate the possibility of identifying the equipment based on the recordings
- toll restriction tables have been established which restricts dialing to 900 numbers and other known cost associated numbers
- maintenance terminal for the voicemail/pbx is located in a locked cabinet in the server room
- unneeded features, such as Automatic Call Distribution (ACD) and Dial In System Access (DISA), have been disabled for any accounts that do not require it. Class of Service restrictions are applied to ensure that access is granted as needed, and implicitly denied otherwise
- PBX and voice mail logs are routinely monitored for inconsistencies, and especially for continual extension dialing from the same open phone line, which may indicate an attempt at reconnaissance

- The “Directory” feature (where a user can randomly dial a number and see whose extension it is) has been removed. This feature was regarded as a security risk, allowing individuals to do reconnaissance and get a list of company names and in some cases their positions or departments.
- During office hours, the phones are answered by a human being, which allows for a more efficient screening of calls, contributes to the personal touch that the company is known for, and prevents random dialing of extensions

While most of the maintenance on the voice systems is beyond the scope of this document, these steps are indicated as best practices for basic security of data.

Network Diagram

The following network diagram is the dataflow to/from the Internet, and excludes the voice discussion of above as well as the flow of data between offices as called out in the previous diagram. Brand and versions are called out in the following text, for ease of use in viewing the diagram.

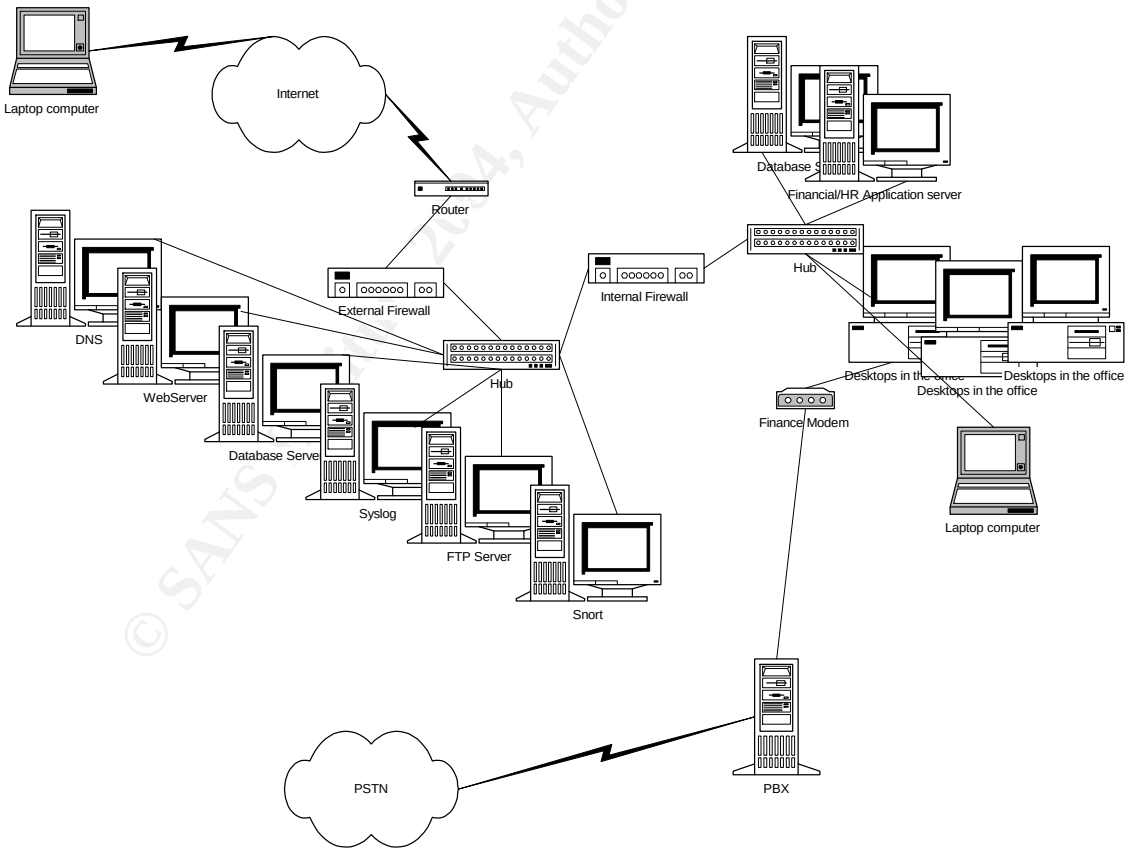

External-facing

This zone is accessible from the Internet to customers ONLY through the web server, which requires authentication, or through the VPN. This is effectively the DMZ. It is not publicized, it has no URL, it is accessible only through IP Address. This means that customers or fortune-contractors that are known to GIAC access this area.

Router
A Cisco 1760 connects GIAC Enterprises to the Internet. This router is equipped with 64 MB Flash SIMM Factory Upgrade. This Cisco is in a pretty standard configuration, including ACLs (Access Control List) that block source address from the internal watch list, and blocks access from multicast source address. The Cisco has been upgraded to the latest IOS release. Anything above 1760 would have been too many ports for the applications needed.

This router functions as a layer of security between the ISP (or other office) and the firewall, and ensures offices can route traffic to the Internet. It's duty is to secure the perimeter of GIAC and enforce restrictions on data packets. In this case, one Ethernet port and dual serial ports: one to ISP and one to inter-office T1. The router can be upgraded in the future, and a simple configuration has been applied – this is to simplify management and ensure that routing problems, a common source of vulnerabilities within the enterprise, are not complex.

However, best practices have been followed and additionally, configuration will reflect the following:

- ✓ Router IOS version up-to-date and patched
- ✓ Documentation of policies on routers exists
- ✓ Router configuration kept off-line
- ✓ Unneeded services disabled
- ✓ Access restrictions imposed through access lists that limit traffic
- ✓ Logging enabled, includes time information, logs sent to server and reviewed often

(These are modified from the NSA Router Security Configuration Guide, Executive Summary Card, Page 2.)

This router has not been configured with NTP (Network Time Protocol) and that is a recommendation for implementation, certainly before correlation technology implementation takes place.

Firewall: Symantec Enterprise Firewall and VPN

When the time arrived to implement a firewall, GIAC turned to a vendor that they were familiar with: Symantec. In addition to the standard firewall implementation, the firewall is used to implement VPN technology. The firewall has two ports: one to the router, and one to the network. It serves as the main protection for all of the internal resources. It can statefully track connections and

policy – like the router, a simple rule set and ingress and egress filtering to limit legitimate traffic has been applied– VPN connection allows full use of firewall to control access.

The system it has been installed on has a hardened OS, and the VPN terminating here allows for effective IDS deployment , all other inbound connections are dropped, with the exception of mail transfers to/from the ISP mail server. The choice for Symantec was clear: a single device, single point of expertise, and the Symantec name.

Web server
HTTP/S access. Uses SSL site to give access to the fortune database when authorized partner, access to the order database when authorized customer. SSL was chosen over other encryptions such as SET and TLF because it is industry-accepted standard and the customers seemed comfortable with it. SSL was compiled as part of monolithic Apache build, so that the organization can control the upgrade paths – applying security patches as necessary and version upgrades only if necessary. (Bahadur & Shema) SSL certificates are owned and read by root only, sample files have been removed, additional unneeded services have been turned off and the compiler removed.

This webserver is technically Internet facing, and it needs to be hardened, both the Operating System and the web server. It was only purchased recently, and so was able to migrate from the Windows platform to a more secure Apache on SuSE 9.x. Apache has proven to be flexible, has a widely installed user base, and it seems to provide bug fixes in a timely manner. All of the web components are reviewed every few month for vulnerabilities. In order to maintain a positive security posture, technical and configuration issues are reviewed because these components are often what leads to applications being insecure.

Database Server
This server houses two separate databases in the DMZ – the fortune database and the order database. This is a Dell PowerEdge Server running IBM DB2. It is a Windows 2000 server with the latest patches applied, and has the following hardware specifications: 4 GB RAM, 120 GB hard drive and a 360 GB RAID array. The RAID array is where the main database applications reside. All of the steps indicated in Server Best Practices have been applied.

E-Mail Server
The E-Mail Server is another fairly recent acquisition. Until recently, e-mail accounts were housed at the ISP, and users had web access to these accounts. The burgeoning growth of their e-business however required a solution that was maintained in-house. The E-Mail Server is another Dell PowerEdge Server with the same hardware specs as above. The E-Mail system employed is Lotus Notes, as the recent deployment of the acquisition enabled the move away from the less secure Microsoft Applications. In addition, the workflow routing of Lotus

was appealing to management. Notes has been patched to the most recent release. In addition, this server is running Symantec's Mail Security for Lotus Notes. This allows anti-spam and anti-virus to be applied to incoming/outgoing e-mail. All of the steps indicated in Server Best Practices have been applied.

Syslog Server

This server is critical as it houses the logs from many of the security devices. It is a high-end Dell Workstation running Windows 2000 with the latest patches applied. It has 360 GB of RAID 5 storage, and is regularly backed up and purged so that logs can continue to be gathered. It is running the Kiwi Syslog version 7.x.x., the licensed version. All of the steps indicated in Server Best Practices have been applied.

DNS Server

The DNS Server in the DMZ is using SSH for administrative processes and has all other services except DNS disabled. The hardware is another Dell Workstation, but this is a fairly newly implemented system and was able to migrate from a Windows platform to a more secure Red Hat Linux running Bind 9.2.3. Checks are made frequently for updates/patches to Bind. All of the steps indicated in Server Best Practices have been applied, and additionally, zone transfers originating from the public Internet have been refused according to the configuration.

FTP Server

This server is critical as it file transfers and critical data to be exchanged between employees, contractors, and customers. It has been separated from other functional servers to eliminate the possibility of hacking into one system and holding the keys to the entire system. Because it has the possibility of storing highly confidential data and data that would be extremely valuable to the competitors, it runs a scheduled purge of files 15 days or older from the system. In addition, employees are encouraged to encrypt the data on the system. Like the other systems, it is a high-end Dell Workstation running Windows 2000 with the latest patches applied. It has 360 GB of RAID 5 storage, and is regularly backed up, and all of the steps indicated in Server Best Practices have been applied.

Hubs

Since management of security devices is already complicated enough, the network is connected with NetGear DS516 hubs. These hubs have been in place for awhile, and they feature lots of ports, they are rack mountable, and require no management. When possible, GIAC has attempted to ensure that they do not overmanage the network, and they do not get management capabilities on parts of the network where management is not required – there is

plenty to manage in any configuration without adding to the headache by including unnecessary management interfaces.

NIDS/AV Server
For the NIDS sensor, Snort was implemented almost a year ago. The decision to choose Snort was easy: it is free and considered high quality. It also guarantees a separation of vendors: while largely a Symantec shop, the company recognizes that a vulnerability in one Symantec program may apply to all programs, especially when they use common methodologies such as LiveUpdate. Therefore, proper implementation of defense-in-depth dictates cross-vendor pollination of the security environment. The machine is Intel based with two interfaces, one on the DMZ segment and one on the internal segment. This is a Dell workstation with 1GB RAM, 40GB hard drive, and Snort is configured to log to the syslog server. In addition to the security provided by a network IDS, additional protection is provided by Symantec Client Security on each host. Snort as freeware is therefore the best option in this environment.

The version of Snort implemented is Snort 2.1.2 for Windows (available here http://www.snort.org/dl/binaries/win32/)

This machine is also a high-end Dell Workstation running Windows 2000 with the latest patches applied. It has 360 GB of RAID 5 storage, and in addition to Snort it runs the server portion of Symantec Client Security. Symantec Client Security has been configured to provide AV Protection across the enterprise and provide basic IDS functions in conjunction with a host-based firewall at the client. The server helps provide centralized management of the clients.

Internal Configuration

Internal Firewall
Cisco Pix 501 running version 6.2. In order to reduce duplication of vulnerabilities, a separate firewall vendor was selected to ensure that if a vulnerability is discovered in the external firewall, the internal firewall would not be vulnerable as well. Cisco is a noted leader in the firewall industry, and some knowledge of the Cisco systems had been implemented at the router. Although this firewall has VPN technology, the VPN technology is not being used, because of the earlier stated implementation of Symantec and also because there should be no inward flow on the internal network except that expressly requested from within. The Pix firewall has two Ethernet ports: one connecting to the DMZ and one connecting internally. Although the Pix is currently implemented as a primary line of defense for the critical interior infrastructure, the rule set is relatively simple – but the Pix lends itself to a scalable architecture, and in the future it may provide a more robust role, such as segregating internal servers and workstations, or further routing to the router.

Database Server
This server marries the two separate databases into usable formats – the sellable fortune database tracks which customers have purchased what, and the purchase order database tracks who is paying for those orders. This is the most critical database and great care is taken to ensure the data is always available and kept confidential. This is a Dell PowerEdge Server running IBM DB2. It is a Windows 2000 server with the latest patches applied, and has the following hardware specifications: 4 GB RAM, 120 GB hard drive and a 360 GB RAID array. The RAID array is where the main database applications reside. All of the steps indicated in Server Best Practices have been applied.

Laptops
Centrally managed personal firewall: Symantec Client Security – disallow everything except IPSec when not connected to the internal LAN. End-user configurations logged and verified at connection before connecting to internal network. Windows 2000 is the OS of choice, in the process of being replaced by XP as the depreciation cycle allows.

Desktops
The hardware of choice is predictably Dell. While some of these systems are slightly older, the majority have been rotated out recently due to the depreciation lifecycle of the machines, and they are currently standardized on the Dell Precision 450 Workstations with 2.4 GHz processors, 512 meg of RAM and 40GB hard drives. The Operating Systems vary depending on when the computers were acquired: most are running Windows 2000 Professional SP4 with all of the latest patches applied, and a select few acquired in the past six months have been installed with Windows XP Professional, again with all of the latest patches applied.

