

Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

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

An Exploration into Biometrics, Security Architecture Design, and Security Policies

Thomas Shepherd

GCFW Practical Version 4.1

© SANS Institute 2000 - 2005, Author retains full rights.

Date: April 9, 2005

© SANS Institute 2000 - 2005 Author retains full rights.

An Exploration into the Use of Biometrics 1

Abstract 1

Introduction 1
- Fingerprints 1
- Hand Geometry 2
- Voice Recognition 2
- Face Recognition 3
- Retina Scan 3
- Iris Scan 3
- Handwriting signature 4

Where Biometrics Fit in DID 4
- Biometrics as Identifier 5
- Biometrics as Authenticator 5

Biometrics and Security Administrators 6
- Examples 8

Biometric Innovations 9

Conclusion 10

Security Architecture for GIAC Enterprises 11

Assumptions 11

Network Security Design 12
- Seattle and Miami 12
- London and Sydney 13
- Denver 15

Group Access Requirements 16
- Public 17
- Customers 17
- Suppliers 17
- Partners 18
- Remote Users 18
- Internal Users 19
- Intersite Communications 19

Equipment List 20
- Border Routers 20
- VPN Concentrators 20

© SANS Institute 2000 - 2005, Author retains full rights.

Firewalls 21
IDS Sensors 21

IP Addressing Scheme 22

Conclusion 23

GIAC Enterprises Router and Firewall Security 24

General Security Stance 24

Filtering Router Policy 24

Access Tables 25

Primary Firewall Policy 28
General Firewall Configuration 28
Ruleset on the Public (Outside) Interface 29
Ruleset on the Internet Service Network Interface 29
Ruleset on the Intranet Service Network Interface 30
Ruleset on the Internal LAN Interface 31

Conclusion 33

Bibliography 34

© SANS Institute 2000 - 2005, Author retains full rights.

© SANS Institute 2000 - 2005 Author retains full rights.

An Exploration into the Use of Biometrics

Abstract

What are biometrics? Biometrics are unique physical characteristics or traits of a person, such as fingerprints, voice, and eyes. These biometrics are used to both identify a person and to authorize them access past some security measure. Security Administrators need to pick the right security solution to ensure that people use the security and that the solution provides the needed security application. While biometrics are not new, the use of them for authentication is, and advances are being made all the time. Biometrics – it's the future of security in the IT industry.

Introduction

Security administrators walk a fine line between open-access for users, so that they can accomplish their job, and locking down a network infrastructure. A security administrator can control many aspects of security, such as physical security and loaded server software. Factors that are rarely under the control of security administrators are users, especially when it comes to passwords. Simple passwords, in today's computing environment, are easy to crack. Security administrators must therefore make it harder to crack a password or obtain illegal access. They do this by making passwords expire more quickly, making passwords longer, or by requiring strong passwords. The more security administrators increase the complexity of password requirements, the more likely it is that users will forget their password, or simply write it down. Another way that security administrators have sought to solve this issue is through the use of some third-party hardware, such as certificate keys/USB tokens or smart cards. While this can be effective, it requires the user to carry something tangible with them that can be damaged, lost, or left at home. An effective solution to this problem is to authenticate users based on something unique that they will not forget, give away, or leave at home on their dresser. This is accomplished through the use of personally identifiable physical characteristics or traits called biometrics.

Haircuts, clothing styles, body builds, and even general appearance can be imitated. People who are good at imitation can even approximate speaking like another person. While people may look alike or act alike, as in the case of twins, actors, etc., there are still unique aspects of a person's physiology or behavior that makes them an individual. Biometrics attempt to utilize these unique aspects of a person's physiology or behavior to either identify who they are or to authenticate that someone is who they say they are. The most common biometric measures include fingerprints, hand geometry, voice recognition, face recognition, retina scan, iris scan, and handwriting signature; however, any unique aspect of physiology or behavior can be utilized.

Fingerprints

Fingerprints, which many say was the first biometric measure, have long been used as a method of authenticating that a person is who they say they are.

© SANS Institute 2000 - 2005 Author retains full rights.

Fingerprinting was codified in the 1800s, which made classification and searching of fingerprints much easier and faster. Fingerprints were broken down into 10 distinct categories, which made fingerprint matching easier. Recently it has been recognized that within a fingerprint, there exist "minutiae" that further distinguish one fingerprint from another that is similar.

The ridge patterns of a finger are imaged to produce a fingerprint scan. Fingerprint matching applications originally attempted to match a picture of one fingerprint to a picture of another fingerprint. With a small database of fingerprints, this was acceptable and usually produced results. As the database of fingerprints grows, the more likely it is that fingerprints will be similar in nature and harder to find a match. For this reason, a new method of matching fingerprints was developed. New algorithms were developed that store a numeric value for a fingerprint rather than a picture of the print. Fingerprint records became relatively small and are typically between 512 to 1000 bytes. This makes it easier to search a database to match a fingerprint.

Hand Geometry

Hand geometry is similar to fingerprints in that it attempts to match characteristics of the hand to a known database; however, it is not as complex or comprehensive, which means that it is not as accurate. Many people have adopted it because of the speed of implementation, both entering new users and matching existing users. These applications measure the geometry of the hand, for example, the length of the fingers, the width of the palm, etc., and then store a numeric value for this information.

Hand geometry records are typically 9 bytes. While this makes it easy to store many records, it makes it hard to split the database into small, easily searchable chunks of data. Also, with only 9 bytes of data, it is easy to get false negatives and false positives.

Voice Recognition

Many people often mistake voice recognition with speech recognition. Speech recognition is the application of converting speech to words, and is often used in word processing applications. Voice recognition looks at aspects of a persons voice and attempts to match it with a known person in the database. It is typically not important what a person says, but how they say it that allows the application to make a match. A user is often given a particular phrase to say so that results will be more consistent, and thus easier to match.

While voice recognition is fairly robust, it is not without its problems. Certain situations can affect a person's voice. Depending upon the age of the person, puberty will dramatically alter a person's voice. It is also found that speaking in a different language will often produce irregular results. Another situation that makes voice recognition difficult is when a person has a cold or other ailment that affects the nose or throat. Most commercial applications will not do well in these situations, but much research is being done to improve this technology.

© SANS Institute 2000 - 2005 Author retains full rights.

Face Recognition

People have always been able to identify an individual by their face, which is why criminals always try to hide theirs and law enforcement uses sketch artists to attempt to draw a face. Using how a person's face looks is an indispensable way to identify them, this method of identification and authentication is a relatively new biometric. Until recently it was thought that a computer could not differentiate well enough to tell people apart; however, newer algorithms have shown that it is possible. The algorithm identifies and measures various facial features and then stores a numeric value identifying that face.

Like fingerprinting, face recognition does not store a picture of a person's face, just a numeric value. The distance between the eyes, the placement of the eyes in relation to a person's nose, and the width of a person's mouth are all examples of face recognition metrics. Although it is getting better, facial recognition is hampered by several factors, including lighting, camera angle, facial expression, and simple aging. Recent tests have shown that sometimes it is possible to use a photograph of a person to fool this type of measurement.

The typical record size for face recognition is 1300 bytes, which makes it easy to break the database up into smaller chunks to search.

Retina Scan

During a retina scan, a user places their eye close to the biometric sensor. The sensor then uses an infrared light to heat up the retina at the back of the eye. The blood vessels of the eye will heat up faster than the rest of the eye and will show up. The sensor then takes a scan of the pattern of the blood vessels and attempts to match distinguishing characteristics of the pattern.

This method, although shown to be highly reliable, is the least well liked by users. Users typically argue that it is difficult to use, since you have to put your eye close to the sensor and hold still. Users say that it is too invasive because coherent light is being shone directly into the eye. Many also express concerns that extended use may cause damage to the eye. Most users do not want to go through this trouble or take the time to use this method. Since this method requires a user to hold their eye still and, depending upon the application, to not blink, this method also has the highest error rate.

The typical record size for a retina scan is 35 bytes, which makes it easy to store a large number of scans, but is not as well suited for breaking the database into small searchable chunks.

Iris Scan

Since it does not have to reach the back of the eye, an iris scan is not as intrusive as a retina scan. An iris scan also does not require that a coherent light source enter the eye. For this reason, users are more open to iris scans. An iris scanner takes a picture of the iris from approximately 18 to 24 inches away. It then

© SANS Institute 2000 - 2005 Author retains full rights.

attempts to match distinguishing characteristics of the iris. Like fingerprints, an iris changes very little over the course of a person's lifetime.

Iris scans are not without their limitations. Iris scans may be affected by the amount of light entering the eye. If the light entering the eyes is brighter, the pupils contract, thus increasing the area of the iris that can be scanned. If the light entering the eyes is dimmer, the pupils dilate, which shows less of the iris, thus reducing the area of the iris that can be scanned. In order to obtain the best results, the amount of light shining into the eyes must be regulated. Certain eye ailments, particularly eye infections, can also make it hard to image the iris. For example, ailments can alter the eye either through changing the amount of pupil dilation or by covering the eye with a film.

The typical record size for iris scans is between 256 to 512 bytes. This makes it ideal for storing a large number of scans while still being able to break the database up into easily searchable chunks.

Handwriting signature

Like voice recognition, the handwriting signature is often mistakenly described. It is not, as many think, the comparison of one's signed name to a scan of a previously signed name or the comparison of letters after writing has finished. The handwriting signature biometric measures the process of handwriting while a person is writing. Similar to voice recognition, where it does not matter what is said, it does not matter what is written in the handwriting signature biometric. The sensor measures the manner of writing and compares uniquely identifiable features of the handwriting process.

Having someone sign their name is the most common method because a person is accustomed to writing their name, and it can be used as a legal signature for authentication. However, it has been demonstrated that the manner in which a person writes may change depending upon their mood. The more relaxed that a person is, the more their writing tends to loop and flow. The more stressed that a person is, the more their writing tends to compress and have sharp edges and points. Another factor that may limit handwriting signatures is physical ailments, particularly ones that affect the hands or fingers, such as carpal tunnel syndrome or arthritis.

