Global Information Assurance Certification Paper

Computer Security

What is it Anyway?

SANS Security Essentials (GSEC) Practical Assignment

Version 1.4b Option #1

Brian Piirala

March 16, 2003

Introduction

This document is intended to be an introductory discussion of computer security for information technology professionals. Quite often, IT administrators and network operations personnel have also been assigned or have assumed the role of security manager for their particular environment. Security initiatives typically take a back seat to revenue generating activities within the business and security is frequently the last line item to be allocated budget within the information technology department. With this document, my goal is to help IT professionals define what computer security is at its core and how to implement a sound security infrastructure. I'll explore the concept of Defense in Depth (DiD) as a layered approach to security and I'll discuss the various technologies within each of these security layers including firewalls, policies, patch management, hardening, etc…

When beginning a corporate security discussion, I'm frequently reminded of the IBM television commercials that take place in a corporate boardroom. The team is discussing some great new device that will solve all of their problems. The commercials always end with the revelation that "there is no magic x machine." The same is true here. There is no magic security machine.

Defense in Depth

Our goal in information security is to achieve three primary objectives. (CIA)

- Confidentiality - information should be kept private and accessible to only the appropriate personnel
- Integrity - information should not be altered by any unauthorized people
- Availability - information should always be accessible by authorized personnel

To achieve these objectives we use a Defense in Depth strategy. This is a security philosophy taken from the military which basically says that security must be applied in layers at multiple enforcement points throughout the enterprise network. The idea behind this philosophy being that there is no single device or tool or individual piece of software that can stand independently as a sound security infrastructure.

We see this concept in our every day lives as well. A bank for example has multiple layers of security that you must traverse before you can get inside the vault. In our homes we may keep our valuables in a locked safe that may then be hidden in a closet or inside of a wall. Then we lock our doors and windows and may even add an alarm system in addition to that. We may even live within gated communities that require identification verification before entry is allowed. Any one of these security measures is a good practice but no single one of them can be considered a comprehensive security solution.

A good Defense in Depth strategy will accomplish two things. First, it will make it extremely difficult for an intruder or hacker to gain access to our data. Even if an attacker is able to break or bypass one of our defense mechanisms, it's quite unlikely that he or she will be able to break all of our layers of defense. Secondly, just having these defenses in place makes for a huge deterrent to potential attackers. In most cases, the bad guys are simply out looking for easy targets. They will scan the internet looking for vulnerable systems to attack using known exploits. This is similar to a burglar who is looking for a house to rob, he's much more likely to pick a house with an open window than one that's all locked up and armed with a security system.

Policies

The most overlooked facet of security is the human factor. Many times we find ourselves rushing to implement some new security technology or install some cool new software or we start re-configuring security settings in operating systems and applications. What we tend to forget is that all of this technology is run by people. No amount of technology can assure absolute security as long as the technology is run by humans. Before we rush into any technological security solutions, we need to first ask ourselves what exactly it is that we're trying to accomplish. We need to put our goals and objectives on paper. We then need to publish policies and procedures within our organization and train all personnel on these policies.

A good resource for human security awareness is http://www.humanfirewall.org The Human Firewall Council "is a consortium of professionals who have come together to help educate the public on the human issues involved in information security, as opposed to a strictly technical approach."[i] They conduct surveys on security awareness and they also provide a benchmarking system known as a Security Management Index which "is a free survey that enables security professionals to benchmark their security management efforts based on the global ISO 17799 standard. You get a score and report that compares your security management practices with others in your industry and peer group."[ii] This is a great starting point to gauge the security awareness within your company.

Once you have an accurate understanding of the security awareness within your company, then you'll know where to begin developing good policies to form the foundation of your security infrastructure. A great starting point for written policies is to use some of the templates provided by SANS at http://www.sans.org/resources/policies/#template . You'll find everything from acceptable use policies to password policies there. Of course there are even companies out there charging big money for their written policies too. You'll have to decide what's best for your particular organization.