Part 2: Security Policies

"When you give him the milk he'll probably ask you for a straw. When he is finished he'll ask for a napkin." -Joffe

Border Router - Cisco

The border router is using Cisco IOS (Internet Operating System) 12.2 (More information on 12.2 can be found at http://www.cisco.com/univercd/cc/td/doc/product/software/ios122)

The router is the first line of defense into the network, and its main role is like that of a traffic cop: it directs data packets to where they need to go. This direction is based on mapping to internal tables the addressing within the packets and directing packets accordingly. In addition, the router filters some of the traffic, turning away some data traffic and allowing other data traffic to pass through.

While filtering is very important as a first line of defense, it is not a substitute for the more advanced functions of a firewall.

Configuring interfaces that are used

Ethernet
int e0
ip address 192.168.11x .6 255.255.255.0
ip access-group internal-outbound-1 in
ip access-group external-inbound-1out
ip nat inside
no shutdown

Serial
Interface serial 0/0
service-module t1 clock source internal
service-module t1 timeslots 1-24 speed 64
service-module t1 framing esf
service-module t1 linecode b8zs
ip address 198.100.1.1 255.255.255.0
ip access-group external-inbound-1in
ip access-group internal-outbound-1 out
ip access-group
encapsulation ppp
fair-que
no shut

Interface serial 0/0
service-module t1 clock source internal
service-module t1 timeslots 1-24 speed 64
service-module t1 framing esf
service-module t1 linecode b8zs
ip address 192.168.11x*.3 255.255.255.0*
ip access-group internal-gateway in
ip access-group internal-gateway out
encapsulation ppp
fair-que
no shut

Secure the password

The first step with any system is to protect the passwords and login rights. The following commands have been issued(commands appear in *italics*):
service password-encryption

enable secret 5 XXXXXX (where XXXXXX is the password for enabled mode, and the 5 indicates the password is stored in MD5 hash)

Protect the password further by issuing *no enable password*

Establish administrator accounts and hostname

After securing the password, GIAC defines its administrator accounts. Each administrator has their own login so that the log correctly tracks who has changed what. Accounts are reviewed quarterly to ensure that administrator information is up to date.
username meg password 75-bf1at
username meg privilege 1
username joe password j0shua
hostname Mainoffice in Europe this entry reads *Europeoffice*
ip name-server 192.168.11x.13 192.168.11x.13 (where the x represents the European or American digit respective to their locations, and the other entry represents their non-location)
ip domain-lookup
ip domain-name giacfortune.org
The final entries here set up the hostname and DNS entries, as well as the default DNS name. According to the GIAC IP Address table, the same entries are used in the European office as the American office, excepting the second to last octet. In order to backup their DNS entries, the respective countries serve as backup DNS to their counterparts.

Define a banner

As stated in the security policies of GIAC, wherever possible banners are presented at login, an important factor should legal proceedings come into play should the network be compromised. Banner was established through legal.
banner login <banner text> where text is the banner text established by legal.

Define loopback interface

Defining a loopback will allow the router to correctly identify traffic generated by itself.
interface loopback0
ip address 192.168.11x.5 255.255.255.255

Disable other entry points to the router

The following commands should be used connected to the console port, logged into enable mode, start with *config t* and end with *end*
line aux 0
transport input none
login local
exec-timeout 0 1
no exec
The above commands disable the auxiliary port, which GIAC is not using.
line vty 0 4
transport input none
logging local

exec-timeout 0 1
no exec
The above commands disable virtual terminal lines. Remote administration is not absolutely necessary: the routers will be administered by office personnel from within the computer room.

Move administrator services to privileged level

privilege exec level connect
privilege exec level 15 show ip access-lists
privilege exec level 15 show access-lists
privilege exec level 15 show logging
privilege exec level 15 rlogin
privilege exec level 1 show ip
Note: the last line must be the last line, which moves the show command to lower level 1. All the other commands are moved to operate in privileged mode only.

Disable Unnecessary Services

As stated in the security policy, unnecessary services are disabled or denied. The easiest way to ensure this is through denying unnecessary services at the border router. Most of these services, if enabled, can aid in reconnaissance of a network, or present liabilities by leaving open access points to the network. Unless expressly required, unnecessary services should be disabled. Therefore the following services commands have been issued to the router (commands appear in *italics*):

no service config
This command disables network auto-config loading from a TFTP server.

no service finger
This command denies a hacker reconnaissance information available when an individual performs a query on the router. The finger service may provide useful information such as logged on users, or information regarding network configuration. Unless finger is explicitly required, secure networks will have it disabled at their borders.

no service pad (disables all packet assembler/disassembler (PAD) commands)

no service tcp-small-servers and *no service udp-small-servers*
TCP and UDP services used in reconnaissance such as echo or chargen should not be running on the router. The incoming packets will be discarded.

no ip bootp server (disables attackers ability to download IOS)

no ip source-route (disables service which can be used in tunneling attacks or mapping of the network)

no ip identd (disables service which is unprotected against unauthorized queries regarding the host TCP port)

no ip http server and *no ip http server-secure* (disables the ability to use http to manage Cisco routers. See BugTraq 2936 , OSVDB ID 578 or this Cisco Advisory)

no cdp run (this will disable Cisco protocol used to identify other Cisco routers)

no ip domain-lookup router should not be sending DNS name queries

no snmp-server enable traps and *no snmp-server (*disables SNMP to address vulnerabilities. See BugTraq 4132 , OSVDB ID 3664, or this Cisco Advisory)

On each interface

Use *config-if* to configure the interfaces with the following

no ip directed-broadcast (this is the default)

no ip redirects (IP redirects offer an attacker an opportunity to perform reconnaissance, as the system tries to be helpful and offer back alternate routes)

no ip proxy-arp (internal addresses should not be revealed)

no ip unreachables(disables the ability to use ICMP response messages to map a network)

no ip mask-reply(disabled by default, can aid in IP address mapping)

ntp disable (As NTP (Network Time Protocol) has not been implemented, it should be disabled)

Configure Logging (include timestamps)

service timestamps log date msec local show-timezone
logging trap information (identifies severity level for logs sent to the syslog server)
logging 192.168.11x.11 (identifies syslog host)
logging facility local 0 (identifies name of syslog facility storing router messages)
logging source-interface loopback0 (source interface sending the syslog messages)

(These guidelines follow recommendations NSA's best practices page 131)

Access Control Lists

Perimeter router should block local and private addresses from the outside. The ordering of statements should ensure that that unwanted traffic is filtered for legitimate addresses as well as other addresses, so permitted traffic needs to follow the denied traffic.

The following access lists are created:
ip access-list external-inbound-1 Applies to traffic from the external network inbound
no access-list external-inbound-1 to make sure it is clean
deny ip 192.168.113.0 0.0.0.255 any log
deny ip host 198.100.1.1 host 198.100.1.1 log - This will be used to defend against an attack where the router receives a packet with the same IP address in source/destination address and port fields.
deny ip 192.168.114.0 0.0.0.255 any log
deny ip 10.193.111.0 0.0.0.255 any log
deny ip 10.193.112.0 0.0.0.255 any log
deny ip 127.0.0.0 0.255.255.255 any log
deny ip 224.0.0.0 0.0.255.255 any
deny ip host 0.0.0.0 any log
The above commands deny any inbound IP packet that contains addresses from the internal network or local host, as well as multicast and broadcast. This is an anti-spoofing measure. (Refer to RFC 1918)
deny icmp any any echo log
deny icmp any any redirect log
deny icmp any any mask-request log
permit icmp any 192.168.11x.0 0.0.0.255
deny udp any any eq 2049 log
deny udp any any eq 31337 log
deny udp any any range 33400 34400 log
These will effectively block inbound ICMP traffic echo and redirect, which can be used to perform reconnaissance or denial-of-service attacks. It also restricts inbound traceroute for the same reasoning.
deny udp any any eq syslog log
Log udp syslog traffic from outside the network
permit udp any eq 53 192.168.11x.0 0.0.0.255 gt 1023
permit tcp any198.100.1.0 0.0.255.255 established
This should allow TCP traffic requested from the inside to be allowed back
permit udp any 198.100.1.22 eq dns
permit tcp any 198.100.1.20 smtp
allows the Internet to access our systems with proper addresses: the DNS only DNS requests, and mail server only SMTP requests.

ip access-list internal-outbound Applies to traffic from internal (local) network to outbound or to the router
ip access-list internal-outbound-1
no access-list internal-outbound-1
deny tcp any any range 135 139 log
deny udp any any range 135 139 log
Filter outgoing NetBIOS traffic please
deny udp any any eq 69 log
Filter outgoing TFTP and protect configurations
deny udp any any eq 514 log
Filter out Syslog and don't send it over the Internet
permit icmp any any echo
permit icmp any any parameter-problem
permit icmp any any packet-too-big
permit icmp any any source-quench
permit ip 192.168.113.0 0.0.0.255 any
permit ip 192.168.114.0 0.0.0.255 any
permit ip 10.193.111.0 0.0.0.255 any
permit ip 10.193.112.0 0.0.0.255 any
These allow any internal address to send packets out.
permit udp any any range 33400 34400 log
This is for the internal interface and allows any internal address to send packets out.

ip access-list internal-gateway
no access-list internal-gateway
permit icmp any any echo
permit icmp any any parameter-problem
permit icmp any any packet-too-big
permit icmp any any source-quench
permit ip 192.168.113.0 0.0.0.255 any
permit ip 192.168.114.0 0.0.0.255 any
permit ip 10.193.111.0 0.0.0.255 any
permit ip 10.193.112.0 0.0.0.255 any

These commands allow traffic between the two network segments to flow.

Larger ACLs would have been difficult to maintain, and for GIAC Enterprises purposes, these seem to suit them alright.

Primary Firewall 1 & VPN Symantec

Firewall Configuration

The following usage is intended for the Symantec Firewall:

- ✓ Allow external access to the web
- ✓ Allow SMTP from internet to mail server and from mail server to internet
- ✓ Allow DNS outbound from DNS server to Internet

Remember throughout the following, whenever entries are indicated from the management console, “Save and Reconfigure” is implied to save the changes and apply them to the firewall.

Base machine

Windows 2000 server has been installed, and all unnecessary services have either been removed or made inactive. This includes things like IIS. All of the latest service packs and security patches have been applied.

Setting the route

Set default gateway:
On the Win2K machine where SEF is installed, the default gateway for the external interface is set to 192.168.11x.6 (Internet router). There is no default gateway specified on the internal interface.

Static route for internal is configured with the following command at the command prompt:
route –p add 192.168.114.0 mask 255.255.255.0 192.168.1.3
route –p add 192.168.113.x mask 255.255.255.0 192.168.113.1
route –p add 0.0.0.0 mask 0.0.0.0

This tells the firewall that anything bound for the European office should be directed to the inter-office T1 router interface, and the –p argument makes the route persistent.

The next line tells the firewall that anything destined for the 192.168.113.x network needs to go through the internal interface.

The third line indicates that all other traffic needs to be sent to the external interface (and on to the router).

Securing passwords

Rempass or remote management passwords are disabled. Configuration of the machine must be done from the machine itself.

DNS Records

The firewall is not being used as the DNS server. From within the management console, DNSD has been disabled by selecting "Proxy Services" from under "Access Controls", and defining the properties of DNSD by unchecking "enable DNSD".

In addition, we have configured two Generic Serviced Protocols for port UDP/53 with a source port of 1024-65535 for DNS Resolver requests. One is defined as dns_udp with protocol UDP and source port 1024-65535, Destination port 53 (Request) and one is defined as dns_udp_res with protocol UDP and source port 53, Destination port 1024-65535 (Response).

This is done by selecting "Protocols" from the "Base Components" folder. A "New Protocol" action is selected, and names and ports are entered as indicated above.

Network entities

Symantec recommends defining network entities – the groups of hosts that will be passing data through the firewall. We have the following network entities defined:

- ✓ SMTP Gateway
- ✓ WWW Server
- ✓ DNS Server
- ✓ Router
- ✓ Syslog
- ✓ Database
- ✓ FTP
- ✓ Internal firewall
- ✓ Internal workstations
- ✓ Internet

Configurations of these entities are done from within the management console, from within the "Base Components" folders.

Defining Redirects

There are non-routable addresses within our DMZ, so we need the firewall to redirect traffic that is meant for these services. This is done through the "Access Controls" folder. Right clicking on "Redirected Services" will allow you to define new redirected services. We will define our services as follows:

Service: SMTP

Requested address: *198.100.1.20*
Address Mask: 255.255.255.255
Redirected Address: 192.168.11x.10
Redirected port:25

Service: DNS
Requested address: *198.100.1.22*
Address Mask:255.255.255.255
Redirected Address: 192.168.11x.13
Redirected port:53

Rule Applications:

In order for the rules to work well, they must be specifically ordered: those that apply to hosts should be applied prior to those that apply to a subnet. The subnet rules should be followed by domain rules, domain rules followed by general Internet rules.

Rules will be defined as follows (for all of the following, expand "Access Controls", right click on "Rules" and select "New"):

Rule 1: Internal addresses can be allowed to browse the Internet

For connections coming in via: DMZ
From source: Private Network User Group
Destined for: Universe
Coming out via: External
Allow or Deny: Allow
Services: http (choose configure)
(Special configuration) Allow HTTP/Allow HTTP over valid SSL on Standard Ports
Time Range: Anytime
Authentication:None
Apply rule to: everyone
Alert Thresholds: not configured
Log normal activity: Y
Application data scanning: Y

This rule does the following – anything from the DMZ interface card coming out on the external side of the firewall is allowed http services – regular HTTP and HTTP over valid SSL. This rule applies anytime, and does not require any authentication. Alerts are disabled, but logging is enabled as well as application data scanning: we verify that the request is valid prior to sending it out.

