GIAC
CERTIFICATIONS

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.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials (Security 401)"
at http://www.giac.org/registration/gsec

Low Budget Network Security for The Small Enterprise – The Basics and Beyond
(A Practical Guide for Managers with a Limited Technical Background and Resources)

Abstract

For the small enterprise, network security usually takes a second seat whenever managers establish and/or evaluate their business priorities. Often times this is due to the fact, whether perceived or actual, that cyber security is too costly, too complicated, and really not necessary. After all, who would be interested in compromising the computer-based assets of, for example, a local hardware store anyway? After a short initial evaluation, a strong business case will be presented for having some varying degree of cyber security. This document will lead the non-technical manager/owner, in a step-by-step fashion, through a three phase process (evaluation, implementation, and review) needed to secure computing assets in an efficient and economical (read low or no cost) way. Each of the three implementation stages builds upon the previous one, and brings with it an enhanced level of security. Along the way, the manager is given, in a clear and understandable fashion, the justification for the process, and the tools needed to accomplish this goal. It is only by understanding the threat and the resulting cost of inaction that the decision-maker can fully appreciate the need for a proactive approach to protecting the enterprise's computer-related assets.

Introduction

Hardly a week passes without a media headline detailing some sort of catastrophic cyber event affecting a well-known business or a governmental entity. [1] The details of the incident might include such things as corporate data corruption, theft of confidential customer records, or secrets being leaked to competitors or enemies. Losses are often estimated to be in the millions, or even in the billions of dollars. Fortunately for some, that only happens to large corporations and similar targets! A point of fact? Or a popular misconception? Consider for a moment the thousands upon thousands of small business entities in the US alone that depend upon some type of computing platform to conduct all or part of the daily activities. It would be very difficult to find an environment that is totally immune to a computer security-related incident. While dollar losses might be on the order of six figures for the big guys, the cost might represent only a small percentage of their total net worth. However, losses amounting to 1, 2, or 3 orders of magnitude less could be devastating, if not disastrous to the little guy. So, how important is cyber security to a small business? This question and many others will be discussed in detail in the following guide, as well as ways to minimize risk with limited resources.

Evaluation Phase

For the sake of this discussion, a small enterprise will be loosely defined as any entity that has less than 50 employees, that uses computers and/or computer networks to perform any part of their charter, and that has little or no budget or staff for IT-related processes. Examples of small enterprises might include such things as a SOHO (Small Office-Home Office) environment, a local municipality, a mom-and-pop storefront, a public research institution, a non-profit, or a rural school district. This group is unique in

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

that the bulk of their productivity depends on the efforts of just a few employees, and they are generally owned or operated by someone or something with regulated or limited financial resources. Managers often question the need for any computer security outlay, especially in light of today's narrow operating margins, decreasing budgets, and shaky economy. They feel that they can't afford it, both in terms of capital outlay and personnel costs. The premise of this discussion is that they simply can't afford not to implement some level of computing security. Why? Let's start the evaluation process by asking ourselves a few important questions.

First, **can we identify those business functions that depend partially or entirely on a computing infrastructure and can we quantify the time spent using those items?** Of course, we immediately think of such daily activities as sending and receiving email, writing a grant proposal or a test, entering data into spreadsheets for accounts receivable, checking inventory, looking for business resources on the web, and keeping an up-to-date list of our customers. All of these functions are usually performed on the venerable desktop or laptop computer. But what about the POS (point of sale) cash register, the credit card reader, the data logger on the lab experiment or machine lathe, the fax machine or postage meter with dial-up capability, the cell phone, or the system controlling the building heating and cooling? These, too, are computer based, and are, most likely, connected to a network of some sort. Many other examples abound, depending on one's particular situation. Stand alone or "dumb" business devices are a thing of the past. Each manager should make an inventory of their cyber resources, and consider their role in the overall business strategy.

Second, **would we be able to continue operating if most (or even some) of our computing resources were compromised?** Scary as the thought might be, reports of just such occurrences can be found affecting every part of the small enterprise community. Threats to the security of the business computer assets come in many and varied forms. For example, we have all heard about some of the high profile break-ins making the news lately. They are certainly noteworthy, and they have wide ranging implications for the entities involved. But what about an email virus that sends pornographic messages to the first 50 entries in your computer's address book? Imagine the irreparable damage done to your associates or customers. Think about a RAT (Remote Access Trojan) application, unknowingly installed when one of your employees visited an untrusted, non-business web site during their lunch hour. That hidden application could allow an outsider to take control of the system, sending your customer database and credit card numbers to a remote user for fraudulent or identity theft purposes. Then there is the "ugly" piece of code (unknowingly installed with some "borrowed" software) that has the ability to wipe out your entire hard drive, thereby loosing 3 years worth of data from that research project on which you have been working. Finally, consider this scenario. Because your unprotected computer (or network) is always conveniently connected to the Internet, it has been compromised, unbeknownst to you, in a way that allows it to be controlled by a remote machine. Your "zombie" computer is being used in a DDOS (distributed denial of service) attack against an important military facility.[2] The point is, whether you think you are

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

vulnerable or not, the potential for system compromise is real, and the probability is extremely high.