Where Biometrics Fit in DID

It used to be that a person was tied to the computer at their desk at work, then came telecommuting, where a person was able to work on a company network from home. That is all changing now. There is a growing trend toward "M-Commerce" or a more mobile workforce. Users are no longer limited to a particular location with restricted access. Users now use wireless laptops, PDAs, cell phones, etc., to access company information. Company data is being stored and distributed through a variety of methods, including a recent growth in the use of micro drives or thumb drives. In fact, users are now demanding access to IT resources from anywhere at anytime and through a variety of access methods.

A network's defensive perimeter used to be defined by its firewall and physical security. With a more mobile workforce comes a more complex security infrastructure.

© SANS Institute 2000 - 2005 Author retains full rights.

A network defensive perimeter must now become an IT defensive perimeter. Defenses must be defined in terms of Defense-In-Depth.

> Defense in Depth is one of the overriding principles of information security, allowing layered security to capitalize on the respective strengths of each component while being flexible enough to choose components based on technical, budgetary, and political constraints.[1]

Security solutions must adjust to this growth in complexity. One possible solution is the use of biometrics.

Biometrics are usually deployed in one of two ways. The first way is to identify a person, similar to or in conjunction with a security ID badge. The biometric establishes a verifiable method of confirming that a person is who they say they are. It also adds authenticity to the claim based upon the fact that a person's biometric cannot be loaned to someone else and is nearly impossible to duplicate. The second way in which biometrics are being deployed builds upon the first way. Once a person's identity has been established, they can then be authenticated and granted some level of access. Typically access is not granted based solely upon biometric identification, but rather in conjunction with some other form of access verification, such as a password or smart card.

Biometrics as Identifier

When used to identify a person, biometrics are most often used in physical security; for example, allowing security personnel to match a person's ID photo to a photo in the database. This is especially useful for large companies where security personnel cannot know every employee and do not need to know what security access a person needs. They simply need to know that that person works for the company or organization and has the right to be there. Identity can be confirmed or denied in a manner that leaves little room for doubt.

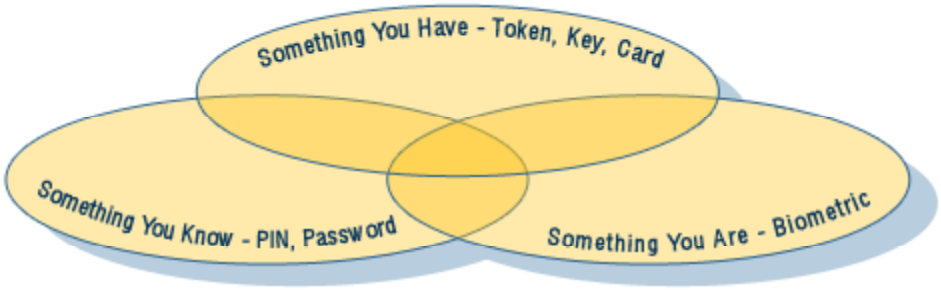
Other than being used by security personnel, biometrics can be used for additional purposes. Since they are unique to each person and a person cannot forget or lose them, they can be used to track attendance or in timesheet applications. Since it is relatively easy to reproduce the results and the results uniquely identify a person, they can also be used as digital signatures. As a digital signature, anywhere that a written signature could be used a digital signature could be used, such as financial transactions. And finally, there are a myriad of solutions for law enforcement personnel, from IDs to fingerprinting to criminal investigations.

Biometrics as Authenticator

Biometrics used to authenticate a person is most often an automated process. The person's identity is established through the biometric and then the appropriate security access is granted. For example, once a person's identity has been established, a door lock or other physical lock is released and the person is allowed to enter. Depending upon the level of security needed, authentication may involve more

© SANS Institute 2000 - 2005 Author retains full rights.

than one biometric or method of establishing identity. More often, though, identity alone is not enough to fully authenticate a person. Biometrics are considered what or who you are. When this is not enough, authentication relies upon something else, like what you have (a smart card) or what you know (a password or PIN) as illustrated in the following image[2].



Combinations of these approaches further heighten security.

A growing trend in the IT industry is the use of biometrics for remote user or non-network authentication. There are various applications of this type of authentication. Typically it is used to authenticate a person logging onto a laptop. It may also be used to authenticate a remote access user on a network. For those who are more security conscious, biometrics can also be used to authenticate users on PDAs, cell phones, and even micro drives. Some biometric solutions require authentication against a central repository of biometric information. And, given that biometric signatures are not all that large, some solutions store the authentication information locally with the scanner.

Biometrics and Security Administrators

The study of biometrics is difficult and highly complex; thankfully, the application of biometrics is not as problematic. The application of biometrics does not require one to fully understand the inner workings of what makes a biometric measure unique and usable. Most biometric solution companies have made it easy for a security administrator to deploy an effective solution. A security administrator needs only to select the appropriate biometric measure and then research possible solutions. For security administrators, the best part about a biometric measure is that a user can't lose it, share it, or loan it to anyone else. This will also help with accountability and audit trails.

Security administrators attempt to be vigilant against two distinct security issues; attacks attempting to bypass or get around security measures and attacks attempting to get through security measures with known information. Biometrics can help with both of these issues. First, biometric scanners can be embedded within hardware, or in the case of computers, within the BIOS, thus making it harder to bypass. Depending upon the application, a biometric scanner may need to know only a few biometric signatures, which can also be embedded with the scanner. Secondly, passwords may be observed, sniffed, or cracked, and a smart card or ID card may be imitated, duplicated, or stolen. It is much harder to intercept and decode a biometric

© SANS Institute 2000 - 2005 Author retains full rights.

scan, if biometric information is passed along the wire at all, and, as stated earlier, biometrics cannot be lost, stolen, or given away.

This should not be interpreted to mean that biometrics are infallible. They are not. Depending upon the sophistication of the solution, it has been shown that a recorded voice can fool some voice recognition scanners. Still others have shown that photographs can fool face recognition scanners. Latex fingers have been used to fool fingerprint scanners; although most fingerprint scanners will now detect these. A relatively recent development called Gummy Fingers, however, has shown that a person's fingerprint can be lifted from some other surface, including the fingerprint scanner itself, and used to fool a scanner. The only biometrics that are relatively difficult to duplicate or fool a scanner are the retina and iris scans.

When deciding upon a biometric solution, a security administrator must also consider the negative aspects of possible solutions. Depending upon the biometric chosen, there might be physical, cultural, or social reasons as to why the biometric measure may not be able to be deployed. In certain situations of blindness, a retina or iris solution may not be usable. If an employee is mute, a voice recognition solution would not be usable. There are some cultures where a person's face must remain covered; in this situation a face recognition solution may not be usable. In those situations were a person has not learned how to write, a handwriting recognition solution may not be usable. Many users will also object to any biometric that is too intrusive, such as the retina scan. These types of issues must be considered by anyone looking to deploy a biometric solution.

The following table[3] shows a comparison of the various biometric solutions.

Table 1. Comparison of biometrics

Characteristic	Fingerprints	Hand Geometry	Retina	Iris	Face	Signature	Voice
Ease of Use	High	High	Low	Medium	Medium	High	High
Error incidence	Dryness, dirt, age	Hand injury, age	Glasses	Poor Lighting	Lighting, age, glasses, hair	Changing signatures	Noise, colds, weather
Accuracy	High	High	Very High	Very High	High	High	High
Cost	*	*	*	*	*	*	*
User acceptance	Medium	Medium	Medium	Medium	Medium	Medium	High
Required security level	High	Medium	High	Very High	Medium	Medium	Medium
Long-term stability	High	Medium	High	High	Medium	Medium	Medium

***The large number of factors involved makes a simple cost comparison impractical.**

There are varying levels of biometric solutions depending upon how much money you want to spend and how much security is needed. A company can spend a great deal of money on a biometric solution and not be any better off if all they need is a simple fingerprint reader. At the same time, a company that needs a great deal of security will need to spend more money on a more complex solution, possibly involving more than one biometric measure. A security administrator should examine each biometric to determine which one is both cost effective and provides enough security.

© SANS Institute 2000 - 2005 Author retains full rights.

Examples

Let's look at an example. The ABC Company does not currently have any biometric infrastructure. The CEO wants to provide more security for field sales staff. He is worried that company information may be stolen off of the staff's laptops. He also wants greater security for field staff that is connecting back to the corporate LAN via VPN. At the same time, he wants it to be as simple and cost effective, i.e. low cost, as possible.

Being field staff, any solution that is chosen must be relatively portable. This would most likely rule out hand geometry and retina scans as these solutions are large and bulky. It is also not likely that a handwriting solution would be viable. The CEO also wanted a low cost solution. This would most likely rule out face recognition and iris scans as these solutions can be costly. Lastly, security must be provided for both logging onto the laptop and logging onto the corporate LAN. The possible solutions could include one or both of fingerprints and voice recognition. Once the solutions have been narrowed, further examination is required. Voice recognition will most likely be software based, which means that a person will already have access to the laptop with the software running. Also, voice recognition was one of the harder to manage solutions.

Fingerprints seem to be the best solution for this situation. It could be recommended that USB fingerprint readers be purchased for all field staff. This way they can use it to log onto both the laptop and the corporate LAN. And, if there is a greater security concern, logging onto the corporate LAN can require re-authentication. As laptops become outdated, they can be replaced by newer laptops that have integrated fingerprint readers.

Now let's look at another example. The XYZ Corporation is extremely security conscious and has already implemented a fixed place biometric solution of retina scans for access to sensitive areas in corporate buildings. The corporation wants to update existing biometric infrastructure to be more secure, and to use a biometric solution for logging into the corporate LAN that can be extended to field staff. Corporate executives are also concerned about any device that might have corporate information on it. They are also sensitive to the fact that users are reluctant to change and don't like the current retina scan system; they feel that it is too intrusive. The project must be cost effective, but security is more important than cost.

First, let's address the existing biometric infrastructure. Since the corporate executives want to improve the user's experience by making the biometric measure less intrusive and easier to use, the retina scan system needs to be replaced. Handwriting signatures are probably not a good solution considering that a person's writing changes with their mood. Face recognition can probably also be dismissed; the technology has not matured enough yet to provide adequate security. This leaves fingerprints, hand geometry, voice recognition, and iris scans as possible solutions.