Rule 2: Deny unnecessary services

For connections coming in via: External
From source: Universe
Destined for: Universe
Coming out via: DMZ
Allow or Deny: Deny
Services: select all services except the following: dns_udp, dn_udp_s2s, http, http/s, smtp, sqlnet
Time Range: Anytime
Authentication: None
Apply rule to: everyone
Alert Thresholds: not configured
Log normal activity: Y
Application data scanning: Y

Rule 3: DNS Resolver traffic (Inbound)

For connections coming in via:External
From source: Universe
Destined for: DNS Server
Coming out via: DMZ
Allow or Deny: Allow
Services: dns_udp, dns_udp_res
Time Range: Anytime
Authentication: None
Apply rule to: everyone
Alert Thresholds: not configured
Log normal activity: Y
Application data scanning: Y

Rule 4: DNS Resolver traffic (Outbound)
For connections coming in via:DMZ
From source: DNS Server
Destined for: Universe
Coming out via: External
Allow or Deny: Allow
Services: dns_udp, dns_udp_res
Time Range: Anytime
Authentication: None
Apply rule to: everyone
Alert Thresholds: not configured

Log normal activity: Y
Application data scanning: Y

Rule 5: Tunnel Traffic for Webserver front end
For connections coming in via:ANY VPN
From source: auto-entered
Destined for: WWW Server
Coming out via: ANY VPN
Allow or Deny: Allow
Services: http, https
Time Range: Anytime
Authentication: None
Apply rule to: everyone
Alert Thresholds: not configured
Log normal activity: Y
Application data scanning: Y

Rule 6: Tunnel Traffic to FTP server
For connections coming in via:ANY VPN
From source: auto-entered
Destined for: FTP server
Coming out via: ANY VPN
Allow or Deny: Allow
Services: ftp
Time Range: Anytime
Authentication: None
Apply rule to: everyone
Alert Thresholds: not configured
Log normal activity: Y
Application data scanning: Y

Additional Steps

HTTP URL Pattern protection enabled

httpurlpattern.cf has been configured with these additional lines, as recommended by Symantec, and http.urlpattern has been added to the http rules.

```
# The following would need to replace the ../ and the end of the file.
# Blocks a url which has 220 or more characters after the .ida.
.*\.ida.....................................................................
.....................................................................
.....................................................................
.............*
# Blocks a url which has 220 or more characters after the .idq.
```

```
.*\.idq.....................................................................
.....................................................................
.....................................................................
.................*
# Each '.' (dot) represent a character
# These are to block anything .eml for Code Red worm variants.
# You could be creative on using any regular expression patterns
# as long as it does not block legitimate requests.
# Refrain from entering too many entries, which may increase load on
firewall
.*\.eml
.*//www/*
\.\./
```

Packet headers removed
An http Rule is created and in the Advanced Services tab "http.remove-header.server" is selected. This removed the server header information from packets. (Symantec, pg. 210)

Don't let the system be used as a proxy!

http.noproxy has been added in the Advanced Services tab of the "No proxy" rule.

VPN

The VPN will be used by everyone that needs to connect to the database and servers.

User accounts are created on the firewall. VPN password is different than the database password for the account: an intruder will require both to compromise the database login. Users have been created with appropriate user groups.

VPN is using IPSEC and IKE for key exchange.

Connections on the VPN are automatically routed to the web server front-end for the database login at

The VPN has been configured with IPSEC/IKE protocol, and an MD5 hash for the data integrity protocol. The Data Privacy Protocols have been configured as 3DES (primary) and DES (secondary). Data compression has been set to LZS, and Session timeout has been set to 4 hours. (This limits an intruder's capability on the system by restricting the length of time they can be connected in a single session.) The VPN sessions will timeout within 30 minutes of inactivity. The Encapsulation Mode is via Tunnel, and Data Integrity is applied to the data portion of the packet.

All connectors to the VPN have been configured with Symantec SEVPN Client Software.

Client Configuration

The security gateway has been configured with the following information:
Name: External
Description: External Interface of SEF/SEVPN
Type: Security Gateway
Address: 198.100.1.13
Network Mask: 255.255.255.0

A Group has been set up as follows:
Name: Internal_Network
Description: Access to DMZ
Type: Group
Members: DMZ (defined as the 192.168.113.x subnet)

User group has been set up as follows:
Name: Client_Group
Description: Basic clients
User Distinguished Name: (blank)
Issuer Distinguished Name: (blank)
Authentication Method: NONE
DNS: (Internal DNS Server)

Users are then created as follows:
Name: (Assigned at time of issue: generally client name)
Description: (Assigned at time of issue: generally client name)
UserID: (Auto-filled)
Groups: Client_Group
IKE Enabled
Phase1 ID: (Default of user name)
IKE Authentication: Shared Secret Allowed
Shared Secret password
Primary IKE user group: Client_Group

VPN Policies are defined in Default Gateway (Route) Configuration.

Secure Tunnels

A secure tunnel has been configured named Tunnel_1 – which is the secure tunnel between the two offices. Local entity has been defined for the 192.168.113.x subnet, with the local gateway being defined as External. Remote entity has been defined as 192.168.114.x subnet, with the remote gateway being defined as the interface for the European T-1 connection.

A second secure tunnel has been configured named Tunnel_2 – this is the secure tunnel between the remote users and the main office. Local entity has been defined for the 192.168.113.x subnet, with the local gateway being defined as External. Remote entity is defined as users, so the remote gateway option becomes disabled.

Both of these use IPsec/IKE encryption.

Remote policies

Remote policies have been configured for the users that contain a common installation password, IP address of the Gateway, Client IKE Authenticayion Method and Phase-one ID of both the Gateway and the Client. This is delivered via diskette to the clients upon initial install of the client end product.

Configuring Logging:

Logging is configured to log everything above the notification level.

Primary Firewall 2 Pix (Internal)

The Pix firewall is serving as a layer of defense into the internal network, and its main role is to "put a wall" up between the DMZ and the internal network - it restricts access to the internal network and controls access from the internal to the external network. This direction is based on mapping to internal tables the addressing within the packets and directing packets accordingly.

As a refresher, here is a diagram of the Pix and its position on the network.

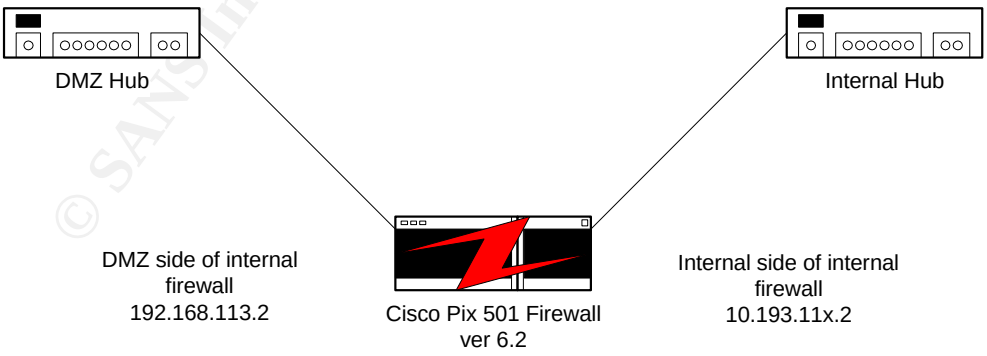

This firewall will participate in protocol & packet filtering, extended logging, and information hiding (where a firewall plays a main role in disguising the sources behind it.)

The protocols and services that are unnecessary should be blocked or restricted if not all systems require access to all services.

In short, the firewall has a main role of preventing unauthorized access, screening traffic and restricting services. It is also providing the important role of logging activity and disguising sources of the internal network.

Configuring interfaces that are used

Since both of the interfaces will be ethernet interfaces, we will configure them together as follows – "outside" will be called "outsideintfw" and "inside" will be called "insideintfw" for ease of identifying the placement on the network throughout the configuration:

nameif ethernet0 outsideintfw security0
nameif ethernet1 insideintfw security100

These lines establish the following: names for the interfaces and the security level for the interface. Cisco has established a traffic flow pattern based on what it terms "security levels". By default, hosts connected to a lower security level interface cannot pass traffic to a higher security level, unless it is expressly configured through access lists to do so – although the higher security level can pass information down if properly configured. (Degu & Bastien, p. 95)

After we name the interfaces, we will assign the ip address that has been established.

ip address outsideintfw 192.168.113.2
ip address insideintfw 10.193.11x.2

Secure the passwords

The next step is to secure the password. In order to offer good security, we want to encrypt the password when it is assigned (where XXXXXX represents the password):

enable password XXXXXX encrypted
passwd XXXXXX encrypted

We have entered two lines: one encrypts the password used to enter "enable" mode – the privileged mode which allows you to change configurations. The second password is securing telnet mode – it is the password required for the telnet access to the firewall.

Secure protocols

As stated at the outset, not every host needs access to every protocol. Additionally, protocols can be abused by attackers, by manipulating ACK bits and other methods. To counter these threats, we will apply Cisco Adaptive Security Services (CASA) for protocols where this may be a threat. Essentially, Pix maintains a stateful connection for the protocols defined here.

fixup protocol ftp 21
fixup protocol http 80
fixup protocol rsh 514
fixup protocol smtp 25

Access lists

We will establish two access lists. One will be filtering traffic exiting the internal network, the second will be to filter traffic entering the internal network. Our goal with each list will be to use as few lines as possible – minimizing confusion and configuration, and yet provide the restrictive flow of traffic for which a firewall is implemented.

Filtering exit traffic

The following commands will implement this model:
traffic to and from the mail server and FTP server are more important than requests to the Internet. File sharing ports that are established in company policy as being counter-productive will be blocked. In addition, noted worm ports will be blocked to prevent an infection from spreading.

access-list frominternal permit tcp host 10.0.0.0 255.0.0.0 192.168.11x.10 eq smtp
access-list frominternal permit tcp host 10.0.0.0 255.0.0.0 192.168.11x.15 eq ftp
access-list frominternal permit tcp host 10.0.0.0 255.0.0.0 192.168.11x.15 eq ftp-data
access-list frominternal permit tcp host 10.0.0.0 255.0.0.0 192.168.11x.10 eq http
access-list frominternal permit tcp host 10.0.0.0 255.0.0.0 192.168.11x.13 eq 53
access-list frominternal deny tcp any any kaazaport
access-list frominternal deny tcp any any limewireport
access-list frominternal deny tcp 10.0.0.0 255.0.0.0 any range 3127 3198
access-list frominternal deny tcp 127.0.0.1 255.255.255.255 any
access-list frominternal permit tcp host 10.193.11x.17 host 192.168.11x.17 1433 log
access-list frominternal permit tcp host 10.0.0.0 255.0.0.0 any eq ssl
access-list frominternal permit tcp host 10.0.0.0 255.0.0.0 any eq www

Filtering entering traffic

The only traffic coming internally will be: a) the database backup traffic and DNS communication, or responses to queries from the internal network. Any inbound traffic with a source address from behind our firewall will be blocked: because of the NAT implementation, any detection of 10.193.xxx.xxx on any other network segment should be deemed suspicious! We also want to log any denied connections, which means we need to specify the implied deny. (Prunoiu, pg.4)

access-list fromdmz permit tcp host 192.168.11x.17 host 10.193.111.17 eq 1433
access-list fromdmz permit udp host 192.168.11x.2 host 10.193.111.20 eq dns
access-list fromdmz permit tcp host 192.168.11x.2 host 10.193.111.20 eq dns
access-list fromdmz deny tcp 127.0.0.1 255.255.255.255 any
access-list fromdmz deny ip any any

Configure logging

We definitely want syslog to record things. We need to specify the level of message which should be recorded, and where the syslog server is – and whether or not the time of events should be recorded with those events.

logging on
logging buffered 4
logging trap informational
logging history critical
logging host outsideintfw 192.168.11x.11
logging timestamp

Configuring interfaces that are used (Stage 2)

Now we want to configure line speeds. To make the configuration simple, we want to use auto-negotiation on the line. The Maximum Transmission Unit (MTU) needs to be set to Ethernet speed, which is standard 1500.

interface ethernet0 auto
interface ethernet1 auto
mtu outsideintfw 1500
mtu insideintfw 1500

The next step is to configure the IP addresses and subnet for each interface.

ip address *outsideintfw 192.168.0.1 255.255.0.0*
ip address insideintfw 10.193.0.1 255.255.255.0

Additional configurations

Whenever possible, take the extra steps to protect against denial of service attacks. If the firewall cannot identify the route the packet took, or if it has an incorrect interface associated, the packet should be dropped. Although this is an internal firewall and supposedly already protected by the router and the first firewall, this extra precaution is an added layer of defense.

ip verify reverse-path interface outsideintfw

GIAC is performing IDS functions through different means, so we should make sure the IDS configurations are minimal. We'll use default configurations.

ip audit info action alarm
ip audit attack action alarm

There is no failover internal firewall. There is a second firewall with the appropriate configurations built in sitting on a shelf, it can replace either this firewall or the one in Europe should it need to. The expense of a failover router for each office could not be approved. As a result, we are turning off failover and specifying empty times and ip addresses.

no failover
failover timeout 0:00:00
failover poll 15
failover ip address outsideintfw 0.0.0.0
failover ip address insideintfw 0.0.0.0

Configuring NAT

In order to maintain the most security within the internal network, we are going to implement Network Address Translation (NAT) on a small scale. Because we are doing this on both the European and American internal firewall, an ip address of 10.x.x.x should raise suspicion if it appears on any other network segment. This also allows the internal networks to use the same address space without conflicting with each other.