Firewall

A firewall is defined as "a system or group of systems that enforces an access-control policy between two networks. It may also be defined as a mechanism used to protect a trusted network from an untrusted network. Firewalls have become a security must have now that so many organizations are connecting their internal networks to external networks such as the Internet."[iii]

A firewall is typically located at the outer boundary of a given network which is known as the perimeter. It is primarily used at the connection point between the Internet and an internal private local area network (LAN). Its job is to enforce access control policies which either allow or deny specific communications between the private network and the Internet.

Before you implement a firewall, you need to establish a policy regarding what types of communication and access that you want to allow and what access you're planning to deny. This is what makes firewalls such valuable tools because they allow you to take your written policy and use technology to enforce that policy. The general rule of thumb when initially implementing and configuring a firewall is to deny all communications originating from machines outside of your firewall attempting to connect to machines inside of your firewall. This philosophy of "deny all" immediately provides a heavy layer of protection to every machine on your network from potential Internet intruders. From this point, you can begin to open up or "allow" certain types of communication to pass through. You can then setup individual rules that will allow or deny specific types of communication only to specific hosts. It is also important to note that these rules can be configured for outbound traffic as well allowing you to restrict your entire user base from certain types of communication.

Firewalls can make a huge impact on the security of your network although they can't protect against everything. They cannot prevent malicious activity from occurring within the network. They cannot prevent sensitive data from exiting the company via floppy disk or other removable media. They cannot prevent data from being transmitted via modem. The bottom line is that firewalls cannot enforce any policy or rules of any kind for communication that does not travel through it. Even some malicious communications that do go through firewalls cannot be protected against. For example, an email message with a virus attached would not be blocked by a firewall.

A firewall is commonly considered the biggest bang for the buck in terms of security. It provides a central point of control for communication in and out of your network. It keeps detailed logs of all activity at the firewall which gives you the ability to compile very specific utilization reports as well as security reports showing potential hacker scanning attempts and similar activity. While a firewall is definitely a vital core component of a good security infrastructure, it's certainly not the magic security machine that solves our entire security needs.

Virtual Private Networks (VPN)

"A virtual private network (VPN) is a private data network that makes use of the public telecommunication infrastructure, maintaining privacy through the use of a tunneling protocol and security procedures."[iv] The most common use of virtual private networks is to use the Internet to allow individual home office employees or other remote company offices to securely access resources within their organization's network. This technique of using the already existing public infrastructure saves corporations millions of dollars each year over the alternative of installing and paying for a dedicated private connection in every remote location and the home of every home based employee. Additionally, this method is portable, since the public Internet is essentially available from just about any place in the world, these remote users can establish a VPN connection from any place that they may be at any given time.

A VPN works by encrypting and encapsulating your data before sending it over the Internet using "tunneling" technology. This keeps your data private even though it's being transmitted across the public Internet. The contents of your transmission are only readable by the recipient at the other end of your VPN.

Virtual private networks are amazingly simple to implement and are an excellent way to provide secure access to your corporate resources from any remote location. You don't even have to bother with any cost-benefit analysis. AVPN is a highly cost effective solution for an organization of 100,000 people or just one guy.

Patch Management

One of my all time favorite layers of defense and methods of hack prevention is simply to keep my machines patched with the current service packs and patches. It's no secret that a large majority of successful hacks are into unpatched systems using a well known exploit. This should not be happening.

There was a time when patch management was fairly easy to do because there were fairly infrequent patches or updates, and for the majority of our systems all we had to do was install service packs without any other updates for many months between. These days however, patch management is a daily job requirement especially for any security personnel. Microsoft alone has released "over 230 security bulletins since the beginning of 2000."[v]

The world today has an ever increasing number of machines and there are certainly a growing number of vulnerabilities being uncovered every day. Increasingly, it is becoming critically important not just to secure our data but also to protect our servers and workstations from being used as pawns to launch further attacks on others around the world. Each new legal case involving an electronic information security breach brings us closer to the day when we may be held partially responsible for someone else's attack on a third party if we are found to have been negligent in our own security practices.