Next, the replacement system must be more secure than the current system. Simply replacing one biometric system for another does not guarantee improved security. Since the users are used to retina scans, it would be simplest to move to iris scans, which are just as reliable. This does not, however, increase security, so it

© SANS Institute 2000 - 2005 Author retains full rights.

would be prudent to add an additional security measure. This could be as simple as a PIN or password, a proximity card, or a smart card, or it could be an additional biometric measure.

As the situation is further examined, biometric solutions may be combined in order to limit how many different types are used. Secure office space access and LAN authentication could use the same biometric solution. Given that LAN authentication must be extended to field staff, the previous example gave us the ideal solution, which is fingerprint biometrics. This solution can also be incorporated with the iris scanning solution above to provide the increased security that is desired.

The corporate executives also expressed concern about any device that may contain corporate information. Fingerprint solutions can be integrated with these devices, and with the corporate biometric infrastructure, to provide the necessary security. Cell phones and PDAs are available which have integrated fingerprint biometric solutions. Laptops with integrated fingerprint scanners can be purchased. In fact, solutions are now being developed to integrate fingerprint scanners with USB micro drives.

Each security solution must be customized to fit the demands of the situation. A poorly devised solution may not only prove to be insecure, but may be costly to replace and complex to manage. Solutions should be picked based upon the level of security needed, the reliability of the solution, the amount of time it takes to enter a new user, and the time it takes to determine a match. User preferences, including social and cultural preferences, should also be considered when evaluating solutions. It is up to the security administrator or security team to balance these issues and pick the best possible solution.

Biometric Innovations

Biometrics have come a long way since their inception. Researchers are finding that there are other parts of the body that are potential biometric measures; for example, tongue-prints. It is said that the arrangement of taste buds on a person's tongue are as unique as fingerprints. The FBI is working on a biometric that measures the way that you walk, asserting that everyone has unique aspects of their walking gait. In fact, many biometrics have replaced genetic profiling as a means to definitively identify a person.

NIST has developed a standard that should enhance the industry's ability to adopt and deploy biometric solutions. The standard is called the "Common Biometrics Exchange Formats Framework (CBEFF)."

> CBEFF describes a set of data elements necessary to support biometric technologies in a common way independently of the application and the domain of use (e.g., mobile devices, smart cards, protection of digital data, biometric data storage). CBEFF facilitates biometric data interchange between different system components or between systems, promotes interoperability of biometric-based application programs and systems, provides forward compatibility for technology improvements, and simplifies the software and hardware integration process.[4]

At the same time, ANSI has adopted the BioAPI specification for its standard.

© SANS Institute 2000 - 2005 Author retains full rights.

> This specification defines the Application Programming Interface and Service Provider Interface for a standard biometric technology interface. BioAPI V1.1 defines an open system standard API that allows software applications to communicate with a broad range of biometric technologies in a common way. As an "open systems" specification, the BioAPI is intended for use across a broad spectrum of computing environments to insure cross-platform support. It is beyond the scope of this specification to define security requirements for biometric applications and service providers, although some related information is included by way of explanation of how the API is intended to support good security practices. BioAPI V1.1 was developed by the BioAPI Consortium which consists of eighty organizations representing biometric vendors, Original Equipment Manufacturers (OEMs), major Information Technology (IT) corporations, systems integrators, application developers, and end-users. NIST holds membership in the Consortium and is a member of the Steering Committee. BioAPI specifies standard functions and a biometric data format which is an instantiation of CBEFF.[5]

The BioAPI is listed as a more specific instance of the CBEFF standard. These standards should make it easier to deploy a biometric infrastructure and upgrade it without tearing out the old infrastructure.

The Optel Company is developing another innovation in biometrics. They are engineering a new biometric solution that utilizes sonic holography technology to improve the security of fingerprint biometrics. One important, and largely beneficial aspect of this solution is that it is a contact-less solution. This means that an intruder can no longer simply lift a fingerprint due to residue left on a fingerprint scanner. The other way that this is an improvement, for those that are extremely security conscious, is that the same technology can be used to determine whether or not the scan is from an "alive" finger, rather than dead or an imitation.

Conclusion

Many people think of biometrics as the ultra-secure solution that costs thousands of dollars and is extremely complex and difficult to manage. Well, it can be that way, but it doesn't have to be. Quite a bit of progress has been made to make biometric solutions available to the mainstream business market. Several large computer manufacturers are beginning to integrate biometric scanners into their products, making biometric solutions cheaper and available to the average person. The possible solutions can be complex, but they can also be as simple as purchasing a USB fingerprint reader. As the industry converges on and implements the new biometric standards, many more possible solutions will become available and become more reliable. In order to be truly effective, security administrators need to stay current with biometric developments and keep their organization's biometric infrastructure up to date. Biometrics are quickly gaining ground as the next step in IT security evolution.

© SANS Institute 2000 - 2005 Author retains full rights.

Security Architecture for GIAC Enterprises

Some security architectures are simple; some are complex. In all instances, it is highly practical to take the time to properly design the security in the first place, rather than to throw something together and then plug all the holes after something goes wrong. Security architectures are combinations of many different pieces of security – hardware, software, policies, procedures, and standards – all working together to compliment each other and to provide multiple layers of security. If any piece of the architecture is ignored, then the entire architecture is at risk of failure. No single architectural design applies to every situation, and a situation does not have a single correct architectural design. Obtaining a good, workable security architecture is a process.

For GIAC Enterprises and this situation, to further the architectural design, the process starts by listing the assumptions that have been made. Next, the physical and logical network designs will be discussed. Then, the details of Group Access Requirements will be laid out, including access tables. Finally, the recommended hardware necessary to complete the proposed architecture will be listed and discussed. This process should result in a viable, well-rounded security solution.

Assumptions

In order to properly address the needs of GIAC Enterprises, certain assumptions must be stated up front. Normally these would be dealt with during meetings with the customer, but that is not feasible at this time. These assumptions are necessary to further not only the security design, but to also prompt discussion on security issues that relate to a global corporation rather than a small locally owned business. While it can be said that any company doing business on the Internet is in reality a global company, since the Internet is global in nature, the issue here is more one of having physical office space distributed in a global nature.

For this design, the corporate offices are located in Denver, Colorado. There are also listed four satellite offices located worldwide, which will be located in Seattle, Washington; Miami, Florida; London, England; and Sydney, Australia. These locations were picked for two reasons. First, since the company has partners that translate and sell their fortune cookie sayings and the corporate office is located in the USA, it is assumed that the company offices are located in English speaking regions. Second, knowing the location of the offices will become important to the discussion about firewalls, VPNs, and secure traffic.

It is assumed that the company does not outsource its sales operations. This is an important distinction to make. If the company outsourced its sales operations, then the security architecture would have to change to reflect a sales partner, including service level agreements (SLA). With sales being handled within the company, the security architecture need only reflect security for sales information within the company. However, since sales are handled within the company, this architecture will reflect the fact that there is a desire to localize sales information. To do this, sales will be conducted via the Internet through three offices, the offices in Denver, London, and Sydney.

© SANS Institute 2000 - 2005 Author retains full rights.

It is assumed that the company does not outsource its security operations. This is also an important distinction to make. If the company outsourced its security operations, then the security architecture would have to be altered to allow for remote monitoring, and possibly remote control, of security infrastructure. Since staff internal to the company will handle security, security information does not need to leave the company network, nor does an outside source need access to security information or equipment.

Network Security Design

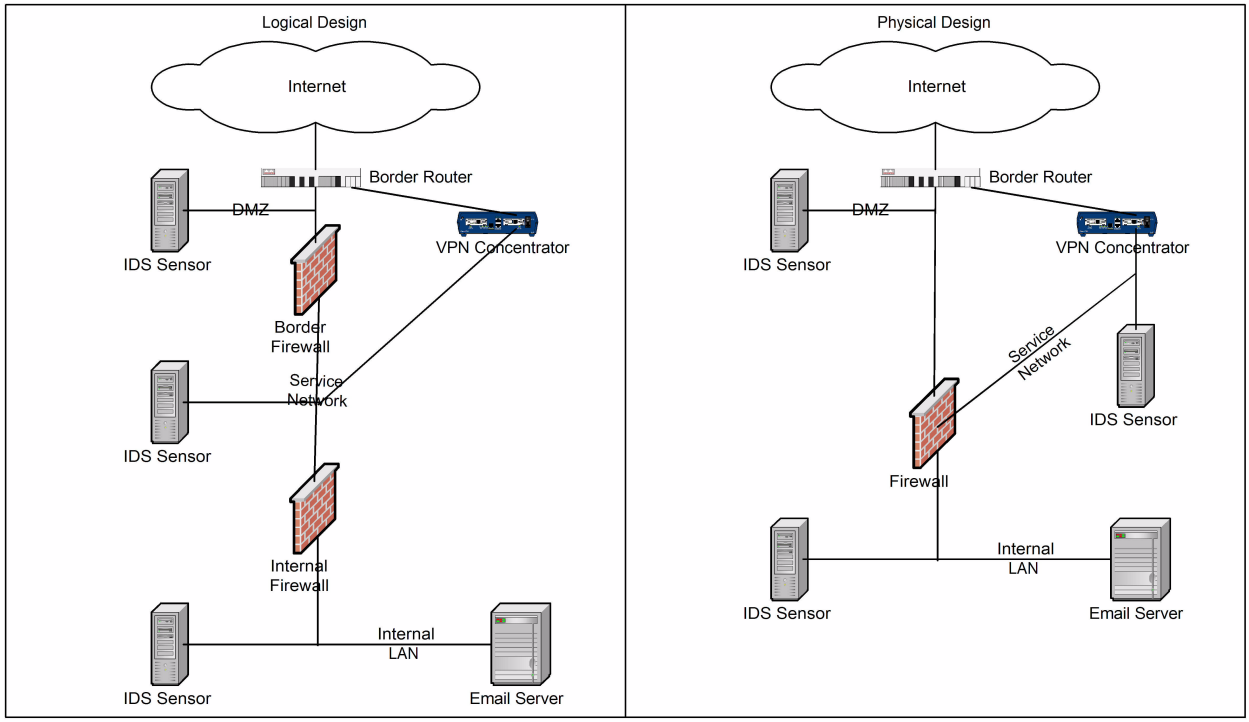
The network security architecture can be broken down into three distinct designs: satellite offices that do not host sales infrastructure – the Seattle and Miami offices; satellite offices that host sales infrastructure – the London and Sydney offices; and the corporate office in Denver. While these designs may have similar features, they each have different security requirements. These designs also do not take into account issues such as redundant data links, offsite backups/mirroring, round robin DNS, etc. as these are issues for proper network design and not security and Defense-In-Depth.