Third, **what would happen to our customer/employee trust if we were unable to assure them that they were not harmed by the event?** Needless to say, years of relationship building could be lost in an instant should the parties involved even suspect a breech of confidence. Their reaction would be akin to the feelings that one experiences after a burglary: uneasiness and violation. Employees, fearful of their upcoming performance evaluations, might worry about whether their email correspondence or online order placement is making it to the appropriate recipient. Those outside of the organization might be concerned about providing such things as additional proprietary research information or internal email addresses for fear of being involved, directly or indirectly, in a future attack. Recovering from those losses of confidence can be extremely troublesome, and can involve a substantial investment in time and effort.

Fourth, **what would be the cost of rebuilding the computing infrastructure?** It might be relatively easy to calculate the expenses associated with such things as replacing hard drives, providing for additional security measures, rebuilding unstable machines, and reloading critical data from backups (assuming that they exist). It would be much more difficult, however, to estimate the cost of overtime labor, the loss of employee productivity for those individuals not having access to computing resources, the loss of business revenue during the downtime, and, of course, the long term cost of the aforementioned loss of trust.

Fifth, **would our insurance cover any of the loss?** Part of the reason for insurance in the first place is to have help in recovering from an accident. Just as auto insurance has exclusions relating to such things as driving under the influence or operating a motor vehicle with intent to harm, most business policies have similar exclusions for lack of "due diligence". Minimum Standard Practices usually state that the enterprise manager/owner should have in place those processes deemed reasonable and prudent to protect the insured business assets. Chances are, an adjuster would not allow a break-in claim if she found that the front door to the office didn't have some sort of lock. Similarly, a claim resulting from a compromised business network that doesn't have in place even the most basic of protection strategies is also destined to be denied. Taking a proactive approach to security can minimize uninsurable losses and facilitate timely adjusting when the need arises. And, like an insurance policy, the real value of a computer security plan is usually realized only when, not if, you have a breech. The rest of the time, it just waits quietly in the wings.

Unfortunately, the answers to these questions are much different today than they were just a few years, or even a few months ago. Computers and computer networks have become, directly or indirectly, an integral part of our business and personal activities. Because any unnecessary loss can indeed be devastating, security considerations are just as critical for the small enterprise as they are for the multinational corporation. Although it is easy for the non-IT person to find numerous technical articles describing

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

the “why” of good security practices, reading these often yields nothing more than a sense of confusion and utter despair. What many of us need is a step-by-step guide detailing what to do, and in what sequence to provide the most ‘bang for the buck”. A good example of this approach is found in basic CPR training. While the trained paramedic or the ER physician certainly has knowledge far beyond the fundamentals, the timing and efficacy of the first responder’s initial ABC (Airway, Breathing, Circulation) technique is without an equal. Granted, it is important to have an appreciation of the theory and technology used to save lives. Far more significant for the lay person, however, is knowing the correct steps to take, and in what order to apply them.

So it is with computer security. We have already looked at the strong business case for having some sort of protection process in place. But unlike the reactive nature of CPR, this process is designed to be proactive in saving the life of the enterprise. You don’t have to be an expert to learn and practice the ABC’s of cyber safety. All that’s needed is an appreciation of the risks involved in doing nothing, and the steps needed to prepare and execute a security plan. Can one be implemented with minimal capital outlay and staff time? Most emphatically, “Yes”. Let us examine the components of a viable security plan, and define the various phases of implementation, each of which is tied to the value of the assets involved, the level of technical expertise available, and the budgetary constraints in place at that time.

Phase One - The Minimum Acceptable

All businesses need, and probably most have, some sort of COOP (Continuity of Operations Plan) or business continuity plan to provide for a disaster, be it simply a broken water pipe, a fire, a break in, or even a total loss of the building and records. This is an excellent place to start developing a **cyber security plan**. In fact, protecting the enterprise computing assets should be an integral component of the COOP, and is just as important as protecting the physical assets. Both parts are needed, and a loss of either one usually means business activities are temporarily or permanently restricted. The primary focus should be twofold: guarding your assets from accidents, and keeping outsiders from intentionally accessing or damaging them.

A significant first step in establishing a security plan is the design of a written security policy. About now you are probably saying to yourself, “Oh no, not another piece of worthless paper”! However, the importance of a basic security policy cannot be overstated. In fact, this document goes a long way in meeting the insurance community’s minimum standard practice guidelines mentioned earlier. The contents must provide a simple, concise, and understandable description of those actions that are, or are not allowable, and the consequences for failure to comply. Included should be such items as acceptable web and email activities, permissible use of the computer resources, both during and outside of work hours, whether or not any non-business software can be installed by the user, responsibility for data backup and control, and consequences if the directives are not followed. The policy statement should be signed

© SANS Institute 2003, Author retains full rights.
© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

by each user, refreshed using occasional reminders, and updated regularly as the business climate changes. The mandatory adherence to the elements of this single document, unimportant as that may seem, can eliminate or minimize a large percentage of the potential security problems in most small enterprise settings.

Accompanying the written security policy should be a handout or booklet detailing the basic operation of the computers or computer-related equipment in use, as well as a review of computer etiquette. It can be quite basic or very detailed, depending on the needs of both present and new employees. Included might be a description of the steps involved in performing an orderly shutdown, a method to follow if the system locks up while in an application, a review of the proper way to safely manage and delete files, and a step-by-step procedure to document a problem. A prudent manager soon learns never to assume the knowledge or skill level of a new, or even a long time employee. It is better to review some established guidelines and run the risk of insulting someone's intelligence rather than to find out after the fact that important data has been lost or compromised by improper handling.