Although patch management is now a much bigger and greatly more complex task than it once was, we now have a variety of tools to help us in this endeavor. Network World Fusion did a review of a few patch management tools recently. Their results are posted at http://www.nwfusion.com/reviews/2003/0303patchrev.html

Patch management systems are available for all environments independent of size, especially for Microsoft systems. In fact, even if you have just a single computer at home, it's a good idea to frequently visit the Microsoft Windows Update website at http://windowsupdate.microsoft.com Windows 2000 service pack 3 and later include the autoupdate client which enables these systems to automatically connect to the Windows update site on a scheduled basis, download and even install critical patches without user or administrator intervention. This is a very good idea for home users or many small office environments.

For slightly larger environments or for situations where systems and security administrators want to maintain a tighter degree of control, Microsoft also now offers Software Update Services (SUS).
http://www.microsoft.com/windows2000/windowsupdate/sus This is essentially the same as the traditional Windows update solution. However, SUS effectively allows you to install your own Windows update server in your own environment and synchronize its contents with the servers at Microsoft. You then have the ability to define which patches you'd like to deploy to your machines and which ones are not required. Then with Group Policy you can point the machines in your enterprise to your newly defined SUS server. For even larger organizations, Microsoft also offers a Systems Management Server (SMS) add on for patch management.

All of the above patch management solutions are extremely useful and certainly well worth the cost ($0.00). However, for most large enterprises, this just isn't quite enough. Windows update solutions will obviously only work for Windows, and the patches that it provides are only for the operating systems. Windows update does not include patches for applications. Your business critical Exchange servers, SQL servers, etc… are not included in these updates. Of course, those types of applications should be closely monitored and maintained by their administrators anyway, although this type of routine maintenance will often go overlooked if not automated. The recent SQL slammer worm is a perfect example.

Large corporate enterprises will typically need some form of centralized patch management software. They will need a system to scale to maximum capacity and provide support for multiple operating system platforms and their applications. Patchlink http://www.patchlink.com has done an excellent job in this area. Their product is very accurate, easy to use, and has good reporting abilities. The fact that the administrative user interface is web based is a big plus too.

Maintaining all systems with current service packs and patches is probably the easiest and most cost effective method of security breach prevention. It is not at all difficult to setup and configure an automated patch management solution, even if you're using one

of the free options. This simple task will vastly tighten your security exposure, even if you do nothing else.

Vulnerability Assessment

One of the biggest challenges facing security professionals today, particularly in large enterprises, is just figuring out what's on their network to begin with. After all, it's pretty difficult to secure something that you don't know exists.

Vulnerability assessment (VA) is basically an audit. Certainly before we can establish any effective defenses, we must first be aware of our own strengths and weaknesses. Taking a good look at ourselves and accurately identifying and quantifying our current level of security can be a very difficult task. We often have multiple divisions or departments or separate business units each responsible for their own resources and not always in communication with each other. It is vital that security across a corporate network be managed as a whole. A favorite tag line in the security industry is that "the risk assumed by one is shared by all."[vi]

Performing these types of security audits manually is not very practical at all in today's networked environments and is probably completely impossible in most large networks. Vulnerability assessments are almost always done with automated tools. There are several commercially available software tools on the market as well as open source or shareware. There are a couple of reviews that can be found at http://www.networkcomputing.com/1201/1201f1b1.html or http://www.infosecuritymag.com/2003/mar/cover.shtml which evaluate these tools comparatively. Some may work better than others and each tool may have its own unique features but ultimately, they are all after the same goal.

"Network-based scanners are excellent tools for evaluating security risks associated with two types: risks associated with vendor supplied software, and risks associated with network and systems administration."[vii] The idea is to have the VA tool scan your network or a particular network segment to first map it out. This is typically done by asking the VA tool to attempt to communicate with every IP address within a given range. The scanner will document every IP address that is in use and it will then try and determine the operating system of each of these hosts and which specific ports are open on each host. At this point, the VA scanner can run hundreds or even thousands of tests against each of these hosts to try and determine their level of security.