Seattle and Miami

The Seattle and Miami offices host only field sales staff, which means that other than email there is no reason for anyone outside of the company to require access to their network infrastructure. The design may be described as follows. Each office will have a border router connecting the office to the Internet. Connected to the border router DMZ are an IDS sensor, a VPN concentrator, and a border firewall. Inside the border firewall is a service network where the VPN tunnels will terminate. Connected to the service network is an IDS sensor and the internal firewall. Inside the internal firewall is the LAN infrastructure, including an IDS sensor/server and a mail server.

© SANS Institute 2000 - 2005 Author retains full rights.

© SANS Institute 2000 - 2005 Author retains full rights.



Through this design, known bad traffic can be filtered at the border router, providing the first line of defense. Traffic that passes the border router will then have to pass the border firewall. The reasoning behind having a border firewall and a service network is that many firewalls do not adequately check encrypted VPN traffic. By terminating VPN connections within the service network and not the internal LAN, all VPN traffic can then be decrypted and scanned by the internal firewall.

The IDS sensors are deployed to give maximum information for security personnel. It is possible, but not as effective, to get by with only a single internal IDS sensor/server; however, if the IDS is only applied internally, the only alerts that you will get will be after the traffic has already entered the network and potentially after the damage has already occurred. A single IDS sensor/server deployed externally would also be limited. The sensor may generate more alerts, but there would be no easy way to tell whether or not any of the traffic made it past the internal firewall. By placing sensors on each segment of the security architecture a security analyst will be able match particular alerts to determine to what level the security is being penetrated.

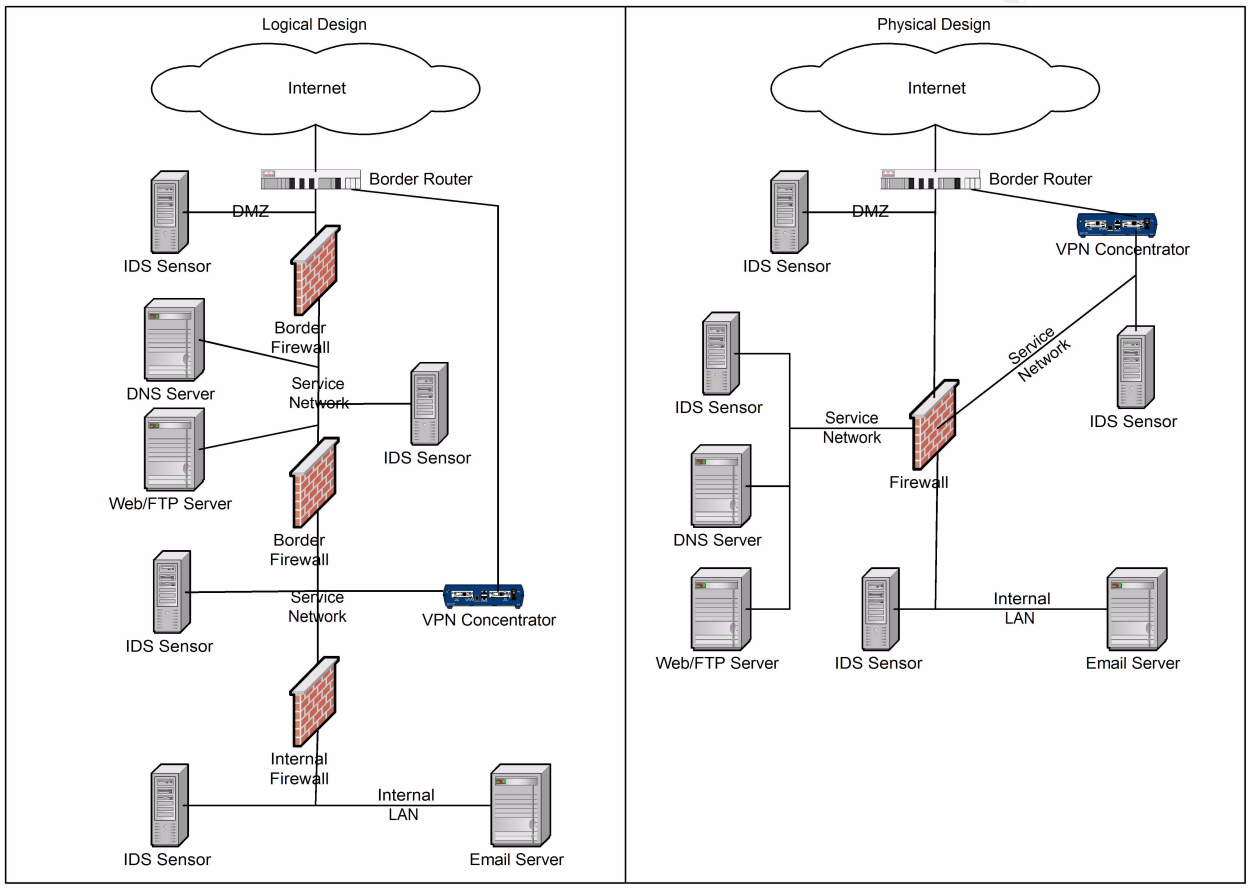
Outbound traffic security is not ignored by this design. Either or both of the firewalls may be configured to provide NAT services for outbound traffic. The firewalls may also be configured to block particular traffic to the Internet while allowing traffic to the other company offices. This is particularly helpful for the company's site-to-site VPNs.

London and Sydney

The London and Sydney offices host not only field sales staff, but also sales infrastructure. This means that besides the standard company infrastructure each office will also host web services. The design may be described as follows. Each office will have a border router connecting the office to the Internet. Connected to the border router DMZ are an IDS sensor, a VPN concentrator, and a border firewall.

© SANS Institute 2000 - 2005 Author retains full rights.

Inside the border firewall is a service network, which is designated as the Internet Service Network. Connected to the Internet Service Network are an IDS sensor, a DNS server, a web server, an FTP server, and a buffer firewall. Inside the buffer firewall is another service network, which is designated as the Intranet Service Network. Connected to the Intranet Service Network are an IDS sensor and the internal firewall. This service network is also where the VPN tunnels will terminate. Inside the internal firewall is the LAN infrastructure, including an IDS sensor/server and a mail server.



Similar to the first design, the border router is the first line of defense and will filter known bad traffic. In this design, more traffic will be allowed past the border router to support sales operations, but the traffic will still have to pass the border firewall.

The addition of the Internet Service Network provides a more secure environment for systems that will be accessible by users on the Internet. It also adds another layer of security for internal infrastructure. By allowing people outside of the company farther into the security architecture, the security has become weaker. The addition of the service network enhances the security to counter the effects of granting people greater access.

The VPN connections are still terminated within a service network in order to provide the security described in the previous design. They are not terminated, however, in the Internet Service Network by design. This is done to protect the integrity

© SANS Institute 2000 - 2005 Author retains full rights.

of the VPN traffic tunnels. Terminating them in the Internet Service Network would leave the entire company VPN system open to attack by a compromised web server. By terminating them one level down a compromised web server would not be able to utilize any of the VPN services to bypass company security and attack other company systems.

Once again, the IDS sensors are deployed to give maximum information for security personnel. In this situation it is not possible to deploy a single IDS sensor/server, either internal or external, and have the alerts be effective. This design necessitates that at a minimum both the Internet Service Network and the internal LAN need to be monitored. The Internet Service Network needs to be monitored due to the potential for compromise due to the large number of people that will access services there. And, the internal LAN needs to be monitored due to the sensitive nature of the information stored there. Again, like the first design, by placing sensors on each segment of the security architecture a security analyst will be able match particular alerts to determine to what level the security is being penetrated.

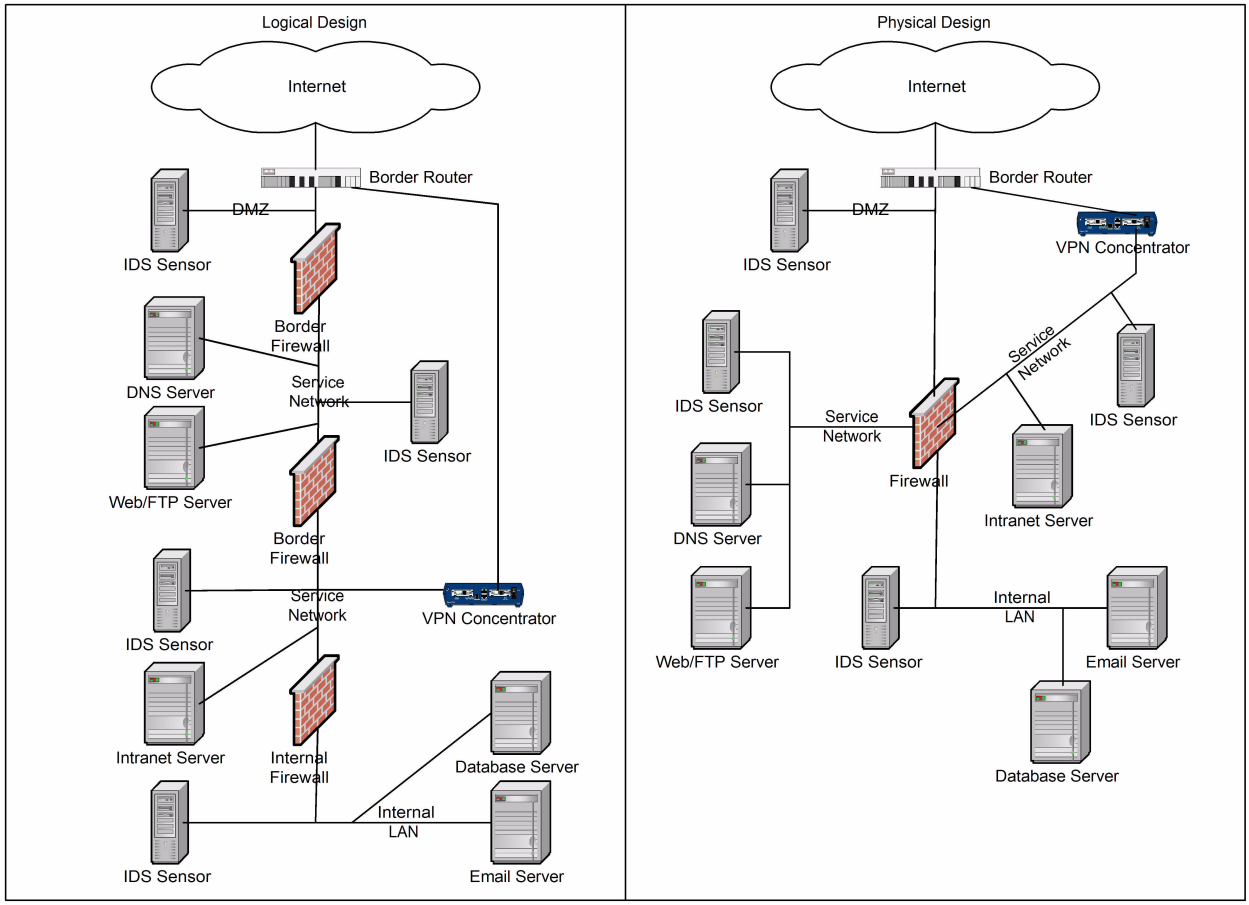
Outbound traffic security is again handled as it is in the first design. NAT services for outbound traffic are configured in order to provide security for internal equipment.