First, we need to assign the address for traffic that is leaving the internal network – this will be bound to the outsideintfw interface. The "global" command instructs what addresses should be translated into. Because every outgoing user from the internal network can use the same interim address in the DMZ, assigning these is a fairly simple proposition.

global (outsideintfw) 1 192.168.11x.18

We then want to identify the networks: all hosts internally should receive a NAT address on exit.

nat (insideintfw) 1 10.193.11x.0 255.255.0.0

We always want the database servers to be able to communicate, and since we have higher security defined on the internal interface, we need to define a "transparent" static translation – it doesn't really do anything other than specify the IP is the same, but is required for our lower security interface to communicate with the higher security one.

static (outsideintfw, insideintfw) 10.193.11x.17, 10.193.11x.17 netmask 255.255.255.255

Assign Access Lists

Use the ACLs we built and assign them to the appropriate interface.

access-group frominternal in interface insideintfw
access-group fromdmz in interface outsideintfw

...and if you don't know where to send the packets, route them to the next firewall.

route outsideintfw 0.0.0.0 0.0.0.0 192.168.11x.1 1

Disable unnecessary service

Disable unnecessary services which may allow others to configure remotely, and make sure default settings on these services are changed in the event of their inadvertent enabling.

no http server enable
no snmp-server location
no snmp-server contact
snmp-server community x(Xf8345@#d
no snmp-server enable traps

Set timeouts

Make sure to set timeout values that are pertinent to your network. These are currently still at defaults for most things.

arp timeout 14400 (number of seconds. This value represents 4 hours)
timeout xlate 3:00:00 (hours until translation slot is available)
timeout conn 1:00:00 half-closed 0:10:00 udp 0:02:00 rpc 0:10:00 h323 (hours until connection slot is available – idle time) (Cisco)
0:05:00 sip 0:30:00 sip_media 0:02:00
timeout unauth 0:05:00 absolute
telnet timeout 5
ssh timeout 5

Final Housekeeping Items

Set the firewall to reclaim resources if it is running low in the user authentication system. This is a default setting (floodguard) where PIX will list messages declaring whether it is out of resources or tcpusers, and resources may be reclaimed depending on urgency. (Cisco)

floodguard enable

terminal width 80 (sets width of display for console connections)

Should we want to, at some point, set up a site-to-site VPN tunnel so the really internal networks can communicate with each other, we would do so using the Symantec guidelines to connect the inter-office WAN links on the SEF to the internal Pix connections. (Symantec, How to set up...)

Part 3: Design Under Fire

"Then he'll want to look in a mirror to make sure he doesn't have a milk mustache. When he looks into the mirror he might notice his hair needs a trim..." -Joffe

Introduction to the Assignment

In this section, we were asked to research, design and execute an attack against a previously posted practical. For this assignment, I will be using Brian Rudzonis' practical from February, 2004, which was posted here: http://www.giac.org/practical/GCFW/Brian_Rudzonis_GCFW.pdf

How this practical was chosen:

In order to complete this assignment, I had to pick a practical that was posted from 6 months of mine. I originally wanted to complete my practical by the end of April, which meant there were very few assignments originally that would have met the requirements. I relaxed my own submission until May, and waited until all other sections were complete to ensure that I had a better selection of practicals. In the end, however, Mr. Rudzonis' was just second in the March expirations, newly posted when I went looking for a "target", so I figured there was less of a chance that someone else was already targeting that practical. There was also the added benefit of using similar technologies within our network.

I am going to assume the role of a "Black Hat" attacker, hired by the competition to effectively neutralize the ordering system of GIAC. In order to accomplish this task, I will assume a couple of things that were undocumented from Mr. Rudzonis' practical. 1) There is a corporate PBX with a voicemail and directory listing in place. 2) Assumptions are made about the availability of information in Whois, as this is not accounted for within his paper.

Following is his network diagram:

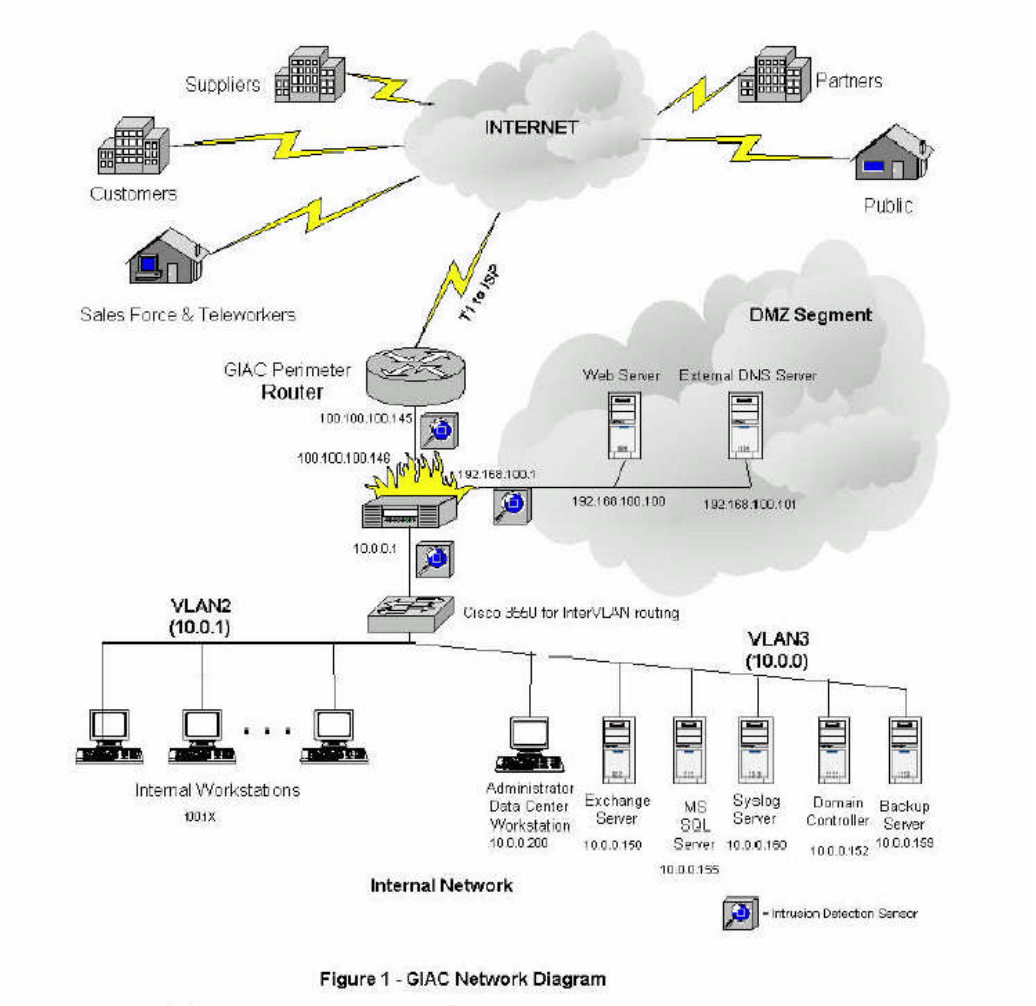

Figure 1 - GIAC Network Diagram

Stage 1: Reconnaissance

There are some things, as an attacker, I already know about my target: the company name and location. That is all I need to get started on finding out what I really need to know. During this stage, I will be making notations of everything that I discover, because it might prove useful later.

Web Search

By accessing the main corporate web site, I assume the role of "the general public." I discover on this site "general company information, marketing, contact information (e-mail and telephone) and the front-end log in..." I make a note of the contact numbers listed on the site. There is an employee listing which

doesn't list phone numbers but does list names under "suppliers" and "sales". I also make note and take screen shots of the front-end to log in that I discover in one of the links. That might come in handy.

For additional information, I'll conduct a whois search. (Please note information on the following screen shots have been modified to match the specifications of Mr. Rudzonis' practical.)

I use the following URL:
http://www.networksolutions.com/en_US/whois/index.jhtml

and type in the name I used to access the web site....giac.com

I get the following screen (please note code has been changed from actual screen):

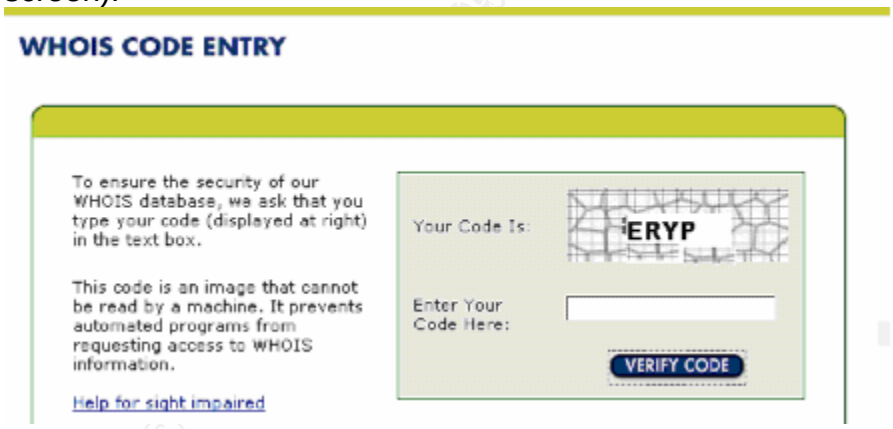

I enter the code in the box below, and am rewarded with the following information, which includes an e-mail address that displays the naming convention used within the organization for e-mails:

WHOIS SEARCH RESULTS

WHOIS RECORD FOR

giac.com Back-order this name

The data in Register.com's WHOIS database is provided to you by Register.com for information purposes only, that is, to assist you in obtaining information about or related to a domain name registration record. Register.com makes this information available "as is," and does not guarantee its accuracy. By submitting a WHOIS query, you agree that you will use this data only for lawful purposes and that, under no circumstances will you use this data to: (1) allow, enable, or otherwise support the transmission of mass unsolicited, commercial advertising or solicitations via direct mail, electronic mail, or by telephone; or (2) enable high volume, automated, electronic processes that apply to Register.com (or its systems). The compilation, repackaging, dissemination or other use of this data is expressly prohibited without the prior written consent of Register.com. Register.com reserves the right to modify these terms at any time. By submitting this query, you agree to abide by these terms.

Organization:
GIAC Enterprises
GIAC
123 Pretend Address
Somewhere, VA 20181
USA
Phone: 703-867-5309
Fax: 703-867-5310
E-mail: jname@GIAC.com

Registrar Name....: Register.com
Registrar Whois...: whois.register.com
Registrar Homepage: http://www.register.com

Domain Name: GIAC.COM

Created on.....: Tue, Jul 03, 2001
Expires on.....: Sat, Jul 03, 2004
Record last updated on..: Fri, Nov 28, 2003

Administrative Contact:
GIAC Enterprises
GIAC
123 Pretend Address
Somewhere, VA 20181
USA
Phone: 703-867-5309
Fax: 703-867-5310
E-mail: jname@GIAC.com

Technical Contact:
Register.Com
Domain Registrar
575 8th Avenue
New York, NY 10018

Under the technical contact, I find the following goldmine of information, and I now know the Name Server of the organization:

```
Zone Contact:
Register.Com
Domain Registrar
575 8th Avenue
New York, NY 10018
US
Phone: 902-749-2701
Fax..: 902-749-5429
Email: domain-registrar@register.com

Domain servers in listed order:

NS1.GIAC.COM 100.100.100.162
```

Phone Reconnaissance

I call the main telephone number listed on the web site first explored in the Web Reconnaissance. "Welcome to GIAC Enterprises..." I hear. I wait through the messages until I hear "for the company directory, hit 3". Upon pressing [3] I hear "Please enter the first three letters of the person's first or last name, followed by a pound sign. For further instructions, press star". I press [*]. I then hear "To select a person by department, press 1." I press [1] and hear "Please enter the first three letters of the department you wish to reach, followed by a pound sign.". I select [4][6][3][#]. "That is not a valid selection, please try again." I select [4][8][#]. I hear "I.T. Department" and then I hear "for Joe Name, dial 324. For Deb Othername, dial 325. For Fred Finalname, dial 326."

Hmmm... Joe Name, at extension 324 is probably the jname@GIAC.com listed in the administrative contact information of WHOIS. I would guess, judging from the I.T. Department phone tree, that he is the head of the department, based on a) the directory is not alphabetical, nor is the numbering, indicating Fred probably was the last hired. b) he is listed first in the phone directory and also listed as the technical contact for the domain name.

Mitigating Strategies

There are two glaring holes in the security so far:

1) Information available through Whois is far too friendly. This often occurs when sites have been established for awhile, because when the Internet was first set up this information was not considered for "evil". To counteract, references to specific names should be removed and proxy e-mail addresses or other e-mail addresses which hide the username and the naming convention used for e-mails internally should be used. Network Solutions offers tips on protecting Whois information here: http://www.internetprivacyadvocate.org/ProtectYourPersonalInfo.htm

2) The PBX is not considered as part of perimeter security. It contains a treasure trove of information, that if used correctly, can turn up all sorts of

company information. The likelihood that it is being audited to detect such reconnaissance is not high, since it was not mentioned in the practical at all. If it was being monitored, the accounting system would need to note how many times the caller changed commands if trying to track someone using the same phone connection to perform reconnaissance – and it would be helpful to have Caller ID, which would identify if the same number called numerous times to perform different tasks.

Stage 2: Scanning the network

While the telephone network may have provided some additional sources of potential vulnerabilities, sweeping and scanning the network proved fruitful with some tools from my bag of tricks.

Since I had discovered the DNS server registered publicly as 100.100.100.162, and that many companies have addresses allocated in order, I am going to try some tools against 100.100.100.160-100.100.100.165. By using a small address space, and tools that are not very intrusive to perform this sort of information-gathering, I may be able to escape detection. Some of these tools I will spread out over time, so that unless state is maintained or some correlation of events is being performed over a period of time, it will further obfuscate my activity.

First thing I want to do is try to see what hosts are alive. In order to do this, I use SuperScan, a graphical tool that is easily downloaded from Foundstone for free http://www.foundstone.com/index.htm?subnav=resources/navigation.htm&subcontent=/resources/proddesc/scanline.htm. They are currently on version 4.x, I was using version 3.x and these screen shots are from the earlier version.