These tests can include anything from verifying that specific service packs and patches have been applied, to actually trying to guess passwords for accounts on those machines, to checking bizarre configuration settings on the OS or specific applications installed. VA scanners are compared against each other by examining a few aspects. Accuracy is critical to determine how well they actually identify the hosts and their vulnerabilities (including the proper remediation procedures). Speed is quite often of significant importance as well. You many not have weeks to wait for the scan results of hundreds or

thousands of hosts. Finally, good reporting can win big points for management. It is very important to be able to document these security scans effectively. More importantly, the ability to show a history of scans with the action taken to correct items identified each time can go a long long way in establishing your level of conscientiousness and due care, especially if you ever end up in court.

Once you're armed with this inventory of information, it becomes a great deal easier to ensure that your systems are identified and in compliance with the corporate defined security policies. This is undoubtedly one of the easiest ways to prevent successful hacks into your corporate resources. A VA scanner can get you way ahead of the game, as "hundreds of new vulnerabilities are being discovered annually, dozens of new patches are being released monthly, and thousands of systems are already behind the security eight ball."[viii] Don't wait to take action until after you've been hacked.

Hardening

So what exactly do people mean when they say that they're going to "harden" an operating system? This is an overused and often misused term. What hardening actually refers to is the process of carefully examining a given machine, removing or disabling any components that are not absolutely necessary for the operation of that machine, and then configuring very strict security settings and controls on whatever is left.

The idea here is to minimize the risk of attack by eliminating as much of the target as possible. For example, if IIS happens to be running on a given server, but that particular server has no reason to act as a web server, then I should definitely remove IIS from that machine and thereby eliminate any possibility of that machine being compromised via any IIS exploit.

This is certainly a very sound plan and seems to be common sense. However, the concept of so-called hardening seems to have gotten a bit out of control. It started with a few security configuration checklists, but now there must be hundreds of these step by step hardening manuals that go into excruciating detail on every possible little configuration setting. The fact of the matter is that by following one of these lengthy checklists, you're much more likely to cause the system to cease to function as required than you are to prevent a security breach.

A vast majority of IT and security professionals simply need a basic list of the most obvious operating system vulnerabilities and specifics on how to protect against these. Luckily, such a list actually does exist in the form of the SANS/FBI top 20 list which is maintained at http://www.sans.org/top20 and is updated regularly. This list identifies the "ten most commonly exploited vulnerable services in Windows, and the ten most commonly exploited vulnerable services in Unix."[ix] If you've implemented reasonable security controls at your other layers of defense, then simply hardening your hosts based on this top 20 list will put you in very good shape. "Although there are thousands of

security incidents each year affecting these operating systems, the overwhelming majority of successful attacks target one or more of these twenty services."[x]

As you get deeper into hardening specifics, there are two resources that are quite good. First is the Microsoft Security Configuration and Analysis snap-in. This is part of the operating system and comes with a number of pre-defined security templates for varying degrees of security. The tool allows you analyze your current security settings and compare them against any of the templates. You can also apply any of the templates to your machine to put all of those settings in affect. Secondly, the National Security Agency (NSA) has posted templates as well as some excellent documentation at http://www.nsa.gov/snac/index.html .

Security Administration

A frequently overlooked aspect of security management has got to be security administration. We tend to spend a lot of time and effort on these new security initiatives and we're almost always focused on how to keep the bad guys outside of our firewall or how to make sure our hosts are as hack proof as possible or we think about implementing complex intrusion detection systems. We spend an awful lot of time defending against an attack that may come someday. That is definitely a great thing to do, but we always seem to forget about putting security into our everyday administrative processes.

Yes I'm talking about our internal bad guys, or more commonly, our internal good guys that make honest mistakes. The truth is that we have a lot of low level people internally that have high level administrative authority within our systems. Think of all of the administrative activity that goes on every day. User accounts are created and deleted or modified. Passwords are reset. Group accounts are created and their members are adjusted. Access permissions on hundreds or thousands of files are modified every day, sometimes in a manner different than intended.

What if someone deletes the wrong user account accidentally? What if a help desk person changes security on a share or directory and inadvertently grants access to half of the company to the corporate personnel files? What if a conscientious system administrator updates an obscure file on the mail server and halts all email for six hours? Then what happens when that internal bad guy gets loose?