Denver

The corporate office in Denver has the same core infrastructure of the field sales offices with the addition of database, accounting, and human resources services. The design may be described as follows. There is a border router connecting the office to the Internet. Connected to the border router DMZ are an IDS sensor, a VPN concentrator, and a border firewall. Inside the border firewall is a service network, which is designated as the Internet Service Network. Connected to the Internet Service Network are an IDS sensor, a DNS server, a web server, an FTP server, and a buffer firewall. Inside the buffer firewall is another service network, which is designated as the Intranet Service Network. Connected to the Intranet Service Network are an IDS sensor, an intranet/application server, and the internal firewall. This service network is also where the VPN tunnels will terminate. Inside the internal firewall is the LAN infrastructure, including an IDS sensor/server, a database server, and a mail server.

© SANS Institute 2000 - 2005, Author retains full rights.

© SANS Institute 2000 - 2005 Author retains full rights.



As in the first two designs, the border router is the first line of defense and filters known bad traffic. The border router does not need special configurations simply because it is at the corporate office.

The VPN connections are still terminated within a service network in order to provide the security described in the previous two designs.

Once again, the IDS sensors are deployed to give maximum information for security personnel. Again, like the other designs, by placing sensors on each segment of the security architecture a security analyst will be able match particular alerts to determine to what level the security is being penetrated.

Outbound traffic security is again handled as it is in the other designs. NAT services for outbound traffic are configured in order to provide security for internal equipment.

Group Access Requirements

The following are the access requirements tables. They start with the group that would have the least amount of access, the Public, and proceed to the greatest amount of access. I have also added one group that was not in the requirements document, which is Intersite Communications. It's important to identify any intersite network traffic that does not belong to a particular user or group of users, but that does affect network security.

© SANS Institute 2000 - 2005 Author retains full rights.

Public

The Public group consists of any user/organization that is not company personnel or affiliated with the company. Simply, they are users who do not belong to any of the other security groups.

Source	Destination	Port(s)/Protocol	Description
General Public	Company Website	80/TCP (HTTP)	Public access to the company's website to view company information.
General Public	Company DNS Server	53/UDP & TCP (DNS)	Access to make queries against the company's DNS server.
General Public	Company Mail Server	25/TCP (SMTP)	Access to send SMTP email to the company's email server.

Customers

The Customers group consists of those users/organizations that are purchasing fortune cookie sayings from the GIAC Enterprises.

Source	Destination	Port(s)/Protocol	Description
Potential Customers	Sales Website	80/TCP (HTTP)	Access to the sales section of the company website with information on what products are offered and how to become a customer.
Existing Customers	Secure Sales Website (MyAccount)	443/TCP (HTTPS)	Access to the secured sales section of the company website so that customers may place orders and make payments using a secure connection.
Customers	Company DNS Server	53/UDP & TCP (DNS)	Access to make queries against the company's DNS server.
Customers	Company Mail Server	25/TCP (SMTP)	Access to send SMTP email to the company's email server.

Suppliers

The Suppliers group consists of those users/organizations that supply GIAC Enterprises with fortune cookie sayings.

Source	Destination	Port(s)/Protocol	Description
Potential Supplier	Company Website	80/TCP (HTTP)	Access to the supplier section of the company website with information on how to become a supplier.

© SANS Institute 2000 - 2005 Author retains full rights.

Existing Supplier	Secure Supplier Website (MyAccount)	443/TCP (HTTPS)	Access to the secured supplier section of the company website so that suppliers can access their account information.
Existing Supplier	FTP Upload Site	20-21/TCP (FTP)	FTP access so that suppliers can upload their fortune cookie sayings.
Suppliers	Company DNS Server	53/UDP & TCP (DNS)	Access to make queries against the company's DNS server.
Suppliers	Company Mail Server	25/TCP (SMTP)	Access to send SMTP email to the company's email server.

Partners

The Partners group consists of those users/organizations that act as resellers for GIAC Enterprises fortune cookie sayings.

Source	Destination	Port(s)/Protocol	Description
Potential Partner	Company Website	80/TCP (HTTP)	Access to the partner section of the company website with information on how to become a partner.
Existing Partner	Secure Partner Website (MyAccount)	443/TCP (HTTPS)	Access to the secured partner section of the company website so that suppliers can access their account information.
Existing Partner	FTP Download Site	20-21/TCP (FTP)	FTP access so that partners can download their fortune cookie sayings.
Partners	Company DNS Server	53/UDP & TCP (DNS)	Access to make queries against the company's DNS server.
Partners	Company Mail Server	25/TCP (SMTP)	Access to send SMTP email to the company's email server.

Remote Users

The Remote Users group consists of company employees who are not physically connected to the company's internal LAN infrastructure. This table lists the most likely ways that an employee could access company systems and is not site specific. The company's Remote Access Policy should determine each employee's remote access abilities.

Source	Destination	Port(s)/Protocol	Description
Remote User	Company Website	80/TCP (HTTP)	Access to the company's website.

© SANS Institute 2000 - 2005 Author retains full rights.

Remote User	Secure Company Website	443/TCP (HTTPS)	Access to secure portions of the company's website.
Remote User	Web Accessible Email*	80/TCP (HTTP) 443/TCP (HTTPS)	Access to a website implementation of an email client.
Remote User	FTP Site	20-21/TCP (FTP)	Access to the company's FTP site.
Remote User	Company DNS Server**	53/UDP & TCP (DNS)	Access to make queries against the company's DNS server.
Remote IT User	Internal IT Equipment	22/UDP & TCP (SSH)	Secure Shell access to internal IT systems.
Remote User	VPN Concentrator	500/UDP (IKE)	Access to establish a VPN connection.
Remote User	VPN Concentrator	IP 50 (ESP)	VPN tunneling protocol.
Remote User	Intranet Server (via VPN)	80/TCP (HTTP)	Access to the company's intranet web server.
Remote User	Database Server (via VPN)	1433/UDP & TCP (SQL Server) or 1521/TCP (Oracle)	Access to the company's database server.
Remote User	Local LAN (via VPN)	88/UDP & TCP (Kerberos)	Local LAN login authentication.
Remote User	Local LAN (via VPN)	445/UDP & TCP (SMB)	Local LAN login authentication.

*This is a generic web email implementation. Specific software may listen on a different port.
**This is a generic implementation of DNS services. Best practice is to split the DNS between publicly accessible records and records only for use on the internal LAN.

Internal Users

The Internal Users group consists of company employees that are connected directly to the company's internal LAN infrastructure. Unlike any of the previous tables that list access from the outside, this table lists the access that internal users would need for outside systems.

Source	Destination	Port(s)/Protocol	Description
Internal User	Internet Websites (via NAT)	80/TCP (HTTP)	Access to websites on the Internet.
Internal User	Secure Internet Websites (via NAT)	443/TCP (HTTPS)	Access to secure websites on the Internet.

Intersite Communications

The Intersite Communications group lists the connections that would be required between the various company sites but that are not specific to any single person.

© SANS Institute 2000 - 2005 Author retains full rights.

Source	Destination	Port(s)/Protocol	Description
Office A	Office B	500/UDP (IKE)	Access to set up site-to-site VPN connection tunnels between the company's offices.
Office A	Office B	IP 50 (ESP)	Site-to-site VPN tunnel encapsulation.
Company DNS	Upstream DNS Servers	53/UDP & TCP (DNS)	Access to allow the company's DNS servers to query upstream DNS servers for DNS information.
Office Network A	Office Network B	Network specific ports and protocols (via VPN)	Ports and protocols that transfer network specific information between servers or between servers and workstations.

Equipment List

There are many different ways to put together the equipment that would work with this security architecture. The following list does not necessarily represent the ideal solution. An ideal solution would be a match between functionality and budget, leaning more toward functionality rather than budget. It is based on a ground-up approach rather than assuming that there are already pieces of equipment in place that will have to be used in the design. The list is also based on an assumption that the company has a policy that puts forth the standards for any IT equipment in the company.

Border Routers

The router selected for each office is a Cisco 2600 series router.

Cisco was chosen because it is an industry leader. While it may be a little more costly than some competitors, there is an established knowledgebase for support and maintenance. Also, given the small size of the company, it is prudent to chose equipment that is well known and will be easy to maintain by any of the company's IT staff.

The border routers are the first line of defense for each office. They can filter traffic before it ever reaches the company's networks. By blocking known bad and unwanted traffic, some of the load is lifted from the firewalls, which can then be configured for more advanced filtering. The advantage of filtering at the router level is that traffic is filtered before it ever reaches the company network. One disadvantage is that the more complex the filters become, the harder it is to track and manage. So, to mitigate this, the router filters should be kept simple and the more complex filtering should be done by the firewalls.