Find some hosts

SuperScan is a speedy Windows port scanning tool. I double-click on my SuperScan icon and enter the following information into the window: I configure it to Ping only, and if it gets a reply to resolve the hostname. I use my dial-up connection to give me a connection to the Internet.

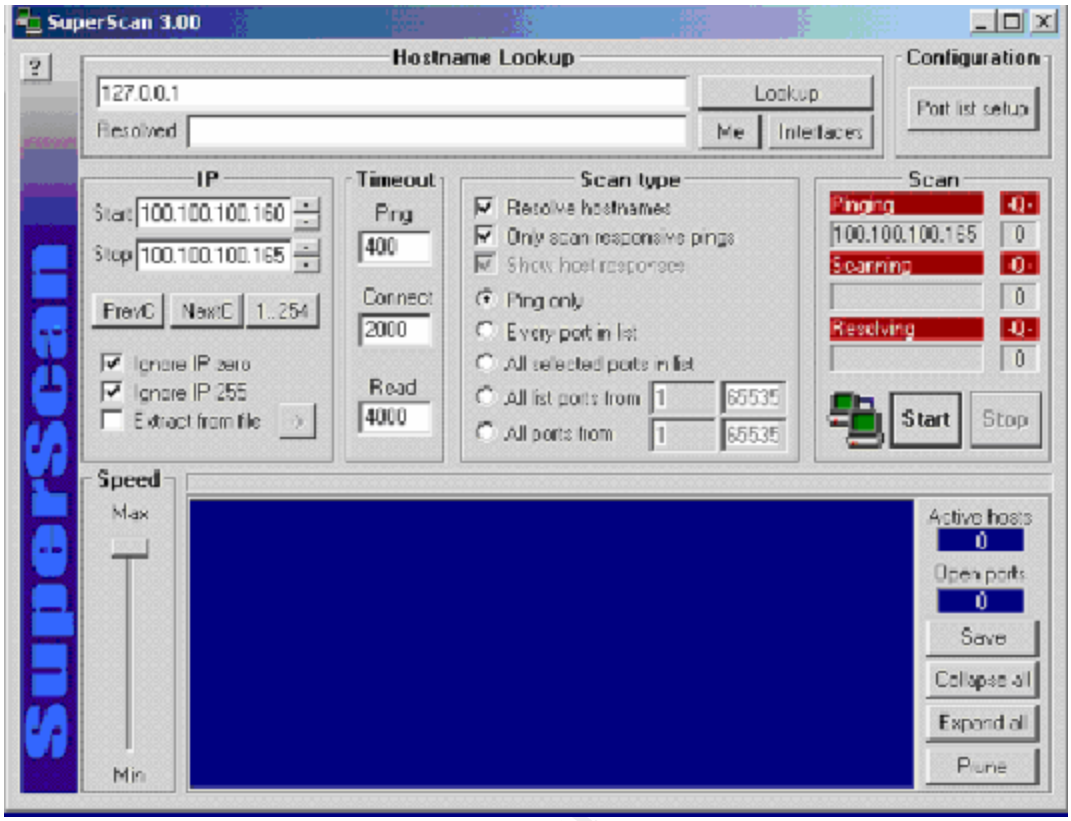

This will not give me open ports, but will tell me which hosts in this range are active and if it is active, what the host names are.

My scanning returns the following addresses:
100.100.100.161 (Web Server)
100.100.100.162 (DNS – but I already knew this!)
100.100.100.163 (Mail server)
100.100.100.164 (Syslog)

Network Topology

Now I know what hosts are associated on the network, I probably want to traceroute these addresses. Traceroute will tell me the general network topology, what gateways, routers, etc. take me to where I want to go. In order to escape detection, I switch to a cable modem through a different network provider.

I go to the command prompt and type in tracert 100.100.100.161 (please note the responses below have been simulated for this exercise, and identifying network information has been blacked out.)

```
C:\>tracert 100.100.100.161

Tracing route to 100.100.100.161 over a maximum of 30 hops

  1    <10 ms    <10 ms    <10 ms  1  .6 .9 .3
  2    <10 ms     10 ms    <10 ms  1.   6 . .1
  3     *         *         *      Request timed out.
  4    <10 ms    <10 ms    <10 ms  6 .9 .1  .1
  5    <10 ms     10 ms    <10 ms  6 .9 .2 .1
  6  6 .9 .2 .1     reports: Destination host unreachable.

Trace complete.
```

While I receive “Destination host unreachable” messages, this is not from the target network. I enter the IP into “Whois” and discover this resolves to my ISP. For this “whois” search, I use ARIN at http://www.arin.net/, because Network Solutions does not easily resolve IP addresses.

Right at the top of ARIN’s home page you can enter an IP address to get the associated information.

Access through Ports

I may not be able to tell what the topology is, but that won’t stop me! I return to SuperScan (and my dial-up connection) and select different tools – I choose to try a list of ports. I want to know which ports are open on the hosts I have selected.

The results reveal the following (please note these are theoretical results based on Mr. Rudzonis’ document :

HOST	Port open	Usage?	Reference
100.100.100.161	80	http	**RFC 1945**
100.100.100.161	443	ssl	Internet Draft
100.100.100.162	53	dns	**RFC 1034, 1035**
100.100.100.163	25	smtp	**RFC 821, 822**
100.100.100.164	514 (UDP)	syslog	RFC 3164

Well, now I know which ports are open on these hosts. Hmmm.

Vulnerability Scanning

If I wasn't sure before, I am pretty sure now that 161 is the webserver and 163 is the mail server. At this point, scanning for vulnerabilities and configuration errors would be a logical step. There are a couple of different tools that I can use to do this, and from my little Windows box I find that N-Stealth does a nice job of scanning for web vulnerabilities if you know how to configure it. It's free and pretty fast, and has a number of different configurable options, depending on the mission.

I open the program and enter the WebServer IP and be sure to select Top20 from the ScanMethod drop-down. According to the N-Stealth Website, ...

> "This method will scan the web server for the top 20 vulnerabilities list published by SANS/FBI (www.sans.org). It is a very fast security check but it will certainly produce superficial results. It is recommended for brief security checks."
> http://www.nstalker.com/products/nstealth/features.php

Complete scan is not enabled in the Free Edition.

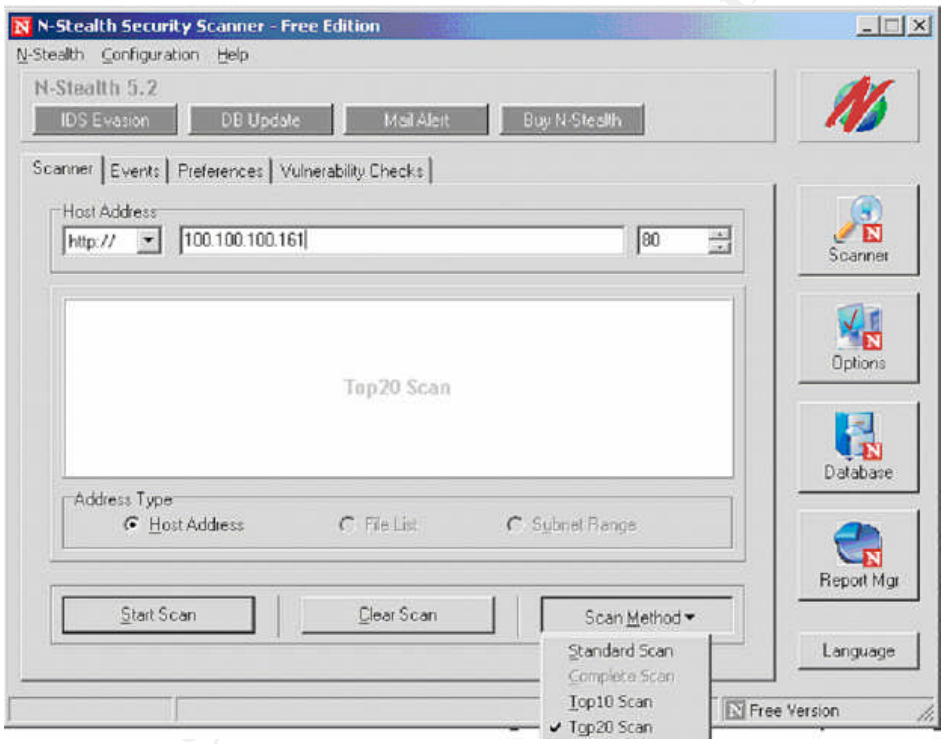

The following security checks are then enabled by default:

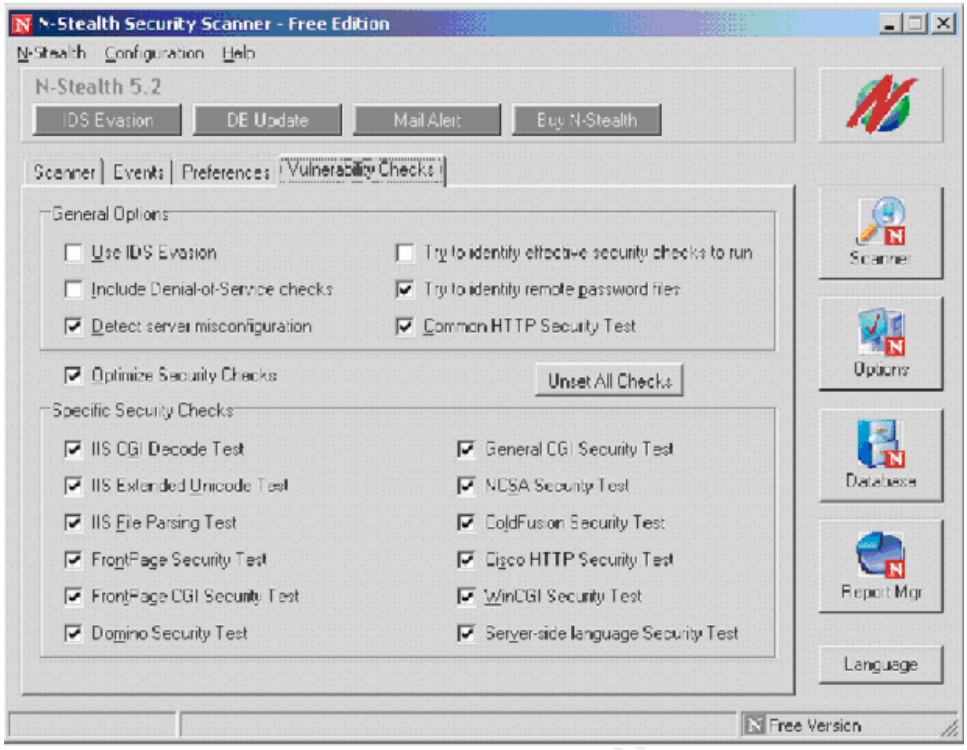

and I try some basic evasion to avoid being detected:

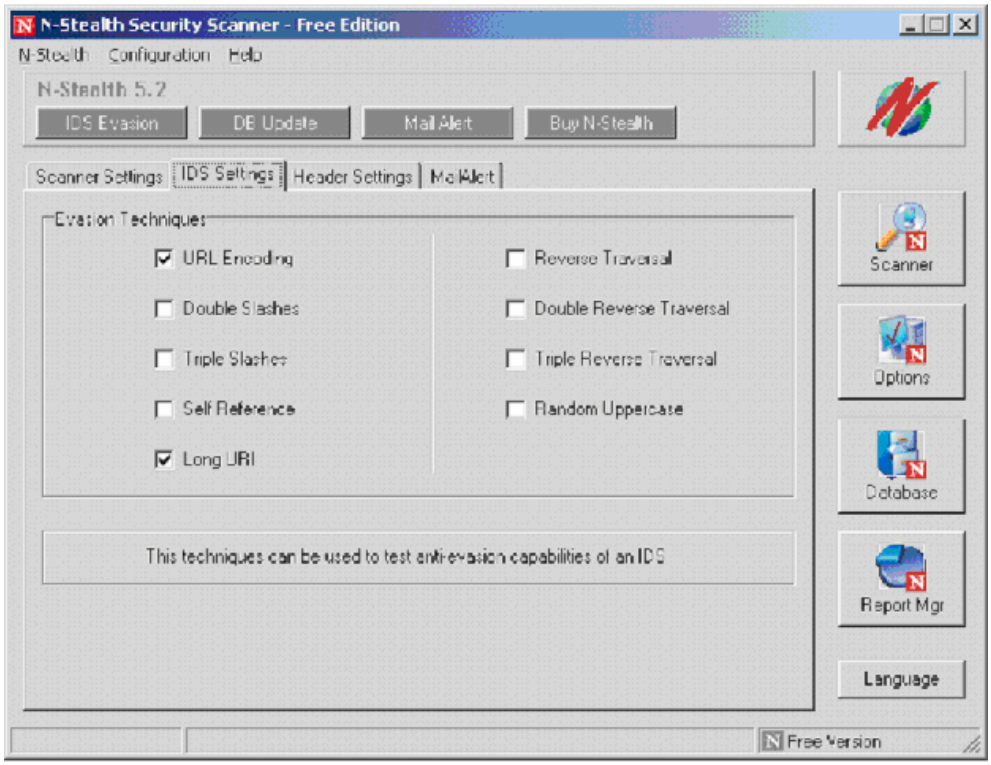

Since it is not possible to run this scan in my environment, I can only predict of the documentation from Mr. Rudzonis' practical whether or not vulnerabilities would have been detected, and the documentation does not specify whether any

tools were used to harden this server, although it does say the following: "the web server has been hardened since it is one of the primary attack targets for hackers...A custom policy template has been created...A minimum of services are running..."