Having dealt with these two essential tasks, the rest of the steps necessary to achieve a minimum level of security can be instituted in almost any order. Early in the game, however, it is imperative to provide for virus protection, and for some sort of barrier between the local machine(s) and the outside world. Antivirus software should reside on every piece of computer equipment possessing an identifiable OS (operating system), i.e. Windows, Linux, Unix, Macintosh, IBM OS2, etc. It is relatively inexpensive, and can help protect against some of the most prolific and devastating computer system hazards. More important than having the software installed on the machines, however, is having a defined process to keep the virus definitions up to date. Because new viruses and worms are released regularly (often weekly or daily), overlooking the update step can lead to a false sense of security while leaving you vulnerable. Old definitions are almost worthless. Most commercially available antivirus applications have provisions to retrieve updates from a corporate server, but they normally require that the user do so on a regular basis. Again, this step should be detailed in the handouts given to all employees.

When it comes to protecting the family home, one often installs some sort of perimeter security system to guard all of the doors and windows from intruders. Because a house has many points of entry, just having a lock on the front and back doors is insufficient. Consider, for example, the sliding patio door, the flimsy basement window, or the inexpensive garage door opener. Each provides the skilled burglar easy access. Computers, too, have many points of entry. On the surface, it appears that the sole office PC or the entire network has only a single link to the outside world, usually in the form of a dial-up line or a high speed Internet connection. It's similar to a two-lane bridge leading to an island with a city full of streets and buildings. Once the single stream of traffic makes it to the island, each vehicle, in the absence of a good map, may end up traveling to many different addresses in the town before the correct one is located. In reality, there are approximately 65000 openings (virtual doorways) into the heart of today's computer systems. Just as easily as prying open that basement

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

window, an attacker could use many of these ports to gain access. When it comes to protecting the enterprise's computing infrastructure from outsiders, the recommended practice is to install some type of programmable router or a firewall. It serves to establish control over those points of entry based on a set of guidelines and on rules that you define. While all of this sounds extremely complicated and costly, a very inexpensive firewall router ($50-$300) can provide reasonable protection for most small business environments. By limiting or totally blocking access from the outside world, and by controlling when and why certain ports will allow traffic in, the firewall, usually in the form of an appliance, provides an element of isolation and measurable defense. It isn't necessary to have extensive technical training to implement a firewall. Most come pre-configured to provide a moderate level of protection. Final configuration, based on local networking requirements, can be accomplished by following the setup wizard included with the devices. For single computers, free or inexpensive local (host based) software firewalls can provide excellent protection. It isn't prudent to permit any computer or network to have direct contact to the hostile world of the commodity Internet. It is just an accident waiting to happen.

Firewalls go a long way to protecting the network or the individual computer. But, they are not perfect. When a trespass does occur, wouldn't it be nice if the intruder couldn't find anything to exploit. All major commercial and open source operating systems, most programming languages, and some hardware control environments have inherent weaknesses. Some of these can be taken advantage of by an unscrupulous person, often with disastrous consequences. Many of these vulnerabilities have been identified, fixed, and the repair released to the computing public in the form of patches and security updates. By far, one of the best and least expensive ways to protect your computing assets is to regularly and systematically apply those patches from the hardware and software vendors. Most operating systems, many applications, and some hardware devices come with routines that provide for automatic or manual updates. All you have to do is identify those systems that can be updated, and schedule the process. Because new vulnerabilities and exposures are constantly being identified, sometimes on a weekly or daily basis, patching has to be just like any another regular maintenance task. Failure to do so is equivalent to handing over the keys to the castle.[3] Not a good idea!

Most of this discussion has been centered on keeping the bad guy out. But what if the intruder was coming from the inside?[4] Not realistic you say. Think again. Statistics clearly show that cyber security breeches are more frequently associated with internal employee/user actions rather than with external attacks. It could be something as simple as an after hours cleaning person sitting down at an unprotected system to visit a malicious web site, a curious employee trying to access personnel files without the appropriate authorization, or a disgruntled worker attempting to sabotage an important database. These and many other scenarios can yield as much devastation as an attack from outside the network. Also, consider those situations where, even though no malice was intended, someone accidentally overwrites or alters critical files they should not have been able to get to. What we need is a simple yet extremely effective way to

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

protect computer-based systems, one that works on either side of the wall (firewall, that is).

Enter, stage left, the almighty and infamous password. Although the technology press has both praised and maligned password usage, the fact remains that it is still one of the most efficient and powerful tools at hand to safeguard systems. More sophisticated tools and techniques are currently available (biometrics, tokens, authentication servers, certificates), but they are either too costly or too complex for the small enterprise setting. The strength or weakness of a password policy depends entirely on the implementation scheme.[5,6] With the proper tricks and tools, most employees have little trouble embracing the concept. The value of the resource should dictate how and when a password is used. For starters, all machines should have some type of password protected screen saver with realistic time values. Most operating systems, including the early versions of Windows, have such a provision. This prevents quick access by the casual passerby. Because it can be easily circumvented, the logon password on Windows 9X machines is of little or no value, and should not be considered as a security measure. A more robust solution can be found in Windows NT, 2000, XP, and most flavors of Unix, and should be implemented to the fullest extent possible. Logons, whether local or network, and access to files can be securely controlled with proper password usage. Likewise, routers, switches, data loggers, and other computer-based devices should be configured to require password logon. Many excellent treatises have been written on the care and feeding of passwords. Details of such, however, are outside the scope of this discussion. Suffice it to say, strong passwords and a strict password policy are at the heart of a low or no cost enterprise security plan.