VPN Concentrators

The VPN concentrator for each office is a Cisco 3005 VPN concentrator.

© SANS Institute 2000 - 2005 Author retains full rights.

Again, Cisco was chosen because it is an industry leader and because of the support base. Consideration was also given to the fact that Cisco is available worldwide, making it possible to standardize all of the offices. Consideration was also given to the United States' restrictions on the exportation of encryption technology to other countries, and fortunately Cisco VPN technology is available in all the countries where the GIAC Enterprises has offices. Cisco was also chosen in order to maintain/show the company's standard equipment policy.

The VPN concentrator is used to provide a secure channel of communication from site to site and for remote access users. Some administrators tie the VPN concentrator to a firewall. This weakens the firewall since many firewalls pass encrypted VPN traffic through the firewall without inspecting the payload of the packet. This then becomes a security hole that can be exploited. While combining the concentrator with some other piece of equipment may work in some instances, it is much better to have the concentrator be stand-alone. By engineering it as stand-alone, the concentrator can be placed at the appropriate point in the architecture that creates the best possible security scenario.

Firewalls

The firewall selected for each office is the Sidewinder G2 Security Appliance firewall; model 410 for the non-sales field offices, model 510 for the field offices with sales, and model 1100 for the corporate office.

While it would seem that Cisco would be the choice in this situation given the company's standard equipment policy, Sidewinder was chosen to show that not all equipment in a standard equipment policy has to be from the same manufacturer. Selecting different vendors is also important because there may be times where if a flaw is found in one piece of equipment, other equipment from the same manufacturer may also have the flaw. Multiple manufacturers are chosen to reduce the risk that a single vendor issue could make a major portion of the network insecure. Sidewinder was also chosen because of its ability to properly handle multiple network connections off of a single piece of equipment, which reduces the number of machines being administered and ultimately the cost of security.

Firewalls perform complex packet filtering, and to a certain degree packet payload filtering. Firewalls have long been the one security device that cannot be left out of a security architecture design. While many consider firewalls to be perimeter defense, they can be deployed anywhere in the security architecture for increased security between two networks, including service networks. They can also be deployed as individual units or as multiple interfaces on the same unit, as is done is this paper.

IDS Sensors

The IDS sensors selected to be placed in each office are a standard PC – any manufacturer – running Fedora Core 3 Linux and Snort version 2.3.2. The sensor information is rolled up into a central database at each office running on the internal IDS sensor. Each office's information is then rolled up into a central database at the corporate office running on a server that provides analysis tools.

Snort was chosen to be the IDS sensor due to its widespread acceptance as an

© SANS Institute 2000 - 2005 Author retains full rights.

outstanding security tool and because of its robust nature. While Snort can be run under Microsoft Windows, it is more configurable under Linux. Fedora Linux was chosen because of its widespread use and support base. As a matter of opinion, Linux is much easier to scale back the operating system and support services to a level that makes it ideal for an IDS sensor. This configuration can also be set up to consolidate sensor alerts into a central database that can then be analyzed, rather than having to analyze the alerts from each individual sensor. Snort and Linux were also chosen because of the obvious cost issue – Free.

IDS sensors are needed to assist in the analysis of network traffic that appears to be normal traffic but that may actually be malicious in nature. The biggest issue facing IDS sensors is the issue of false alerts – legitimate traffic that triggers an alert. In a security architecture that has only one IDS, determining false alerts or what course of action to take for alerts is difficult. To mitigate this problem, multiple IDS sensors are placed in the architecture in such a way that security alerts can be correlated and an appropriate response can be initiated.

IP Addressing Scheme

The internal network IP addresses for each of the offices will be a portion of the 10.0.0.0/8 private address space. Class C address spaces are chosen to make working with the IP addresses easier.

Denver	-	10.10.10.0	255.255.255.0
Seattle	-	10.10.20.0	255.255.255.0
Miami	-	10.10.30.0	255.255.255.0
London	-	10.10.40.0	255.255.255.0
Sydney	-	10.10.50.0	255.255.255.0

The Intranet Service Network IP addresses for each of the offices will be a portion of the 172.16.0.0/12 private address space. Class C address spaces are chosen to make working with the IP addresses easier.

Denver	-	172.16.32.0	255.255.255.0
London	-	172.16.42.0	255.255.255.0
Sydney	-	172.16.52.0	255.255.255.0

The Internet Service Network IP addresses for each of the offices will be a portion of the 192.168.0.0/16 private address space. Class C address spaces are chosen to make working with the IP addresses easier.

Denver	-	192.168.55.0	255.255.255.0
Seattle	-	192.168.65.0	255.255.255.0
Miami	-	192.168.75.0	255.255.255.0
London	-	192.168.85.0	255.255.255.0
Sydney	-	192.168.95.0	255.255.255.0

NAT will be used to provide routable IP addresses for all internal IP addresses. Static NAT will be used for IP addresses that need to be accessed from outside the company, such as a web server or VPN concentrator. Dynamic NAT will be used to minimize the number of routable IP addresses needed by the company. The following

© SANS Institute 2000 - 2005 Author retains full rights.

list provides an example of a possible routable IP address scheme.

Denver IP Addresses (64)	155.70.58.64	255.255.255.192
Static NAT	155.70.58.65 – 155.70.58.79	
Dynamic NAT	155.70.58.80 – 155.70.58.126	
Seattle or Miami IP Addresses (16)	155.70.172.144	255.255.255.240
Static NAT	155.70.172.145 – 155.70.172.150	
Dynamic NAT	155.70.172.151 – 155.70.172.158	
London or Sydney IP Addresses (32)	212.28.221.0	255.255.255.224
Static NAT	212.28.221.1 – 212.28.221.10	
Dynamic NAT	212.28.221.11 – 212.28.221.30	

Conclusion

This is one possible solution for GIAC Enterprises. The security architecture presented here provides a layered approach to security, some of it logical and some of it physical. Both the logical and physical network designs have been presented. Access tables showing who needs access to what services has been included. And, the equipment needed to build the presented security architecture has been listed. There is no single piece of equipment that if it were compromised would grant access to the entire network. At the same time, the security is not overly restrictive in denying access to resources for customers, suppliers, partners, and employees. The equipment listed was chosen to provide the best possible solution based upon budgetary and political considerations and company standards.

This solution does not mark the end of the security issues that need to be addressed. The major network security components were presented; however, upon implementation the security architecture will need to be tweaked to adjust for the minor issues that surround day-to-day IT operations. This architecture will also require the support of management in the form of policies and procedures that are enforceable. Other issues that will need to be addressed that are not addressed here include information security, an account policy, a password policy, single sign-on issues, anti-virus, and physical security.

© SANS Institute 2000, Author retains full rights.

© SANS Institute 2000 - 2005 Author retains full rights.

GIAC Enterprises Router and Firewall Security

In order to be truly effective, security must be planned and properly implemented, not added as an afterthought. The more complex the security, the more it should be planned out in advanced and tested. This is especially the case with router access control lists (ACLs) and firewall rulesets due to the fact that misconfiguration could cause a widespread outage or leave the network entirely vulnerable. Care should be given to address both incoming and outgoing traffic.

General Security Stance

The general security stance of GIAC Enterprises is that they want to be as open and accessible as possible, in order to attract business. Conducting that business should be as secure as possible, though. The product they sell is not confidential or regulated so they want security on product distribution that is both secure enough to protect their business yet is easily manageable and does not require a great deal of effort on the part of their associates. To accomplish these goals they have added web servers that provide content to the public and potential associates. They have added secure services to their web server to provide account management. And, they have added FTP servers to allow for swift and secure distribution of their product.

For employees, GIAC Enterprises wants to provide a robust working environment that allows employees to do their job effectively without granting them more access than they need. GIAC Enterprises wants to secure all communication between their offices. They want to ensure that employees can access the company LAN from wherever they happen to be, especially the field sales staff, while still maintaining IT security. And, they want to allow employees on the company network to browse the Internet, but not to utilize other time consuming activities such as Internet Chat. To accomplish these goals they added site-to-site VPNs between the offices. They have added VPN remote access for employees that are out of the office. And, they have added egress filters to limit what employees on the company network can access outside the company.

Filtering Router Policy

The router chosen is a Cisco 2600 series router. On this router, ACL entries are processed in a top-down order, which makes it important to know which entries need to come before other entries. It should also be noted that entries added to an ACL are appended to the end of the list. So, in order to insert an entry, the entire list needs to be deleted and recreated.

Also on this router, once a packet has matched an entry the packet is not compared to any more entries. The action associated with the matching entry is carried out and the packet is either routed or dropped. For example, on an ingress filter, if a permit packet entry is placed before a deny packet entry and a packet always matches the permit entry, then nothing will get denied because processing of the packet will always stop with the permit entry. This is important to know so that permit and deny entries are placed in their proper relationship based on whether the filter is an ingress or egress filter.

© SANS Institute 2000 - 2005 Author retains full rights.

Another property of the router that should be pointed out is that while many ACLs may be created, only one ACL may be associated with a particular direction, incoming or outgoing, on each port. This means that each port on the router may have one ingress filter and one egress filter associated with it. This makes it easy to update an ACL and then switch it out for the ACL on the port rather than worry about messing up an active connection.

Access Tables

The following table represents entries that are to be placed in an ingress filter that would be applied to the port of the router that connects it to the Internet. The ACL works better as an ingress filter on this port rather than an egress filter on the port that connects to the company's DMZ. By setting it as an ingress filter, packets can be dropped without taking up buffer space in the router, after being accepted by the incoming interface. The sections are placed in top-down order to show an appropriate order to filtering rules.

Known "Bad" IP Addresses (Lines 1-2)

The filtering router should be configured to deny packets inbound from the Internet that are from a source that the company determines is a security risk. This category includes IP addresses that are listed as associated with known hacker organizations, known attackers, and blocks of IP addresses associated with a country or region of the world that is actively engaged in malicious traffic. This list may be quite dynamic and will require monitoring to ensure that valid traffic is not being denied.

Unused IP Addresses (Lines 3-7)

The filtering router should deny packets inbound from the Internet that have a source IP that falls within the range of IP addresses that have not yet been assigned. These IP addresses may be obtained by visiting the IANA website at <www.iana.org/assignments/ipv4-address-space>. While an administrator who didn't know better may have picked and used these IP addresses, there is no way to verify that this is the case. It is much better to simply deny these packets access.