Mitigating Strategies:

At this stage, I have discovered another problem in the configuration that could have been prevented.

a) Filtering for ICMP Echo Request packets should be applied at the gateway. This will make sure that answers to pings cannot be used to perform the reconnaissance illustrated here. Traceroute can also be thwarted by filtering ICMP Time Exceeded messages. (Skoudis, pg. 200) While this can be inhibitive to network management and potentially the ISP, routers can be figured to accommodate these needs.
b) Disable unused ports/services. On critical systems, take this a step farther and remove or rename references to the programs that are associated with those services.
c) The testing done on the network was not vulnerability specific, and these scans should be run regularly to ensure there are no rogue services/ports opened that may be in conflict with organizational policy.
d) Use a hardening tool to ensure that the web servers are as secure as possible.

Stage 3: Compromise an internal system

Based on the information I have collected, I have narrowed down my target: I specifically want to access the database of the Fortune Cookie Saying. This will involve compromising someone's logon on the web server and then see what that gets me on the MS-SQL server inside....

Setting the stage

When I travel to the web site, I am faced with the following information:

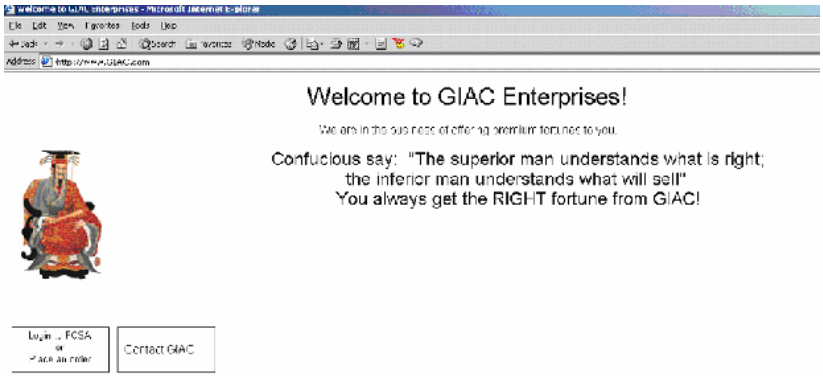

and if I click the first button, I get the following:

There's a lot I can do with the information I have gathered so far!

Compromise the systems

From here it becomes relatively easy to set up what is commonly called a "phishing" scam. (Note: a lot of the following example is modeled off of examples in K.K. Mookhey's SecurityFocus article)

Using the list of names available from the web site, I create an e-mail from jname@anotherGIAC.com that says:

> "We've been having problems with the database system and logins have been unexpectedly failing. Please click on the link below and confirm your information and that your login is still working properly. Thanks!
>
> http://www.GIAC.com:ac-piUq3027qcHw003nfuJ2%01@anOtherGIAC.CoM/3/?W4x6

Joe Name, I.T"

I send this to the people listed on the website using the naming convention firstinitial lastname@GIAC.com - i.e. jname@GIAC.com as indicated in the reconnaissance information I had obtained. I am careful not to send it to anyone who was identified as an IT Department member during our phone reconnaissance! I also make sure that I send it blind carbon copied in small batches, so that others cannot see who else it was addressed to, and by sending in small batches I hope I won't trigger any spam filters that may be in place. I send this e-mail on a Friday night. This is a good time to send it: perhaps people will access their e-mail on the weekends, and not bother I.T. even if they think that something is suspicious until Monday.

In this e-mail I am using a pretty basic e-mail attack: maybe not everyone will click on the link, but maybe some people will. Although the link looks legitimate, according to HTTP rules, it will connect the user to my website anOtherGIAC.CoM – which will display the following:

Looks familiar, doesn't it? Even if one person clicks the link and enters login id and password (which will send the information to me instead of to the system) – I have one login id and password that I can use to look around on the system. In addition, I ensure that once they click Login on my crafted page, they are answered with "Thank you for verifying your details."

The link I provide takes advantage of two things:

- ✓ Because most end-users are not familiar with HTTP rules, they do not know that this link, although it starts with GIAC.com, will actually go to the site indicated after the @ symbol
- ✓ There is a bug in Microsoft's Internet Explorer that if Internet Explorer sees %01 (a non-printing character representing 0x01) an attacker can hide the real location of the page because IE will not show anything after that in the title bar – so that my page appears to be from the legitimate domain. (More details on this vulnerability can be found here

http://www.zapthedingbat.com/security/ex01/vun1.htm, or in articles by SecurityFocus, found in the references section at the end of the paper.)

By obtaining access in this way, I have

1. compromised the internal system of the user who accesses the e-mail and clicks the link. They are now unintentionally sending to me their information.
2. compromised the web server by using the information that is sent to me and
3. compromised the internal database by using the information that is sent to me.

Mitigating strategies

Education, policies, and more education. These are the strategies that can be used to thwart attacks that depend on end-users not knowing enough to avoid them.

1. Put a policy regarding hyperlinks in e-mail in place. Ensure users know NEVER to click on hyperlinks arriving in their e-mail. If possible, filter or ensure that they are not active when they arrive at the user. Many users will not take the time to type in the long URL if it is not active, and so the risk of such an attack becomes less.
2. Educate end users on some of the common attacks and ensure they know how to spot “suspicious” URLs.
3. Where possible, encrypt your e-mail, at least internally. If instructions are issued from an IT Department via e-mail, have some way for the end user to validate those instructions and who the sender actually is.
4. Make sure your end users know NEVER to release their password to anyone. Implement an IT policy that IT will NEVER ask for a users password or send users information requesting that they reverify their information.

Stage 4: Retain access

At this point, to loosely quote Tennessee Williams, I am depending on the kindness of strangers.

There are a couple of measures I can take to help retain access and avoid detection.

First, provided I have multiple login names now available as a result of my phishing expedition, I will use these multiple names for a number of things, so that all my actions can't be traced to one user name.

I will also work fast to look around and see what might be interesting to me now that I have the access.

Once I have performed the additional reconnaissance, I might install a backdoor, such as a Reverse WWW Shell (van Hauser) I will also open a few accounts of my own, in case the accounts are not regularly audited for authorized accounts.

I probably will try to install a sniffer somewhere on the segment so that I can monitor other passing traffic and see what other interesting things are there.

Mitigating Strategies

1. To ensure that the files on the main file machines are as intended (and things like root kits or sniffers are not installed), Tripwire or a similar file-integrity-checker and reporting system should be considered.
2. User accounts should be regularly audited to ensure that they belong to valid users, and have appropriate levels of authorization.
3. Should an account be compromised from an unknown source, all accounts should change their passwords and take appropriate safety measures.

Conclusion

By not adequately protecting a crucial entry point to the network (the PBX), reconnaissance and achieving the end goal of additional information becomes relatively easy. PBX technologies have often been in place in organizations for a long time, and they are often not on the same replacement/upgrade cycles as the computers and desktops are. Voicemail and PBX vulnerabilities have been linked to critical information leaks in the past – witness the trouble that existed around the HP merger a few years ago (Singer) , and with the continuing convergence of technologies, careful attention needs to be paid to this additional entry point to the corporate network.

Part 4: A Recommendation for Event Collection and Correlation Technology

"... he'll probably ask for a pair of nail scissors. When he's finished giving himself a trim, he'll want a broom to sweep up. He'll start sweeping...." -Joffe

Introduction

Event collection and correlation technology is a relatively new-to-market technology which has the potential to greatly affect enterprise networks and the way in which events on the security devices are acted upon. In many enterprises, such as GIAC Enterprises designed in this paper, logging is done but not reviewed often enough to catch an attacker in the act. The sheer volume of events from the different devices make the logs long and cumbersome to review, and often it is difficult to discern which are the important events amid the network noise. Events that identify a threat to an organization's security are often hidden among other events and can be difficult to identify. Different products have their own methods for logging and alerting – some have their own consoles, some have their own proprietary logging formats. In Federal Computer Week, Yasin notes how "so much data flowing in from firewalls, intrusion-detection systems and other devices...some threats and potential attacks...not noticed." He calls this "security data overload". (Yasin)

Correlation technologies available today purport to sort through the messages for the enterprise, and link together those events deemed critical into an alert, allowing the administrators to see the important events and act upon them in a swift manner. Technologies such as these have the potential to increase security in the enterprise by allowing for a more rapid response to the attackers, and by establishing an overall security vision of the network at any given time. In the best of technologies, correlations are performed cross-platform, cross-vendor and cross-device, allowing for events from a firewall to correspond to that of a NIDS and that of an operating system. When this happens, an overall view of an attackers progress through the network can be gained. This correlation ideally will take place in near-real-time, alerting staff to potential threats in a more proactive timeline then manual review of logs or data mining can ever provide. In this section we will review one such product, look at the underlying technology and implementation, and evaluate its features and functionality for the larger enterprise.

Correlation technology is a step above the Enterprise Management systems from the past. While the initial trend for ESM systems was to get the information from different products into one console and database for consistent reporting and views, correlation technologies take this a step further by grouping events together into meaningful patterns and only displaying those events that are

important to the established security posture of an organization. Over the past few years, several products offering this technology have been offered, and the offerings continue to improve. Larger companies like Symantec and e-Security offer products alongside focused companies such as ArcSight, NetIQ, GuardedNet and NetForensics. Since these products are entering a highly competitive space, the information on products and their features is sometimes difficult to gather. A good source of introductory information on the products available is the SANS site. For additional insight into ArcSight, refer to Timothy Muchow's GSEC practical. Information on netForensics, GuardedNet, and e-Security can be found in Kevin McIntyre's practical. Additionally, I have had personal experience with some of these products – when product comparisons are called out, they will be from either the product information on the company web-site, additional papers within SANS, or personal experience and notations from working with the products.

Correlation: How it works

Before establishing an overview of the product, the technology and its value should be explained. There are different approaches to correlation, but in order to be effective, all products must provide the same basic components.

- ✓ Normalization: congregation of events into a standard format
- ✓ Collection: a method of collecting events from different security products
- ✓ Event Database: for storing events for data trending or forensics capability
- ✓ Correlation system: the methodology of establishing separate and unrelated events into an actionable item or alert

Normalization

In order for data to be viewed on a consistent basis, relevant information must be normalized. This means that data can be viewed on an apples-to-apples basis, and the relevant details of an event are not only represented in a standardized schema and standardized fields, but also in standardized formats.

Normalization can occur on three levels, and different products may apply this in different ways:

- ✓ schema normalization: fields broken out are reported within the same schema structure consistently
- ✓ message normalization: messages are normalized to a point – through perhaps the application of categories
- ✓ event normalization: events are normalized to the point that there is a consistency of information available for that event across point products – the description, the severity, the what it means and how to patch

For the sake of illustration, let's look at a simple event declaration.

Consider the following event from Snort:

```
Aug 19 07:39:34 xx.xxx.xxx.xxx    snort: [1:716:5] TELNET access [Classification:
Not Suspicious Traffic] [Priority: 3]: {TCP} xx.xxx.xxx.x:23 -> xx.xxx.xxx.xxx:4512
```

and the following from a Pix log

```
Sep 04 09:37:51 xx.xxx.xxx.xxx %PIX-6-307002: Permitted Telnet login session from
10.0.1.3
```

Both of these events indicate telnet access, but their formats differ greatly. Normalization would ensure that not only are the ip addresses and ports always located in the same place, but things like dates and protocol information are recorded consistently.

An example of normalization would be to reorder the critical information in these events into a standardized schema as follows:

```
Aug-19 07:39:34 xx.xxx.xxx.x 23 xx.xxx.xxx.xxx 4512 snort 716 TELNET
Sep 04 09:37:51 10.0.1.3 23 xx.xxx.xxx.xxx null PIX 6-307002 Telnet
```

This simple normalization of records puts the information in consistent fields and consistent formats.

```
(Date) (time) (source ip) (source port) (target ip) (target port) (vendor) (vendor alert)
(event)
```

Let's further apply message normalization:

```
Aug-19 07:39:34 xx.xxx.xxx.x 23 xx.xxx.xxx.xxx 4512 snort 716 TELNET
Sep-04 09:37:51  xx.xxx.xxx 23 xx.xxx.xxx.xxx null pix 6-307002 TELNET


where (date) is represented in the following format: Aug-19, time is represented
hr:min:sec, vendor is in small letters, event is all caps.
```

Why are both of these steps so important? Consider if the second event above was normalized just by schema, and had not applied the consistent formatting to the fields within the event.

To run a report on telnet events, a single report may not be able to pick up the second event because TELNET may not be considered the same as Telnet. In addition, overall views may get confusing to someone trying to sort on the information if it is represented inconsistently.

To apply a third level of normalization, we need to further drill into the events. Event normalization would allow the end user to be confident that a Telnet event from Snort would carry the same severity and basic meaning as a Telnet event from Pix. Event normalization may look like this:

```
Aug-19 07:39:34 xx.xxx.xxx.x 23 xx.xxx.xxx.xxx 4512 snort 716 TELNET  SEV3 Telnet is a
remote control program commonly used in Web server maintenance, and may violate corporate
policy.
Sep-04 09:37:51  xx.xxx.xxx 23 xx.xxx.xxx.xxx null pix 6-307002 TELNET SEV3 Telnet is a
remote control program commonly used in Web server maintenance, and may violate corporate
policy.
```

Now, the user knows what happens, what the event means, and why they care about it, and can investigate easily whether this system is one in which telnet should be allowed, or even whether it is running on the system. The end user does not need to be a firewall administrator or Snort expert to understand these events.

Different products use different methods to normalize. NetForensics, for example, converts event data into XML at the agent and then converts the XML into their own format. (Leadston, p.7) ArcSight achieves normalization at the manager level where conversion occurs into their proprietary format (Muchow, p.4). In future implementations of products, we may see adherence to some form of emerging standards, such as the DMTF Security and Protection Management standards, which will at least apply schema normalization to events. (DMTF) However, most security products have not fully embraced such standards at this point, and schema normalization alone does not fully meet the needs of the end user. End users need to evaluate products based on what level of normalization they require.