Even the best of security plans are only as strong as their weakest component. Maybe that is why a significant percentage of computer system compromises can be traced back to human error, either in judgement or in training. Tricking a person into giving away sensitive information about the enterprise environment has become known as “Social Engineering”. Often this can be in the form of a request for a “lost password”, the IP or port address of a critical server, or the brand of database program which stores customer profiles. Each bit of information helps the would-be intruder. By taking advantage of specific probing tools and scripts gleaned from a myriad of web sites, and by utilizing the CVE (common vulnerabilities and exposures) list mentioned earlier, the perpetrator can tailor an attack on a system in an extremely short period of time. It is a little like the “What’s My Line” quiz game that was popular in the early days of television. If the answers to the questions are in the negative, the contestants have a very hard time narrowing in on the person’s occupation. But, just one positive response can trigger a series of directed questions that quickly lead to the correct answer. Social engineering can take many forms, but the results are almost always the same: valuable information in the wrong hands. Employee awareness and training, an easy no-cost step, can all but eliminate the threat. This is yet another justification for a written security policy and the employee handouts.

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

Like a lot of things in life, we only fully appreciate something after we have lost it. Because Murphy's Law seems to be pervasive in the world of computers, one must learn to expect losses, be they intentional or accidental. At the heart of all cyber systems is digital code, in one form or another. Examples are spreadsheet files, letters to the graduate school, pictures from the lab picnic, important business proposals, configuration files for the firewall, the software applications that make one's life more efficient, and, of course, music to soothe the savage beast. Since any or all of these items could easily disappear in a heartbeat, we need a way to protect and restore them when disaster strikes. Enter, stage right, the time-honored backup. Rebuilding systems after a cyber accident has been likened unto a religious experience. We can develop an appreciation of backups, either by "getting the faith", or learning from the mistakes of others. Even though we recognize the importance of regular backups, we often relegate them to the bottom of the to-do list because they seem too complicated or too time consuming. Granted, getting started is hard. But, the process can be broken down into small, easily digestible parts.

Nearly all of our digital assets are found in two main forms: those that change infrequently, and those that are modified on a regular basis. Once most physical computing resources are configured and placed in service, they change very little. Typical examples include a PC with its operating system and all necessary software applications installed, a network printer with a hard drive holding extra font sets and emulation code, a firewall with ACL's (access control lists), and a data logger programmed to process inputs from multiple lab instruments. Making a full backup, a digital snapshot if you will, of a stable, working system needs to be repeated only when the functional components are modified, updated, or upgraded. Once the backup or image is taken, often using tape or optical media, that part of the process is finished until something changes. If an accident or compromise disables the device, it can be quickly and easily restored with a minimum of downtime. For major backups, best practice standards suggest using 2 forms of media or 2 different devices, the thought being that if one recording device malfunctions without our knowledge, the other would provide the necessary coverage. This step has a secondary benefit. Moving one copy to offsite storage helps to meet one of the requirements of the COOP discussed earlier. The key point is that a lot of work goes into setting up the hardware that is used to process and store our intellectual efforts, and protecting it just can't be overlooked. Again, free or inexpensive software (<$40) is readily available for copying or preferably imaging the overall system setup.

The other portion of our digital assets, that which is subject to frequent additions or modifications, is referred to as "the fruit of our intellectual talents", and is known simply as "data". Because of the dynamics involved, a tailored strategy is needed to ensure that our creative efforts are backed up at a frequency commensurate with their value. Initially, all of the data on the system needs to be replicated using a process known as a full backup. While this can be rather large and time consuming, it needs to be done only infrequently. Then, at an interval to be determined by established evaluation criteria (such as criticality or ease of replacement), only the files that have been changed (emphasis on change) since the last full backup need be copied. This process

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

is usually very quick. Depending on the business needs, those files might need to be replicated fairly frequently. Like most strategic processes in our lives, the ones that are scheduled are the most likely to be completed. Think about that house payment which is withdrawn monthly from your checking account, or that sprinkler controller which ensures efficient lawn watering. So it is with backups. For those systems that have some sort of schedule function (most PC's do), performing regular backups is relatively straightforward. For other devices, manual backup is usually the only option. But, like the dentist's reminder postcard, posting a reoccurring note in a PDA, day-timer, or office calendar program is the key to success. Lastly, as with any good emergency plan, regularly practicing the restoration operation and verifying the data ensures success when the actual need arises.