Just like in every other layer of defense, our goal is to minimize the potential threat as much as possible. We do this be eliminating any unnecessary vulnerabilities. We must restrict the number of administrative personnel to an absolute minimum. When we do grant administrative authority, we must grant only the minimum level of authority that is absolutely required to perform those duties. The rule of least privilege is extremely important.

The same rules apply when granting access permissions to any data files or other resources. Unless a person has a specific requirement to be able to edit a particular file, then they should have read only access to that file.

Achieving a very high level of administrative security can be a challenge. You have to weigh the benefits of very tight control versus a distributed workload and come to a happy medium. Once you find that medium, it is crucial to include these rules in your written policy.

Intrusion Detection Systems (IDS)

Intrusion detection systems are a relatively new layer of defense and IDS is still a fairly immature market overall. These systems go beyond just defining a set of communications traffic filters or just applying some specific configuration settings to a given machine or set of machines. The goal of IDS is to actively detect intrusions once they've occurred or even detect intrusions in real time as they are occurring.

Intrusion detection systems have a collection of knowledge built into them. They are programmed with varying degrees of intelligence to then use this built in knowledge to be able to recognize known patterns of attack from potential intruders. These systems are almost always separated into two classifications of IDS technology. There are host based intrusion detection systems (HIDS) and network based intrusion detection systems (NIDS). These are two very different and complimentary technologies. Jamie French describes a thorough overview of IDS at http://www.whitehats.ca/main/members/Malik/malik_ids_overview_files/frame.htm which also covers risk analysis and the pros and cons of each type of IDS technology. I agree with his assessment that any plan to implement an IDS solution must first start with a risk assessment. "Consider what resources are accessible and what the impact or cost might be if this were to be compromised or destroyed."[xi]

Network based IDS is done by actually watching network traffic on the wire and examining these packets for known patterns of malicious activity. NIDS is the most common form of intrusion detection and is often implemented as the first form of IDS in a given network. This is because NIDS is relatively easy to implement, it can monitor one or more subnets without having to go through a complicated installation and deployment process, and it's invisible to attackers and more cost effective than HIDS. NIDS is certainly not without its limitations though. It cannot see any activity at all that is not on the wire. A server or workstation may be compromised directly without the NIDS sensor ever being aware of the activity. Also, encrypted network traffic is troublesome for NIDS. NIDS may also be limited in the amount of traffic that it can analyze at a constant rate. High bandwidth may overwhelm a NIDS sensor. Probably the biggest problem that most organizations face with NIDS today is its tendency to trigger a large number of false positives. That is, it will raise alerts and identify normal network activity as a potential attack. This can often lead to the security administrator losing

confidence in the system's ability to accurately detect intrusions and may eventually cause correctly identified intrusions or intrusion attempts to go ignored.

Host based IDS provides a very welcome compliment to NIDS. A HIDS solution implements its IDS monitoring techniques at each individual host. In this scenario a piece of software is installed and running on every host that is being monitored. Since the IDS software is running as an application on the host, it can be keenly aware of all aspects of activity on that particular host. It will pay close attention to the operating system as well as other applications running on that system. It understands user behavior and can monitor their individual behavior as well. One of the very popular uses for HIDS is to monitor operating system log files for known items that may indicate suspicious activity. HIDS is typically much more accurate and much less prone to false positives because it tends to look for specific items and not necessarily just patterns. On the other hand, HIDS is much more difficult to deploy, configure, and manage because you have to maintain it on every monitored system. Also, since HIDS runs as an application on each of these systems, it can be attacked and compromised itself, just like other software can.

Together, both network and host based IDS are an extremely powerful combination. It would be rather difficult for an attacker to bypass both IDS systems. Each layer of defense greatly increases our security strength and using two levels of IDS within the intrusion detection defense layer goes a long way toward protecting our corporate assets.

Conclusion

I have certainly not covered every possible layer of defense. There are definitely other layers that provide significant benefits as well such as anti-virus software or personal firewalls. Of course I am also assuming that physical security has been addressed first. By definition, any machine that is physically accessible by an unauthorized person is not secure.