Collection

Just as different products achieve normalization in different ways, different products collect event data through different methods. For the purposes of this paper, we will call the collection sources "agents". These agents have two primary roles: they control the data collection from the security device, and they control the transport of the collected data to the management server. The transport may be configured on a "push" scheme (data is pushed to the management server by the agent, either as it is received or on established time intervals) or on a "pull" scheme (data is requested from the agent by the management server, usually on requested time intervals.) Let's look at how the method used to collect data can affect the scalability and portability of the solution as a whole.

Some products require an agent to run on the devices where the information is being collected from. In this configuration, the agent would need to be lightweight and unobtrusive so that it does not result in additional system load and does not affect the critical security functions of the device. Otherwise, the system could not be deployed in large environments without negatively impacting the security posture. neuSECURE from GuardedNet varies from other offerings in that it does not need an agent to run on the system that is being monitored. (McIntyre, p.9) Other products may have agents that collect from logs, management consoles, or product databases, depending on the vendor.

In addition to being lightweight, agents need to be flexible enough to be deployed in a variety of configurations. If a Snort installation is configured to log to a syslog server, and that syslog server is a Windows Kiwi server (as is the case in GIAC), having an agent that only reads from Linux or Solaris syslog offerings will

result in a need to reconfigure the network, probably an unacceptable overhead when implementing a solution that is supposed to make things easier.

Collection agents may or may not have roles to play in the correlation process. Some agents may do some preprocessing: normalizing or applying mini-correlation processing to the events received. Through this process, event roll-up (sending one event in place of multiple events with the same data) or the ability to discard uninteresting events may reduce network traffic by applying filtering and ensuring only interesting events get sent to the database.

Typically, the agent will be configurable to the user's environment. Since it is in charge of both the collection and the transport, it defines how "near real-time" the solution is. For instance, if the agent is configured to collect data as new lines are entered into a log, but only sends that information to the management server every ten minutes, then there will always be a timelag before the event is used in correlation.

The communication and transport between the agent and the management server is important. Some systems use this layer as one-way communication, and only to send the event data. Other systems have developed two-way communication, where the agent sends event information to the manager, and the manager can send configuration information back to the agent. Ideally, a product will use some form of encryption for the data and some form of authentication for data sources to preserve the security and integrity of information presented to the user from the management platform.

Event Database

The Event Database serves as storage for the security events, and it has two critical roles: it allows for data mining/after the fact correlation (associating an event that has long since passed with events occurring today) and as a tool for forensic analysis in the event of a security breach. In addition, it can be used for trending analysis and evaluating the overall ongoing health and security stance of an organization. Events that are stored in a database with a standard format, such as IDMEF or MIB, will also allow for 3rd party tools to analyze them (Sharan, p.8).

Because this database has the potential to process and store large amounts of data and many events, some features are critical to maintain it's longevity in the enterprise. The database must be robust, and respond quickly to queries in a highly volatile (non-static) database environment. In order to reduce the amount of storage required in the database, aggregation of events prior to storage may help. In addition, a strong data management utility and structure will help the user partition the data/archive the data in customizable ways, allowing the data to be stored, backed up, and pruned in a way consistent with an enterprise's security requirements. The enterprise may require a certain method/timing for

data retention, or that may be mandated by the industry in which the enterprise operates. Whatever the reduction and retention requirements are, the correlation solution must be able to be managed accordingly. Finally, the database of the system should be scalable and flexible, as the network and number of security devices on a network grows, the database must also grow to accommodate the new influx of events. There are many different database architectures out there, and a database administrator of DB2 may not have the time or resources to devote to learning and managing an Oracle database. In order to meet the database requirements, implementations of database solutions may differ.

Correlation Methodologies

The real center of a correlation product will be the method it uses to correlate events, and what it does with events once it correlates them. Correlation technologies may perform correlation in a variety of ways.

Rule-based correlation

In rules-based methodologies, events are evaluated against a set of rules, which may be stock rules (out of the box), or may be written by the network administrator to correlate events together. Rules will look for patterns occurring in the event stream, create associations on those events, and alert the analyst accordingly.

Field-based correlation

When field-based correlation technologies are applied, fields within events are evaluated and if those fields meets specified criteria, and then an alert occurs. For example, if a system is configured to alert when it sees a connection to a specific ip address, as long as the data is being reported in a normalized format (so that the ip address is reported in a consistent field), then an alert would occur, regardless of which security device detected that event.

Context correlation

A system that allows the end-user to populate environmental variables and then correlates on those variables allows for a consistent representation of the network security stance overall. For instance, if the end user can populate the system with asset information such as which IP Addresses would signify WebServers, or all WebServers are Apache, then correlations could be completed based on that information. (A NIDS just detected a signature which exploits a vulnerability on IIS? That is less relevant than this one which detected a signature which exploits a vulnerability on Apache and is further evaluated to be detected on our WebServer farm.) This context correlation can also be applied to integration with vulnerability assessment products: being able to detect events and correlate them to the Operating System, patch level, open

ports on a system and ONLY if they are relevant to the organization's context allow them to raise an alert would eliminate many false positives and a lot of background noise that can occur in an Security Event Manager. McIntyre refers to this type of correlation as *statistical based.*

Aggregation and filtering correlation

Through aggregation and filtering correlation methods, data can effectively be reduced, eliminating network traffic noise from the displays of relevant events on the system. This is generally applied through normalized categories or fields. Aggregation can also be accomplished by assessing the event stream for duplicate events reporting the same information, perhaps varying only by time. Aggregation correlation would then eliminate the second event and increase the count of occurrences in the first event to accurately reflect the system status. This assists an analyst by allowing for meaningful views of the event stream and a reduction of data to be analyzed.

Anomaly correlation

By evaluating the event stream and allowing correlations on non-standard traffic (hey, there's an FTP to this server which NEVER had an FTP event before), near real-time detection of previously unknown attacks would be a critical benefit to a security analyst.

Post-occurrence correlation

Data mining for related events also offers a value to an analyst. The ability to search a database and correlate events that occurred days ago after a pattern is detected would be critical in establishing the reconnaissance phase of attacks or in establishing the actual state of a system. For instance, if a weekly scan of a system with an anti-virus program installed on it revealed the presence of a rootkit, the ability to then return to the database and correlate activity that may have led to the installation of the rootkit in previous days would be invaluable.

Correlation to a standard dictionary of information

In the best of products, an analyst will be able to look at an event and know various things: what does the event mean, why it is important within their own environment, what patches need to be applied, what measures should be taken internally to ensure it does not occur again. Integration to dictionaries and databases such as CVE, BugTraq, or the newly established OSVDB would allow an analyst insight to events without having to resort to the vendor documentation or desktop of different point products, and then sorting out what the event means on the standard level.

Symantec's Incident Manager

GIAC has already established an attempt at collecting security events across the enterprise by establishing a syslog server that collects events from different security devices. In time, they will want to consider taking this to the next step: reducing the amount of data that needs to be reviewed and correlate the meaningful events into actionable incidents within their enterprise. Since they have established a level of trust with Symantec, Symantec's Incident Manager product should be explored. In addition to their proven experience with the Symantec organization, their security products stand ready to integrate with Incident Manager. (See the Implemented Devices Table)

The goal of Symantec's Incident Manager version 3.x is to provide "a unified view of a company's information security issues and response activities, and helps organizations measure the effectiveness of existing security investments and proactively manage enterprise security." (Symantec Incident Manager Key Features). This product claims "near real-time correlation" that allows events from cross-vendor products to be correlated into incidents. It goes a step further by maximizing Symantec's acquisition of SecurityFocus and provides intelligence, including BugTraq information, on the events detected. If Symantec Vulnerability Assessment is installed, it also will correlate those events with the state of the system, indicating whether or not the system is vulnerable to the attempted attack.

SESA Architecture

Symantec Incident Manager is built on the Symantec Enterprise Security Architecture, which "allows multiple security products from disparate vendors to be 'plugged in' and report events in a common format to a centralized DataStore." (Symantec: Managing Security Incidents...) Through SESA, the end user does not need to switch between different vendor consoles to get an event stream which provides the schema normalization. Products based on the SESA Architecture "can be managed from a common, Web-based interface." (Roberts)

SESA is an open, multi-platform, standards based architecture. This architecture allows for the management of the security events from multiple point products, enabling a centralized view of the enterprise without needing to look at different consoles or different point product reports. It also enables consistent security event data collection and normalization, providing schema and message normalization, and then storing those events in a centrally managed database. Because SESA is an open architecture, third party product integration is achieved rapidly. In addition, the common architecture of SESA allows unified event information to be retained from cross-functional security realms, tiers and products. Symantec integrates information from firewalls, intrusion detection

systems, and anti-virus systems, as well as information across gateways, servers, and clients. (Symantec, Symantec Enterprise Security Architecture, p.2)

Further, SESA implements industry standards within their architecture, including:

- ✓ SSL (utilizing HTTPS on port 443 through anonymous SSL or full SSL means that SESA utilizes well-established ports and protocols which security administrators are familiar with)
- ✓ the CIM specification (The Common Information Model is language/method for describing management data and SESA incorporates the CIM schema to normalize management data)
- ✓ WBEM (Web-based Enterprise Management has been developed to unify enterprise computing environments, leveraging available Web technologies
- ✓ XML (eXtensible Markup Language is standard web language for exchanging information, and SESA standardizes its data into XML format

These are just a few of the standards listed in the Symantec Enterprise Security Architecture document. For more information on any of these, I refer you there.

SESA Architecture: event collection

The SESA Architecture is the basis for important components of the Incident Manager system. All integrating products report through a collector, or an agent.

> The SESA Agent is a Java application that runs at each managed endpoint to connect all Symantec Enterprise Security products to the SESA Manager. Symantec wrote the SESA Agent as an extension to an open-source implementation, which is based on industry standards.
> (Symantec, **Information about the components of Symantec Enterprise Security Architecture (SESA)**)

Event Managers provide an essential element to SESA. In the following figure, one can see how the variety of security sensors feed into the event managers within SESA. The event then travels the architecture, landing in the repository for the normalized data (the Datastore), and the console.

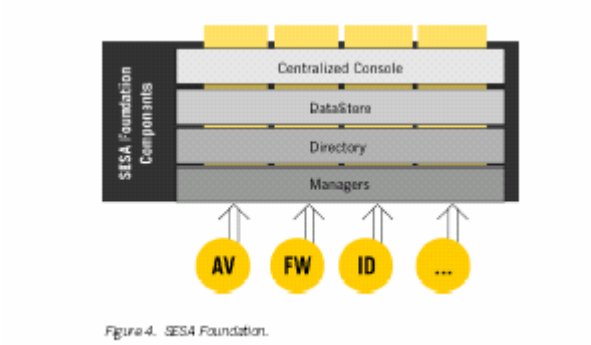

Figure 4. SESA Foundation.

(Symantec Information Security in the 21st Century, pg.10)

Symantec collectors also have the ability to do some level of filtering and aggregation of events on their own. For instance, in the Symantec Event Collector for Dragon documentation, the following is stated: the event collector "rule files contain standard rules that let the Event Collector perform translation and event filtering." (Symantec Event Collector for Enterasys Dragon, p.33)
This reduces the overall event stream and noise that makes it to the management console.

SESA Architecture: Datastore

SESA implements a robust, relational datastore that allows for querying and storage of event information in a highly regarded system. The initial SESA implementation of the SESA DataStore was IBM DB2, and future versions promise Oracle support as well. Further data integration includes the SESA Directory integrated into an LDAP Directory (IBM SecureWay). By using a SQL-based central point for the data collection, Symantec can leverage traditional database management tools, querying tools and data-mining techniques to analyze the data. (Dietz, page 10)

These implementations help to ensure scalability and manageability of the SESA datastore. Different point products will affect the datastore recommendations. For instance, the following recommendations can be found for SESA when implementing the Symantec AntiVirus Event Manager:

- Windows 2000 Server or Advanced Server with SP2
- 1 GHz Intel® Pentium® processor
- 1 GB minimum memory recommended
- 64 MB RAM per integrating SES product
- 80 MB disk space (program files)
- 1 GB disk space (SESA DataStore and SESA Directory program files)
- 128 GB disk space per product instance for SESA DataStore Data (for maintaining one month of data)
- 140 MB disk space for SESA Directory Data
- TCP/IP
- SSL-enabled

-Symantec Event Manager for AntiVirus System requirements

The SESA datastore is an integral part of the Incident Manager solution.

When originally introduced, SESA was established not as a product, but as a platform. It was described as "the rules of engagement" for how security products would report into an integrated view of the enterprise. (Sutton) For Incident Manager, SESA provides the basic architecture and building blocks for normalization, event collection, and the datastore.

Extending SESA's Functionality with Incident Manager

While SESA is the building block for the normalization, it deals with and delivers information on events within a continuous event stream from the point products. Basic data reduction techniques can be used, but the real benefit of implementing a Symantec solution will come in the form of Symantec Incident Manager.

Symantec Incident Manager separates Incidents from the event flow. Defining an Incident as "a set of one or more security events or conditions that require action and closure to maintain an acceptable risk profile," (Symantec, Managing Security Incidents... pg. 5) Incident Manager effectively separates into Incidents correlated events, narrowing the event stream that an analyst needs to evaluate. It then turns its focus to managing the life cycle of that incident. (Aberdeen Group, p. 11) While other tools can be used for sorting, filtering, thresholding, or aggregating the security events, Incident Manager can help answer the important question of why someone should care about the event they are looking at. (Managing it all, InfoWorld)

Correlation
Symantec's rule-based correlation includes logic "to detect known attacks and escalate anomalous activity". Utilizing a ruleset that includes pattern matching across devices and product groups, out-of-the box Incident Manager offers immediate benefit by narrowing the event stream to focus on what is critically important to the enterprise.

Symantec's field-based correlation and aggregation and filtering is enabled by leveraging the SESA implementation and schema normalization. By adding normalization down to the event level, Symantec's Incident Manager can leverage all correlation technologies referenced above. By maximizing the abilities of their relational database, an incident retains all the knowledge of the original security event that led to the incident, and that information does not require an analyst to switch views to access.