A few final suggestions will round out this discussion of the "basic" cyber security implementation. The first has to do with email. Convenient as it might be, business email systems should not be used for personal communications. Each unnecessary message brings with it the potential for problems. SPAM (unsolicited email), viruses, proprietary information leaks, objectionable material, loss in productivity, and resource overload are just a few of them. Employees should be encouraged (or required) to have separate, personal web-based email accounts from another service provider, and to only check them using work machines if allowed or necessary. Minimizing the number of published business email addresses also cuts down on unnecessary traffic. Secondly, just because a network is in place does not mean that all of the computers need to be a part of it. Connect them only if or when they have a need, such as updates, transferring data, or backing up to other systems. Otherwise, disconnect them and keep them isolated. They can't be a target if they are not visible. Finally, casual local hosting of web, email, and FTP (File Transport Protocol or file sharing) services dramatically increases exposure of internal resources. Rather than making internal files accessible to the outside world, consider copying them as often as necessary to inexpensive writeable CD's and distributing them to employees who work remotely using snail mail or overnight courier. Contractor or ISP (Internet Service Provider) hosted offerings provide a distinct layer of separation between the outside world and the internal environment. They also have the tools to do a better job of inspecting and screening the incoming traffic. Evaluate password-protecting a part of the externally hosted web site. This area could be used for low security data that is not necessarily public but is needed for remote employees or for after business hours access. Most of all, do not allow in any traffic other than web and email without taking the steps discussed in the next section. Nice segue, huh.

Phase Two – The Intermediate Approach

The next implementation phase could be called Cyber Security 102. Let's pick up the pace a little because, hopefully, we have a clearer understanding of the inside world, the outside world, and the problems associated with their coming together. Speaking of the outside world, there seems to be a direct correlation between the growth of the enterprise and the increase in use of mobile computing technologies. We

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

need to address the dilemma of protecting the data on laptops, PDA's and other systems when they are not in the office. It is all too easy to overlook the business value and the irreplaceable nature of information found on mobile devices. Likewise it is often difficult, if not impossible, to make timely copies of that information. Of course, the most straightforward approach is to replicate the data, routinely and automatically, on higher capacity removable media such as CD-R or flash memory. Should those options not be available due to the age or configuration of the systems, another popular option is to use a web-based solution. Whenever an Internet connection is available, a software agent installed on the computer will connect to a designated web site and will backup any new or modified data. This has the added advantage of storing a copy of the data in a remote location. A secondary benefit can be realized by using the service to provide additional offsite storage for the in-house systems as well. Online storage can be economical and reliable, but only if the appropriate security precautions are taken. Careful evaluation of the service provider's pricing structure, reliability specifications, and encryption standards is a must.

Loss of availability of the mobile data isn't the only security concern. As was mentioned earlier, a computer security plan has two components: preserving what we have so we can continue to operate, and protecting it from use or abuse by others. What if the portable device is simply stolen or misplaced? Over and above the obvious economic burden from the lost hardware is the cost of dealing with the misuse of the information. Credit card numbers, proprietary research data, online account PIN's, and private correspondence can all wreak havoc if they fall into the wrong hands. Just as in the office, we need to provide for both the physical and the cyber security of the asset. Low cost solutions to the problem range all the way from encrypting the important data, attaching security cables to the chassis, maintaining critical files separate from the device, to guarding access with tokens and biometrics. The value of the assets determines the types of measures to deploy.

While we are still considering the domain outside of the office, let's glance back at that all-important continuity of operations plan. Now is the time to refine it a bit. It has always been important to be able to have access to the business's data should a disaster strike. But it is also important to make sure that it is possible to actually use it in a timely manner. This is especially true if you are forced to operate apart from the primary facility and are unable take your systems with you. If the data was backed up on a device such a tape drive having a proprietary recording format, or on a device that requires a unique interface card, it is imperative that the same hardware be available at another location. A recommended strategy is to have enough similar equipment off site to cover the initial stages of disaster recovery. For example, having a mirror configuration of a critical office system in the home of a remote worker or manager would provide for partial operations continuity, especially if the duplicate records were also stored there.

Now, let's go back inside the business walls and see what can be done to strengthen our basic security foundation. Security issues were relatively unimportant when developers were writing earlier versions of today's popular operating systems.

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

However, as we mentioned before, most have major weaknesses that have been identified and exploited. Because cyber safety has become an increasingly important requirement of the computing community, the newer OS's, such as Windows 2000 and XP, Linux 2.4, and Mac OS X, have been written with higher levels of protection in mind. Not only is the newer code more secure, but it is more likely to be patched in a timely manner. Upgrading or replacing older operating systems is one of the most cost-effective ways to bolster the enterprise.

One of the easiest ways for problems to extend throughout the network is to have open file shares on some or all of the systems. While it is extremely handy to click on the Network Neighborhood icon to look for a particular file that you need, that very same convenience can allow malicious code on one machine to quickly spread to others. Even though the shares would still be visible, simply password protecting them can go a long way toward slowing down or eliminating trouble. Of course, the stronger the password, the safer you are. Perhaps now is the time to strengthen your enterprise basic password policy by requiring longer, more complicated passwords, and changing them more frequently. To add additional security, a more robust OS can be configured not to show any of the shared resources. They allow for the creation of individual protected accounts on remote machines, thereby permitting access only to designated files. Yet another reason to upgrade when time and finances permit.

Another simple step to reinforce local security, be it on an individual machine or the entire network, is to log off or disconnect from the Internet when outside access is not needed. The less time your system is exposed, the greater the level of protection. Some networks can be configured to drop the connection after a defined period of time and reestablish it when a request is sent. This step can also generate a cost saving if you are using an ISDN connection, or are on metered telephone service. Always-on services such as xDSL or cable modem can be manually disconnected at the firewall for after-hours protection, especially if the business computers remain on all of the time. Particularly important when using dial-up modem Internet access, whether primary or backup, is having the auto-answer function disabled. An improperly configured modem can provide a convenient and easily exploitable "back door" into your system(s). Most communications software programs install with the auto-answer function turned on by default. My grandmother always said that closing the door keeps those pesky "critters" out. I guess she was right.