So what have we covered? My goal has been to help you define just what computer security is and how to implement a sound security infrastructure. I've outlined the defense in depth strategy as a layered approach to security, providing a form of fault tolerance. No failure at any single layer will result in a compromise of CIA – confidentiality, integrity, or availability.

As you've likely experienced, computer security can be addressed quite differently, depending on the point of view of the person being asked to secure an organization. If you ask ten people to give you their interpretation of what computer security means, you'll likely get ten different answers. So you may be wondering where to begin. Here is my basic advice:

- Always start with written policies outlining your computing goals and security requirements
- If you don't already have a firewall, get one

- Then start implementing a patch management and vulnerability assessment program
- Next, review all systems on your network against the SANS/FBI top 20 list
- Implement processes that employ the rule of least privilege regarding the daily administrative tasks within your organization
- Consider further hardening steps for critical systems
- Intrusion detection systems (IDS) should be last on your list of security initiatives

Finally, keep in mind that every computer network has the potential of being compromised. There is no magic security machine. There are only security professionals who can help design and implement effective policies to greatly reduce exposure and minimize risks. Remain diligent in your efforts to continually keep a step ahead of the bad guys. An ounce of prevention is worth far more than a pound of cure in security terms.

Resources

http://www.humanfirewall.org
http://www.shmoo.com/mail/ids/may01/msg00058.shtml
http://www.enterprisesecuritysolutions.net/IDS_presentation.ppt
http://www.whitehats.ca/main/members/Malik/malik_ids_overview_files/frame.htm
http://documents.iss.net/whitepapers/nva.pdf
http://www.securityfocus.com/infocus/1518
http://www.networkcomputing.com/1201/1201f1b1.html
http://www.infosecuritymag.com/2003/mar/cover.shtml
http://www.nswc.navy.mil/ISSEC/Guidance/infocon/Conop2-1.doc
http://www.nwfusion.com/reviews/2003/0303patchrev.html
http://www.patchlink.com
http://www.microsoft.com/windows2000/windowsupdate/sus
http://www.nsa.gov/snac/index.html

Bibliography

[i] Human Firewall Council. URL: http://www.humanfirewall.org/hfwcouncil.htm (16 Mar. 2003).

[ii] "Building a Human Firewall." URL: http://www.humanfirewall.org/smiinfo (16 Mar. 2003).

[iii] The Computer Information Center. "firewalls." URL: http://www.compinfo-center.com/netw/firewalls.htm (16 Mar. 2003).

[iv] "Terms Used in VPNs." URL: http://www.vpnc.org/terms.html (16 Mar. 2003).

[v] Andress, Mandy. "Windows Patch Management Tools." 3 Mar. 2003. URL: http://www.nwfusion.com/reviews/2003/0303patchrev.html (16 Mar. 2003).

[vi] "Concept of Operations for the Navy Component Task Force for Computer Network Defense." 29 Jan. 1999. URL: http://www.nswc.navy.mil/ISSEC/Guidance/infocon/Conop2-1.doc (16 Mar. 2003).

[vii] Internet Security Systems. "Network and Host-based Vulnerability Assessment." URL: http://documents.iss.net/whitepapers/nva.pdf (16 Mar. 2003).

[viii] Forristal, Jeff & Shipley, Greg. "Vulnerability Assessment Scanners." 8 Jan. 2001. URL: http://www.networkcomputing.com/1201/1201f1b1.html (16 Mar. 2003).

[ix] "SANS/FBI Top 20 List." Version 3.22. 3 Mar. 2003. URL: http://www.sans.org/top20 (16 Mar. 2003).

[x] "SANS/FBI Top 20 List." Version 3.22. 3 Mar. 2003. URL: http://www.sans.org/top20 (16 Mar. 2003).

[xi] French, Jamie. "Intrusion Detection, Overview of the Technology." 13 Nov. 2002. URL: http://www.whitehats.ca/main/members/Malik/malik_ids_overview_files/frame.htm (16 Mar. 2003).