Symantec's context correlation includes the ability to escalate events representing activity against critical enterprise resources or even from known suspicious sources. Organizations can classify their systems and networks with profiles that enhance the system's view of the organizational security posture at any given time. (Symantec, Incident Manager - Real-time security incident management, p.2)

Leveraging their acquisition of SecurityFocus two years ago, Incident Manager has an important differentiator by offering to BugTraq information. The expert content available within the product includes action recommendations, vulnerability and safeguard data, as well as device-specific knowledge incorporated into the system. (Symantec, Real-time security incident management, p.3)

This information provides a wealth of information at the fingertips of the analyst: they do not need to reference additional resources or outside vendor documentation to assess quickly and accurately what is occurring on the system.

Additional features worth mentioning

Products that correlate on third-party products depend on the stability of those products to be able to offer consistent support. New event signatures that are applied, for instance, to a RealSecure solution need to be reflected within the correlation solution within a timely manner, especially in the face of the new attacks and emerging threats available on the Internet each day. As a result, the ability to implement these signatures across the board in a timely and trusted manner is vital. Symantec's implementation of LiveUpdate to its collectors as well as the SESA Manager and Incident Manager is an update mechanism with which customers are familiar and which they can configure to suit their own organizational policies. This makes the task of product updates more manageable across the board.

GIAC Implementation Recommendations

There are several things within the GIAC Architecture that makes them an ideal candidate for Symantec Incident Manager:

1) They have an established and trusting relationship with Symantec, are comfortable with the product solutions
2) Their distributed products are already integrated with SESA and Incident Manager
3) The building block of SESA normalization allows for different vendor's products to be "seen in the same light"
4) Asset identification within their environment can be applied to allow additional correlation capabilities
5) Their already deployed DB2 solutions imply a knowledge of DB2 and its management requirements
6) Their consolidated logging solution leads them naturally to implementation of the next step: finding out what these events mean. Incident Manager's implementation of an incident view reduces the amount of event data the analyst is faced with, and yet doesn't compromise the ability to drill into events as needed.

By leveraging these components of their established infrastructure, Incident Manager is an ideal candidate when it becomes time to evaluate a Correlation Technology solution.

At a glance: GIAC Implemented Security Devices

Device	Platform	Incident Manager	NetForensics*	ArcSight	neuSecure	e-security	
External F/W	Symantec	Y	Y	Y	Y	N	
Internal F/W	PIX	Y	Y	Y	Y	Y	
NIDS	Snort	Y	Y	Y (v. 1.8.1)	Y	N	
Client protection	Symantec Client Security	Y	N	N	N	N	
A/V	Symantec A/V	Y	Not native	Y	Y	Y	
Mail Security	Symantec Mail Security for Notes/Domino	Y	N	N	N	N	
Router	Cisco 1760	Y	Y	N	Y	Y	
DB implementation	IBM DB2	Y	N	Y	?	?	
Syslog support	Kiwi	Y	?	?	?	?	
OS Implementations	Solaris/Linux/Windows	Y	Y	Y	Y	Y	

Information on NetForensics from
http://www.netforensics.com/documents/pr_devices.asp

Information on ArcSight from
http://www.arcsight.com/product_platforms.htm,
http://www.arcsight.com/product_supported.htm

Information on neuSecure from http://www.guarded.net/supp.html

E-Security information from McIntyre's practical, "Event Correlation Systems..."
http://www.giac.org/practical/GSEC/Kevin_McIntyre_GSEC.pdf.

References

Aberdeen Group. "The Financial Value of Symantec's Security Solutions" http://www.wissensnavigator.com/documents/SymantecWhitePaper.pdf

ArcSight . http://www.arcsight.com/product_platforms.htm, http://www.arcsight.com/product_supported.htm

ARIN http://www.arin.net/

Bahadur, Gary & Shema, Mike. "Improving Apache" *Information Security Magazine,* April 2001: http://infosecuritymag.techtarget.com/articles/april01/features1_web_server_sec.shtml

Berners-Lee, Fielding, Frystyk. "RFC 1945: Hypertext Transfer Protocol -- HTTP/1.0" May 1996: http://ftp.ics.uci.edu/pub/ietf/http/rfc1945.html

Cisco. "Cisco Security Advisory: IOS HTTP Authorization Vulnerability " June 2001: http://www.cisco.com/warp/public/707/IOS-httplevel-pub.html

Cisco. "Cisco Security Advisory: Malformed SNMP Message-Handling Vulnerabilities" Dec 2003: http://www.cisco.com/warp/public/707/cisco-malformed-snmp-msgs-pub.shtml

Cisco. Cisco IOS Release 12.2 Documentation http://www.cisco.com/univercd/cc/td/doc/product/software/ios122

Cisco, "D through F Commands", *Cisco Pix Firewall Software,* http://www.cisco.com/en/US/products/sw/secursw/ps2120/products_command_reference_chapter09186a00801727a8.html

Cisco, "T through Z Commands", *Cisco Pix Firewall Software,* http://www.cisco.com/univercd/cc/td/doc/product/iaabu/pix/pix_62/cmdref/tz.htm#1026093

Degu, Christian & Bastien, Greg. "Getting Started with the Pix firewall", Cisco Press, April, 2003, http://www.ciscopress.com/articles/article.asp?p=31464

Dietz, Lawrence. "Information Security Management in the 21st Century." (c) 2003 http://www.isecurity.ru/technology/management.pdf

DMTF Security Protection and Management Working Group, http://www.dmtf.org/about/committees/spamWGCharter.pdf

Foundstone, SuperScan
http://www.foundstone.com/index.htm?subnav=resources/navigation.htm&subcontent=/resources/proddesc/scanline.htm

Internet Draft "Tunneling TCP based protocols through Web proxy servers", Aug 1998: http://www.web-cache.com/Writings/Internet-Drafts/draft-luotonen-web-proxy-tunneling-01.txt

Kiwi Syslog Information. http://www.kiwisyslog.com/info_syslog.htm

Leadston, David. "Enterprise Security Management – Reducing the Pain..." *SANS Practical,* Oct. 2003: http://www.giac.org/practical/GSEC/David_Leadston_GSEC.pdf.

Lonvik, C. "RFC 3164: The BSD Syslog Protocol" Aug. 2001: http://www.faqs.org/rfcs/rfc3164.html

Margulius, David. "Managing it all." *InfoWorld,* Jan. 2003: http://www.infoworld.com/article/03/01/10/030113fenexttci_1.html

McIntyre, Kevin. "Event Correlation Systems...." *SANS Practical* Feb. 2003: http://www.giac.org/practical/GSEC/Kevin_McIntyre_GSEC.pdf.

Mookhey, K.K. "Common Security Vulnerabilities in e-commerce systems" *Security Focus* April, 2004: http://www.securityfocus.com/infocus/1775

Muchow, Timothy. "Closing in on Intrusion Management." *SANS Practical* Nov. 2002: http://www.giac.org/practical/Timothy_Muchow_GSEC.doc.

N-Stealth, available from N-Stalker http://www.nstalker.com/products/nstealth/ http://www.nstalker.com/products/nstealth/features.php

NetForensics. http://www.netforensics.com/documents/pr_devices.asp

Network Working Group, "RFC 1034: DOMAIN NAMES - CONCEPTS AND FACILITIES" Nov. 1987: http://www.cse.ohio-state.edu/cgi-bin/rfc/rfc1034.html

neuSecure. http://www.guarded.net/supp.html

NSA. Router Security Configuration Guide. September, 2002. http://nsa1.www.conxion.com/cisco/guides/cis-2.pdf

NSA Router Security Checklist. "Executive Summary Card". http://nsa2.www.conxion.com/cisco/guides/cis-1.pdf

Numeroff, Laura Joffe. If You Give a Mouse a Cookie. Laura Geringer Book, 1985.

OSVDB (Open Source Vulnerability Database) http://www.osvdb.org/
Open Source Vulnerability Database. "Cisco Non-IOS Malformed SNMP Message-Handling DoS". OSVDB ID 3664: Jan. 2004: http://www.osvdb.org/displayvuln.php?osvdb_id=3664

Open Source Vulnerability Database. "Cisco IOS HTTP Unauthorized Administrative Access" OSVDB ID: 578 Jan 2004: http://www.osvdb.org/displayvuln.php?osvdb_id=578

PC Guide, RAID Level 5. http://www.pcguide.com/ref/hdd/perf/raid/levels/singleLevel5-c.html

Postel, Jonathan. "RFC 82: SIMPLE MAIL TRANSFER PROTOCOL" Aug. 1982: http://www.cis.ohio-state.edu/htbin/rfc/rfc0821.html

Poulsen, Kevin. "No Relief from Microsoft Phishing Bug" *SecurityFocus* 13 Jan 2004: http://www.securityfocus.com/news/7807

Poulsen, Kevin. "Report: Phishing Attacks on the Rise" *SecurityFocus* 19 Mar 2004: http://www.securityfocus.com/news/8289

Prunoiu, Florin. "Cisco Pix Firewall – Practical Guide," March, 2004, Enterastream Communications Inc. http://www.enterastream.com/whitepapers/cisco/pix/pix-practical-guide.pdf

Rekhter et al. "RFC 1918: Address Allocation for Private Internets" , Feb 1996: http://www.isi.edu/in-notes/rfc1918.txt

Roberts, Paul. "Symantec issues tool for vulnerabilities." *InfoWorld* May 2003: http://www.infoworld.com/article/03/05/12/HNsymantec_1.html.

Sanctum. "AppShield – Symantec Enterprise Security Architecture Solution Brief" , http://www.sanctuminc.com/pdf/AppShield_SymantecSESA_SolutionBrief.pdf

SecureScience Corporation. "Banking Scam Revealed" *SecurityFocus* 13 Nov 2003: http://www.securityfocus.com/infocus/1745

SecurityFocus. "Cisco IOS HTTP Configuration Arbitrary Administrative Access Vulnerability" BugTraq 2936 Mar 2004: http://www.securityfocus.com/bid/2936/info/

Schweitzer, Douglas. "The Enemy Within" , *ComputerWorld* April 22, 2004. http://www.computerworld.com/securitytopics/security/story/0,10801,92510p2,00.html

Sharan, Dhiraj. "IT Security: The need for a Cross-Correlation Platform." *SANS Practical* Mar. 2002: http://www.giac.org/practical/Dhiraj_Sharan_GSEC.doc.

Singer, Michael. "Voicemail Could Topple HP Merger" *SiliconValleyInternet.com.* Apr 2002: http://siliconvalley.internet.com/news/article.php/1008381

Skoudis, Ed. Counter Hack. New Jersey: Prentice Hall PTR, 2002.

Sutton, Neal. "Symantec promises better integration of security products - Symantec Enterprise Security Architecture**."** *Computing Canada* Oct. 2002: http://articles.findarticles.com/p/articles/mi_m0CGC/is_2002_Oct_11/ai_92826159

Symantec AV Center http://securityresponse.symantec.com/avcenter/refa.html

Symantec. *Event Collector for Dragon Guide*, http://www.symantec.com/region/jp/techsupp/enterprise/sesa/sec/pdf/sec_drag_ig.pdf

Symantec. "Event Manager for AntiVirus System requirements", http://enterprisesecurity.symantec.com/products/products.cfm?productID=172&pageID=1791&EID=0

Symantec. "How to configure Symantec Enterprise Firewall to block Nimda and Code Red worms" http://service1.symantec.com/SUPPORT/ent-gate.nsf/docid/2002092413323954?Open&src=ent_hot_de&docid=2002121711565354&nsf=ent-gate.nsf&view=docid/2002121711565354?opendocument&src=ent_hot_de&dtype=corp&prod=symantec%20enterprise%20firewall&ver=7.0.4%20for%20windows%20nt&dtype=corp&prod=Symantec%20Enterprise%20Firewall&ver=7.0.4%20for%20Windows%20NT/2000&osv=&osv_lvl=

Symantec, "How to set up site-to-site VPN tunnel between Symantec Enterprise Firewall and Cisco Pix", http://service1.symantec.com/SUPPORT/ent-gate.nsf/docid/2002022514174154?Open&src=&docid=2002121711565354&nsf=ent-gate.nsf&view=3fcd5fb2fcae709e88256bc1005cd7c9&dtype=&prod=&ver=&osv=&osv_lvl=

Symantec. "Information about the components of Symantec Enterprise Security Architecture (SESA)."*Symantec Support Website:* http://service1.symantec.com/SUPPORT/custserv-

ent.nsf/d13c85baed0cb50888256c09005ab232/1f8454e5bf08564988256cb7005a2cb9?OpenDocument&src=bar_sch_nam

Symantec. "Managing Security Incidents in the Enterprise," *Symantec Website 2002:* http://securityresponse.symantec.com/avcenter/reference/incident.manager.pdf

Symantec. "Symantec Enterprise Security Architecture," http://www.symantec.com.tw/region/mx/sym_en_lam/mexico/enterprise/10039545_SEntSecArch_fs.pdf

Symantec. "Symantec Enterprise Security Architecture Whitepaper," https://enterprisesecurity.symantec.com/Content/displaypdf.cfm?SSL=YES&PDFID=317

Symantec. "Symantec Incident Manager Key Features", *Symantec Enterprise Website:* http://enterprisesecurity.symantec.com/products/products.cfm?ProductID=166&EID=0

Symantec. "Symantec™ Incident Manager - Real-time security incident management for enterprise network environments" http://enterprisesecurity.symantec.com/content/displaypdf.cfm?pdfid=398

van Hauser. "Placing Backdoors Through Firewalls" v. 1.5, http://www.l0t3k.net/biblio/firewall/en/fireback.html

Vulnerability. http://www.zapthedingbat.com/security/ex01/vun1.htm

Yasin, Rutrell. "Security Overload." *Federal Computer Week* 12 Aug. 2002: http://www.fcw.com/fcw/articles/2002/0812/cov-sec-08-12-02.asp