Speaking of doors, even the most secure computing environments have a couple of open ones. Standard firewalls can be set up to block the vast majority of traffic coming into the network. However, some inbound traffic has to be permitted; the most obvious of which is web and email. Many of the hazards discussed earlier enter the network through port 80 (HTTP – HyperText Transfer Protocol or web) and port 25 (SMTP – Simplified Mail Transport Protocol or email).[7] Part of the reason for our security policy placing restrictions on web activities is to minimize the risks associated with indiscriminate and unnecessary surfing. Because the request for an outside connection originates from the inside, most simple firewalls do not block the return traffic, laden as it might be with "critters". Port 80 traffic is relatively unprotected, and is difficult to police

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

without some rather sophisticated tools. Screening HTTP data is best left for another level of security implementation. Many firewalls, either hardware or software-based, do have provisions for allowing connections to specific web sites only, blocking specific URL's, or a combination of both. These settings, in association with the written security policy, can be of great benefit in mitigating port 80 risks.

If the situation and the finances warrant, increased protection can be found in the form of content filtering. An intentional or unintentional visit to a rogue web site often comes with a stiff price, usually in the form of malicious code downloaded to your machine. If you can't get to that dangerous site, you don't have to worry about bringing back a destructive payload. Several firewalls can be licensed to include an integral application which looks at the web request, compares it to a continuously updated list of unacceptable sites, logs the request, and either passes or blocks the traffic. The costs usually range from $50-$500 per year, depending on the number of clients needed. Also, the decision process can be made even more granular by setting time windows when certain web sites may or may not be accessed. This gives the manager much better control of web activity, with the added benefits of increasing employee productivity and ensuring that objectionable material does not make it into the workplace. If the current firewall does not support such an addition, a good option is to replace it with one that does, and use the existing box to provide an additional layer of protection for a critical enterprise asset such as a customer database.

This is an appropriate time to discuss an often neglected but extremely important feature of most firewalls, the logs. Whenever inbound traffic is blocked, an entry is usually made in an internal database. With implementation phase one, we were happy just to have a firewall in place. Now we are ready to take advantage of that extra information to give us a heads up about a possible exploit. Most logs can be exported in a form that works well with a spreadsheet. Having a technically oriented employee periodically review, sort, and archive the data can have a big payback. With a little practice, he or she can learn to separate out the background noise from an attempted or an ongoing attack. Foretold is forewarned. When an actual break-in happens, forensic investigators may be able to extract valuable clues from the logs.

Remember the discussions concerning the imperfect nature of firewalls, the pesky critters, and the dangers associated with port 80? Even with the best perimeter defenses we can afford (virus detection, firewalls, isolation, etc.), some malicious agents will probably make their way to the individual elements or hosts within the network structure. Another no-cost or low-cost tool that can help guard the network (or stand-alone PC for that matter) is the host-based IDS (Intrusion Detection System) or personal firewall.[8] By wrapping each exposed or critical computer with an even more restrictive and finely tuned layer of protection, we can severely limit the power and destruction of such agents. Without going into a lot of detail, just know that host-based IDS sensors can totally or selectively block all traffic, both incoming and outgoing, while logging any attempts to connect. They provide a very powerful and effective method of protecting the key components in the network. As we will see later, the whole process can be automated to ensure efficient use of the tool and to minimize labor costs.

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

One additional tool that can go a long way toward achieving network or workstation stability and safety is a host-based malware (malicious software) scanner. This inexpensive software utility takes over where antivirus programs leave off. Some of the hazards described earlier (RAT's, key loggers, data leaks) can piggyback on free software downloads or interesting screen savers, and are undetectable by normal means. Even the most innocuous and well-intentioned web sites can silently push bits of exploitable code onto our machines. A lot of it is unavoidable. In some cases, these little robots (or simply "bots") can then be externally directed by other sources to perform questionable tasks. If allowed to operate unchecked, they can, for example, track and capture an unbelievable amount of private information (passwords, credit card numbers, web site visits), and leak that information to clandestine sites on the Internet. Other times, system instability is the net outcome. Scanning the individual network hosts to locate and remove this extraneous code is akin to our annual spring housecleaning ritual. We focus on checking all of the cracks and crevices for outside material that doesn't belong there. The big difference is that the digital debris can accumulate at a much higher rate, and can cause a lot more damage. Regularly scheduled malware scans, in conjunction with virus scans, can contribute significantly to the overall security of the computing environment.

Although Cyber Security 102 appears to be somewhat detail oriented, the beauty of it lies in the fact that, once most of the steps are in place, it becomes a set-and-almost forget process. Most of the tools and techniques discussed need be implemented only once and monitored only infrequently. While there are minimal costs associated with some of the later steps, the ROI (Return On Investment) can be high. Probably, by now, you have become just a little paranoid. That's good. And, hopefully, security considerations have or will become an integral part of the way you do business. They are just as essential to the success of the enterprise as are the products or services produced. But, we are not finished yet. There are additional things that can be done to increase the safety factor. The topics in the next section are still governed by the same guidelines as before: you don't want to become a security expert, and you have little or no budget to direct toward the effort! Read on.