Loopback Address (Lines 8-9)

The filtering router should deny packets inbound from the Internet that have the loopback address as the source or destination. The loopback address is used exclusively internal to a computer system and should not be routed across any networks. Many pieces of computer equipment may accept and process loopback packets without any checks on the origin of the packet.

Unknown Destination Address (Line 10)

The filtering router should deny packets inbound from the Internet that have a destination IP that is not a company IP address. While this may seem obvious that a router would not route packets for a destination that it does not know, routers do attempt to honor packets that specifically list the hops that the packet must take. The router may determine that the packet is not destined for a company network, but if the next hop in the source route is a company IP it will try to route it to that IP. By blocking this traffic these types of source-route exploits can be eliminated.

Private IP Addresses (Lines 11-13)

The filtering router should deny packets inbound from the Internet that have a

© SANS Institute 2000 - 2005 Author retains full rights.

source IP that falls within the range of IP addresses that have been reserved for private use – 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/24. According to RFC1918 these addresses are not supposed to be routed across the Internet since they are for private network addressing only. Any packet that contains one of these IP addresses as the source address violates that RFC and should be blocked.

Broadcast and Network Addresses (Lines 14-19)

The filtering router should deny packets inbound from the Internet that have IP addresses matching particular network ID or broadcast addresses. In particular, packets should be denied that have an IP of 0.0.0.0 or 255.255.255.255 as the source or destination. Packets that are addressed to the company's network ID or broadcast address should also be denied. This traffic is usually tied to an attempt to map the company's network as a precursor to a more directed attack.

Permit All Other Traffic (Line 20)

The final entry in this ACL should allow all other traffic that has not yet been denied. This entry is needed because the default action of the router, which is to permit all traffic, changes to an action of deny all traffic as soon as an ACL is created. This means that without this entry, even though specific deny statements were entered, no traffic would be routed by the router.

Line #	Permit/Deny	Source	Destination
1	Deny	Incidents.org Top 10 IPs	Any IP
2	Deny	Country IPs	Any IP
3	Deny	0.x.x.x	Any IP
4	Deny	1.x.x.x	Any IP
5	Deny	2.x.x.x	Any IP
6	Deny	5.x.x.x	Any IP
7	Deny	7.x.x.x	Any IP
8	Deny	Any IP	127.0.0.1
9	Deny	127.0.0.1	Any IP
10	Deny	Any IP	Any Non-Company IP
11	Deny	10.0.0.0/8	Any IP
12	Deny	172.16.0.0/12	Any IP
13	Deny	192.168.0.0/24	Any IP
14	Deny	0.0.0.0	Any IP
15	Deny	Any IP	0.0.0.0
16	Deny	255.255.255.255	Any IP
17	Deny	Any IP	255.255.255.255
18	Deny	Any IP	Company IP.0
19	Deny	Any IP	Company IP.255
20	Permit	Any IP	Any IP

The following table represents entries that are to be placed in an ingress filter that would be applied to the port of the router that connects it to the company's DMZ. The ACL works better as an ingress filter on this port rather than an egress filter on the port that connects to the Internet. By setting it as an ingress filter, packets can be dropped without taking up buffer space in the router after being accepted by the incoming interface. The sections are placed in top-down order to show an appropriate

© SANS Institute 2000 - 2005 Author retains full rights.

order to filtering rules.

Private IP Addresses (Lines 1-3)

As with the above ACL, the same private IP address filters should be applied to any traffic coming from the company network destined for the Internet. While many would argue that these could be filtered out by the firewall, which would make this unnecessary on the router, this is simply not the case. The firewall can be misconfigured and thus route private addresses to the router. In order to be a good Internet neighbor, this filter needs to be in place just in case this situation occurs.

Known "Bad" IP Addresses (Lines 4-5)

Similar to the above ACL, the filtering router should be configured to deny packets outbound to a destination that the company determines is a security risk. This may include destinations that the company does not want to allow on its network, such as instant messaging connections. It should include any sites that attempt to circumvent security policies and practices, such as gotomypc.com. This list will also be quite dynamic and will require monitoring to ensure that traffic to a particular destination should still be denied.

Company Source IP Addresses (Line 6)

The filtering router should permit packets outbound to the Internet that have a source IP address that is a known company IP. Ordinarily this rule would be unnecessary since only company IP addresses should be used; however, this follows the best practice of being a good Internet neighbor. This eliminates the possibility that someone could use the company network to spoof packets or launch source-route attacks. Again, the firewall could be used to perform this packet filtering, but it could still be misconfigured and allow this traffic through.

Deny All Other Traffic (Line 7)

Traffic leaving the company's network should be consistent, requiring fewer filtering rules. The last entry should be to deny all other traffic. Many would question why this entry is needed since once an ACL is created the firewall denies all traffic by default. The reason that this entry is needed is because of the way that the router handles new entries. New entries are appended to the end of the list. If this entry was not here and someone inadvertently added an entry to permit any traffic, it would be added to the end of the list and suddenly any traffic would be allowed through. With this entry added, an appended entry would not inadvertently change what traffic gets filtered.

Line #	Permit/Deny	Source	Destination
1	Deny	10.0.0.0/8	Any IP
2	Deny	172.16.0.0/12	Any IP
3	Deny	192.168.0.0/24	Any IP
4	Deny	Company IP	205.188.7.0 (AOL Instant Messenger)
5	Deny	Company IP	Gotomypc.com
6	Permit	Company IP	Any IP
7	Deny	Any IP	Any IP

© SANS Institute 2000 - 2005 Author retains full rights.

The filtering router could be configured to filter traffic based on source and destination ports, but that would increase the complexity and management of the router ACLs to a level that would not be consistent with a company of this size. Such filtering is better handled by the primary firewall.

Primary Firewall Policy

The Primary Firewall Policy cannot be summed up into a single rule list since there are differing levels of security that the firewall protects. This firewall policy will describe the rulesets of the firewall(s) that protect the Denver corporate office of GIAC Enterprises. To review, the design includes an external DMZ, two service networks, and an internal LAN; however, there is only one physical firewall with four interfaces. Connected to the Internet Service Network interface are a DNS server, a web server, and an FTP server. Connected to the Intranet Service Network interface is an intranet / application server and the internal interface of the VPN concentrator. Inside the firewall is the LAN infrastructure, including a database server and a mail server.

The firewall that was chosen for GIAC Enterprises was a Sidewinder G2 appliance firewall. This firewall is a stateful firewall, which means that it can track the state of a communication session. This is important because complex rules do not have to be written to allow follow-on packets of a session through the firewall filters. It also means that sessions such as FTP, which jump from one port to another port, do not require complex redundant rules to allow them to occur. The purpose of the router is to route all traffic and filter out unwanted packets. The purpose of the firewall is to allow valid packets and filter out all other traffic.

General Firewall Configuration

The firewall should be secured as well as, if not better than, the other security equipment on the network. There are a myriad of tweaks that will improve network security. For example, the firewall should be set to not respond to any ICMP Echo Requests sent to it. It should be set to quench known packet attacks, such as a SYN floods. The firewall should also not respond with RST / RST/ACK packets or with ICMP Destination Unreachable messages. All of these settings protect the firewall and network but do not and should not hinder normal network traffic.

This does not mean that the firewall needs to be totally locked down and act like a “black box.” Valid traffic destined for the firewall needs to be accepted and handled. For example, the firewall should be configured to accept and act on ICMP Source Quench and ICMP Destination Unreachable packets. The company policy allows for remote administration of the firewall so it will have to accept SSH sessions. The firewall should also accept valid management sessions from the internal LAN. These settings allow the firewall to be usable and maintainable without undue security risks.

The following tables present the firewall rulesets based upon the interface to which the rules apply. Like the router, the firewall rules are processed in a top-down order, and once a rule has been matched, processing stops and the packet is handled. This makes the order of the rules important. While it is important to get the permit and deny rules in the proper order, firewall rulesets can grow to be very large so the rules also need to be optimized to get the best performance out of the firewall.

© SANS Institute 2000 - 2005 Author retains full rights.

Ruleset on the Public (Outside) Interface

This is traffic from the Internet that is destined for the company LAN.

Permit DNS Queries (Line 1)

The firewall should permit queries from the Internet against the DNS server. This is necessary to map the web server's name to its IP address. This is also necessary to map the mail server's name to its IP address.

Permit SMTP Traffic (Line 2)

The firewall should permit connections to the Email server.

Permit HTTP Traffic (Line 3)

The firewall should permit requests to the web server.

Permit HTTPS Traffic (Line 4)

The firewall should permit secure connections to the web server.

Permit FTP Traffic (Lines 5-6)

The firewall should permit connections to the FTP server.

Permit SSH Traffic (Line 7)

The firewall should allow SSH connections to equipment for remote control by IT personnel.

Deny All Other Traffic (Line 8)

The firewall should deny all other traffic. While this rule may seem redundant since the firewall blocks all traffic by default, it is added to attempt to catch traffic just in case the firewall is misconfigured.

Line #	Source IP	Source Port	Destination IP	Destination Port	Protocol
1	Any IP	Any Port	Company DNS Server	53 (DNS)	UDP
2	Any IP	Any Port	Company SMTP Server	25 (SMTP)	TCP
3	Any IP	Any Port	Company Web Server	80 (HTTP)	TCP
4	Any IP	Any Port	Company Web Server	443 (HTTPS)	TCP
5	Any IP	Any Port	Company FTP Server	20 (FTP)	TCP
6*	Any IP	Any Port	Company FTP Server	21 (FTP-DATA)	TCP
7	Any IP	Any Port	Company Equipment IP	22 (SSH)	TCP
8	Any IP	Any Port	Any IP	Any Port	Any Protocol

*This rule is only necessary if for some reason FTP sessions are failing to connect.

Ruleset on the Internet Service Network Interface

This is traffic from the Internet Service Network destined for either the Internet or the company LAN. Most administrators don't think about filtering traffic from an internal source traveling to a destination outside of the network; however, if the traffic is

© SANS Institute 2000 - 2005 Author retains full rights.

not filtered then any compromised host can be used to attack other networks on any port that is allowed out. By including this ruleset the company is being a good Internet neighbor and protecting themselves from becoming liable for attacks on other networks.