Phase Three – Making a Good System Better

Human nature prompts most us to deal with the crisis du jour first, and sideline those functions that can wait for another day. You <u>really</u> need to change the oil in the car, but getting the kids to the soccer game takes precedence. Maybe tomorrow. It's just a common faux-prioritization step that we all take. In reality, it is not always the correct decision. Wouldn't it be nice if the car could change its own oil? Maybe someday. Wouldn't it be nice if network security issues could look after themselves? Well, some can. Automation is the one procedure that helps prevent key security processes from being brushed aside when things get crazy. Automation is all about making essential tasks as mindless as possible. Automation is our friend. Let's look at a few of the ways we can use automation to make a good system better.

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

Using a predetermined schedule, many off-the-shelf antivirus software packages provide for the regular updating of virus definition files, and for the scanning of local files. As was mentioned earlier, however, the burden of configuration and maintenance falls on the user. Theoretically, once everything is setup, nothing else is required in order to be fully protected -- until someone disables the scanner to install new software, or until the license runs out. Granted, it is possible for you to schedule a visit to every PC to check on the status of the setup. Not very realistic, though. One solution is to centrally manage the individual systems using a version of software designed for network administration. Most of the major antivirus vendors have such a package, and it can be easily implemented without the need for a dedicated server.[9] One of the system machines with a little more horsepower can realistically act as the management console. Because it is automated, not only is the network approach more efficient and secure, but it is often much less expensive than the cost of buying several stand-alone licenses. All machines are updated and scanned regularly, and system messages can alert someone if there is a problem. The same technique can be used to manage the aforementioned IDS sensors and the malware scanners found on individual machines. In a similar fashion, the network versions of these tools do not require a dedicated server, and do provide for management alerts. If there are more than three or four machines on the network, the logic of this approach becomes readily apparent.

Automation procedures can also help reinforce some other processes. Several PC-based applications are capable of providing more contemporary data replication. These inexpensive COTS (Commercial Off-The-Shelf) software packages can send real-time or near real-time backups of mission-critical data to another hard drive in the machine, to a local file server, or to a remote location. They can be run on individual PC's without requiring any special hardware. Because these programs operate in the background, they continuously monitor designated files or directories. Based on the configuration settings, they can backup those files whenever they change, or perhaps once an hour, once a day, once a week, or anytime in between. This approach fulfills two main cyber security requirements: ensuring that valuable, and sometimes irreplaceable, business data are protected by creating at least two copies, and providing for a mechanism to store the data in a remote location, be it on another machine or totally off site.

Hopefully, by now, you have upgraded most, if not all, of your machines to more secure operating systems. These OS's usually have system event logging as one of the preinstalled services. The nice part is that logging is virtually automatic, the logs are easy to view, and they can provide a wealth of information regarding stability and security issues. The down side is that, without some form of automation, the system logs are of little value, mostly because no one ever looks at the data. One way to circumvent this problem is by using a third party software package to collect the log files from individual workstations and servers, to sort the entries, and to provide a report highlighting any important events or trends.[10]s From a security standpoint, the system logs can be invaluable. By pointing out unusual attempts to connect, newly installed services, and security setting modifications, the log files help to prevent problems and to aid in recovery efforts.

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

One additional benefit derived from reviewing the log files is the opportunity to study the number and the types of services running on the machine. Although some of them are necessary for proper operation, others serve no real purpose in that particular configuration, and can create major security holes. In fact, many of them are enabled by default. The process of removing or disabling unneeded services is known as OS hardening. This procedure is absolutely essential in minimizing the damage from a compromise. Like the individual file shares mentioned earlier, if the perpetrator can't see or manipulate the vulnerable service, he can't inflict any damage. While there are several excellent treatises on the subject, the details of the hardening process are beyond the scope of this guide. Consider hiring an outside consultant to help with this important point. Spending just a little money now can yield a big payback later.

A no-cost step that can have a substantial impact on system security evolves around implementing a stricter password policy. The rule of thumb for passwords is "make them complex, and change them often". Up until now, a password provided only a moderate amount of protection for the enterprise assets. For it to be truly the first line of defense, the password needs to go on steroids! First off, the length needs to be at least 8 characters, but preferably more. And it should contain a mix of letters, numerals, and punctuation marks. And it should not have a sequence of characters that can be found in any dictionary. And it should be changed every couple of months. And ad nausium. About now, you are probably saying to yourself, "Who are you kidding? I can't even remember my mother's birthday!" Fortunately there are several tools and tricks that can make the process more tolerable and enforceable (they don't include writing it on a sticky note stuck to the monitor). Again, the theory is beyond the scope of this guide, but the tools are readily available on the Internet.

Another neat (and cheap) way to keep the bad guy at bay is to eliminate file shares on individual machines. The reason behind this procedural change is to keep a break-in from spreading. Remember that accidents happen, and when they do, it is best to contain the damage to as small an area as possible. Some of the most devastating viruses and worms propagate themselves by jumping from one shared machine to another. Entire networks can be compromised in a matter of minutes. Sharing files is best accomplished by installing a network file server. This doesn't need to be a stand-alone machine with expensive server software. It can simply be a regular workstation configured for double duty. It must be running a secure OS (Windows 2000, XP, or Linux), it must be hardened, and it must be configured with the appropriate accounts and permissions. By eliminating shares on individual machines, the losses can be more easily controlled. A secondary benefit of having a network file server is the provision for an additional location that can be used to backup local data.

Just as separating one's shared files from one's desktop machine helps to stem the spread of a virus or worm, segregating the enterprise resources can go a long way toward providing additional protection for the "family jewels". In IT parlance, this is known as compartmentalization or a layered defense. Using inexpensive NAT (Network Address Translation) firewalls or routers with rule sets, you can isolate various

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

machines by breaking the network into segments based on similarity of function or need, and separating each segment. This adds an extra level of safety to the entire network. In those situations where certain computers need only limited access to the outside world, inter-network data exchange can often be accomplished by using removable optical or magnetic media. Even if you only have two or three machines, guarding the one holding the most business-critical data with an extra layer just makes good sense. Again, a compromise in one section could conceivably be contained without endangering the others. If you feel that you must have wireless network access, the best policy is to compartmentalize it, and only allow authenticated traffic to the interior portion. The more barriers you have, the greater the safety.

This is a probably a good time to revisit the security policy. In it should be a procedure to lock down a compromised machine and totally isolate it from the rest of the segment or network. It should not be allowed back on the network until it has been verified. In terms of productivity and revenue losses, the downtime required to rebuild a key machine is often costly. Another clean computer should be substituted while the original is flushed. This illustrates the need to have spare systems for the key business functions. The spares can also serve an important role in the COOP for offsite operations.

As the business computing environment grows, so does the potential for a compromise. Previously we discussed the problem associated with auto-answer modems. Eventually some businesses find that they just have to allow incoming modem connections. One solution for minimizing internal exposure during those times is to use a network disconnect switch. When the phone line is active, the ethernet cable connection to the computer is opened, effectively isolating that machine. When the call is completed, the switch closes the network connection. The device can even be programmed to force a reboot and a virus scan before the computer is allowed back on. Keep the back door closed! Remote access to the inside network should always be avoided if possible. However, as was mentioned earlier, enterprise growth brings with it an increase in the use of network technology, and along with that use, additional costs. Mobile employees or remote business partners often request access to company records located on internal systems. Several tools with alphabet soup names like VPN, SSL, SSH, PKI can be used to safely provide that access. However, the level of complexity goes up exponentially with their use. Installation of these components is justified only if there is a demonstrated need, and is best left to someone with in-depth security expertise.

A final tool that is extremely useful and informative is the external challenge of your network. It, too, comes with additional cost, but can often be performed when a remote access solution is implemented. The value of attempting to break into your own network cannot be overstated. Wouldn't you rather be one to find the inherent weaknesses in your defense structure? This process should be run on a regular basis to accommodate your changing configurations and the bad guy's stronger assault techniques. Whew! Almost finished!

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

Parting Observations

1. Securing computer systems is a dynamic process. It needs to be continually re-evaluated as enterprise needs change. Keep up-to-date using online resources from authoritative sites such as CERT, SANS, and the FBI
2. It is not "if" you will be attacked, but "when".
3. Cyber security is like a classic odds game. Doing something is so much better than doing nothing. Remember the three phases, and do what you can.
4. There is no final victory when it comes to computer security. The only thing you can hope for is to stay ahead of the bad guy.
5. Don't try to be an expert, and don't try to do it all yourself. Involve or empower interested and motivated employees to help offload some of the responsibility. You might be surprised at the interest they have in new challenges.
6. Having a good security policy not only helps your business, but also helps to protect the United States' computing infrastructure by not becoming an unwilling participant in a major attack.

Good Luck!

References

1. Roberts, Paul. "System Break-in Nets Info on 5.6 Million Credit Cards." Enterprise Security News. Article ID: 1985. 25 February 2003.
   URL: http://enterprisesecurity.symantec.com/content.cfm?articleid=1995&EID=355

2. Fisher, Dennis. "Thwarting the Zombies." eWeek. 31 March 2003.
   URL: http://www.eweek.com/print_article/0,3668,a=39474,00.asp

3. "Mistakes People Make that Lead to Security Breaches." SANS Institute. 23 October 2001.
   URL: http://www.sans.org/resources/mistakes.php

4. Margulius, David L. "Inside Job." InfoWorld. 28 April 2003: 65-67.

5. Wilson, Sam. "Combating the Lazy User: An Examination of Various Password Policies and Guidelines." SANS Institute. 16 September 2002.
   URL: http://www.sans.org/rr/authentic/lazy_user.php

6. "Checklist: Create Strong Passwords." Microsoft Corporation. 2 April 2002.
   URL: http://www.microsoft.com/security/articles/password.asp

7. Costello, Sam. "Server port 80 plagues Internet security." InfoWorld. 3 April 2002.

© SANS Institute 2003, As part of GIAC practical repository. Author retains full rights.

URL: http://staging.infoworld.com/articles/hn/xml/02/04/03/020403hniss.xml

8. Demaria, Michael J. "Defense Starts Here." Network Computing. 20 February 2003: 57-67.

9. Munro, Jay. "Corporate Antivirus Software." PC Magazine. 22 April 2003. URL: http://www.pcmag.com/print_article/0,3048,a=39312,00.asp

10. "Event Archiver." Dorian Software Creations, Inc. URL: http://www.eventarchiver.com/features.htm#overview

© SANS Institute 2003, Author retains full rights.