Permit Queries from the DNS Server (Line 1)

When the DNS server does not have an entry in its cache for the requested query, it must attempt a recursive query or pass the query on to an upstream DNS server in order to fulfill the request. This rule allows that traffic out of the Internet Service Network.

Permit Web Email Traffic (Line 2)

The web access email application on the web server will need to access the Email Server. The port that this traffic uses is typically configurable based on the email application.

Deny All Other Traffic (Line 3)

There are no other applications in the Internet Service Network that will need to initiate a connection, which means that all other connections should be denied. This does not mean that all other traffic will be blocked. Since the firewall is stateful, other traffic from existing sessions will travel to and from this network segment. This is the rule that makes it so that a compromised host cannot be easily used to attack other networks.

Line #	Source IP	Source Port	Destination IP	Destination Port	Protocol
1	Company DNS Server	Any Port	Any IP	53 (DNS)	UDP
2	Company Web Server	Configured Port	Company Email Server	Configured Port	TCP
3	Any IP	Any Port	Any IP	Any Port	Any Protocol

Ruleset on the Intranet Service Network Interface

This is traffic from the Intranet Service Network destined for the company LAN. Like the previous ruleset, this ruleset makes the company a good Internet neighbor and protects them from being liable for attacks on other networks.

Permit VPN Traffic (Lines 1-7)

All traffic coming from the VPN connection should be given appropriate access to the internal LAN. Much of the VPN connection is configurable and the table below gives an example of rules that could be included, but in reality what traffic is allowed or not allowed should be discussed, agreed upon, and tested.

Permit Application Traffic (Line 8)

Applications on the Application Server may need to access the Database Server. The port that this traffic uses is typically configurable based on the application and the database.

© SANS Institute 2000 - 2005 Author retains full rights.

Deny All Other Traffic (Line 9)

As with the previous ruleset, there are no other applications in the Intranet Service Network that will need to initiate a connection, which means that all other connections should be denied.

Line #	Source IP	Source Port	Destination IP	Destination Port	Protocol
1	Branch Office Client IP	Any Port	Company Database Server	1521 (Oracle)	TCP
2	Branch Office Server IP	Server-to-Server Communication Port	Company File Server	Server-to-Server Communication Port	TCP
3	Remote Access VPN Client IP	Any Port	Company Email Server	Email client access port	TCP
4	Remote Access VPN Client IP	Any Port	Company Database Server	1521 (Oracle)	TCP
5	Remote Access VPN Client IP	Any Port	Company Authentication Server	88 (Kerberos)	TCP & UDP
6	Remote Access VPN Client IP	Any Port	Company Authentication Server	445 (SMB)	TCP & UDP
7	Branch Office Client IP	Any Port	Company Database Server	1521 (Oracle)	TCP
8	Company Application Server	Configured Port	Company Database Server	Configured Port	TCP
9	Any IP	Any Port	Any IP	Any Port	Any Protocol

Ruleset on the Internal LAN Interface

This is traffic from the Internal LAN destined for either the Internet or one of the service networks. Like the previous rulesets, this ruleset makes the company a good Internet neighbor and inhibits attacks from company systems against other networks.

Deny Traffic to Specific IPs (Lines 1-2)

The firewall should be configured to block traffic coming from the Internal LAN that is going to any site that the company determines is undesirable. Typically this is a site that attempts to circumvent security, such as gotomypc.com, but it may be any IP that the company determines should be blocked. These rules are placed at the beginning of the list so that traffic will not inadvertently match a permit rule and get out when it should be denied.

© SANS Institute 2000 - 2005 Author retains full rights.

Permit VPN Traffic (Lines 3-5)

All valid traffic destined for the VPN connection should be allowed out. Much of the VPN connection is configurable and the table below gives an example of rules that could be included, but in reality what traffic is allowed or not allowed should be discussed, agreed upon, and tested. Since a great deal of the traffic leaving the company LAN will be traffic destined for the Branch Offices, these rules are placed here to reduce the number of rules that will have to be compared before the traffic is allowed to proceed.

Permit SMTP Traffic (Line 6)

Allow the Email Server to make SMTP connections with other email servers. This rule defines those systems that are permitted to send email. This is useful in that a computer virus may spread through email, and some even install their own SMTP server on a workstation. This rules makes it so only allowed email servers are able to send email, and a computer virus or other user-installed program cannot spread malicious traffic or consume extra bandwidth.

Permit HTTP Traffic (Line 7)

Allow workstations and servers on the Internal LAN to connect to web servers.

Permit HTTPS Traffic (Line 8)

Allow workstations and servers on the Internal LAN to connect to web servers over secure connections.

Permit DNS Queries (Line 9)

Allow workstations and servers on the Internal LAN to query the company's DNS server.

Permit FTP Traffic (Line 10)

Allow workstations and servers on the Internal LAN to connect to FTP servers.

Permit NTP Queries (Line 11)

Allow workstations and servers on the Internal LAN to make NTP queries to Internet time servers. This can be further restricted to a single server if the server is set up to serve time to the rest of the company's LAN.

Deny All Other Traffic (Line 12)

Like the previous rulesets, in order to be a good Internet neighbor and to protect the company from having undesirable traffic leaving the company LAN, all other connections should be denied. This also makes it so that someone can't simply set up another network on the company LAN and have it routed to the Internet.

Line #	Source IP	Source Port	Destination IP	Destination Port	Protocol
1	Company IP	Any Port	Gotomypc.com	Any Port	Any Protocol
2	Company IP	Any Port	Undesirable IP	Any Port	Any Protocol
3	Company Server IP	Server-to-Server Communication Port	Branch Office Server IP	Server-to-Server Communication Port	TCP
4	Company Client IP	Any Port	Branch Office Equipment IP	22 (SSH)	TCP
5	Company Client IP	Any Port	Branch Office FTP Server	20 (FTP)	TCP

© SANS Institute 2000 - 2005 Author retains full rights.

6	Company Email Server	Any Port	Any IP	25 (SMTP)	TCP
7	Company IP	Any Port	Any IP	80 (HTTP)	TCP
8	Company IP	Any Port	Any IP	443 (HTTPS)	TCP
9	Company IP	Any Port	Company DNS Server	53 (DNS)	UDP
10	Company IP	Any Port	Any IP	20 (FTP)	TCP
11	Company IP	Any Port	Internet Time Server	123 (NTP)	UDP
12	Any IP	Any Port	Any IP	Any Port	Any Protocol

Conclusion

This is a good start on the filtering router and firewall policies; however, security is not simply a matter of creating a policy and then moving on. Security is an ongoing process. Once these entries and rules are put in place they need to be tested, and if any valid traffic is being blocked, they need to be changed. Other changes in the company may necessitate changes to security. And, as a general rule, the configurations of the router and firewall need to be regularly audited to make sure that there aren't any inadvertent or unauthorized changes that leave the company's network vulnerable to attack.

© SANS Institute 2000 - 2005 Author retains

© SANS Institute 2000 - 2005 Author retains full rights.

Bibliography

"An Introduction to Biometrics." The Biometric Consortium website. 2005. March 30, 2005 <http://www.biometrics.org/html/introduction.html>.

Alpert, Mark. "Security at Your Fingertips: Fingerprint sensors can guard your computer data." Scientific American. June 2004. March 30, 2005 <http://www.sciam.com/article.cfm?chanID=sa006&colID=11&articleID=0008BE6E-72C4-10A9-A47783414B7F0000>.

Bicz, Wies•aw. "Fingerprint structure imaging based on an ultrasound camera." OPTEL website publication. July 16, 2003. March 30, 2005 <http://www.optel.pl/article/english/article.htm>.

Department of Defense Biometrics Home Page. 2005. March 30, 2005 <http://www.biometrics.dod.mil/>.

Ilett, Dan. "Gates: Passwords passe." ZDNet News website. November 16, 2004. March 30, 2005 <http://news.zdnet.com/2100-1009_22-5454719.html>.

Kotadia, Munir. "Gates predicts death of the password." CNET News.com website. Febrauary 25, 2004. March 30, 2005 <http://news.com.com/Gates+predicts+death+of+the+password/2100-1029_3-5164733.html?tag=nl>.

Lemos, Robert. "This hacker's got the gummy touch." CNET News.com website. May 16, 2002. March 30, 2005 <http://news.com.com/2100-1001-915580.html?legacy=cnet&tag=pt.rss..feed.ne_9914626>.

Liu, Simon and Mark Silverman. "A Practical Guide to Biometric Security Technology." IT Professional website. 2000. March 30, 2005 <http://www.computer.org/itpro/homepage/jan_feb01/security3.htm>.

NIST Biometric Standards webpage. January 8, 2003. March 30, 2005 <http://www.itl.nist.gov/div893/biometrics/standards.html>.

© SANS Institute 2000 - 2005 Author retains full rights.

NISTIR 6529-A, "Common Biometric Exchange Formats Framework (CBEFF)." NIST Publications. April 5, 2004. March 30, 2005. <http://www.itl.nist.gov/div893/biometrics/documents/NISTIR6529A.pdf>.

O'Sullivan, Orla. "Biometrics comes to life." ABA Banking Journal Online. January 1997. March 30, 2005 <http://www.banking.com/aba/cover_0197.htm>.

Ruggles, Thomas. "Comparison of Biometric Techniques." July 10, 2002. March 30, 2005 <http://www.bio-tech-inc.com/bio.htm#Bio_Accuracy>.

The BioAPI Consortium homepage. 2005. March 30, 2005 <http://www.bioapi.org/index.html>.

---

[1] Attributed to Chris Brenton. SANS Local Mentor Program course notes. March 2005.

[2] Department of Defense Overview webpage. 2005. March 30, 2005 <http://www.biometrics.dod.mil/content/content.aspx?navid=2&PageID=168>.

[3] Liu, page 3.

[4] NIST Biometric Standards webpage. January 8, 2003. March 30, 2005 <http://www.itl.nist.gov/div893/biometrics/standards.html>.

[5] Ibid.

© SANS Institute 2000 - 2005, Author retains full rights.

© SANS Institute 2000 - 2005 Author retains full rights.