Core Strengthening Project:

Taking Security to Eleven in Your Microsoft Environment

Chris Coglianese

MSIS 68-595: Information Security Capstone Project

Lewis University

Table of Contents

Abstract

Network security must focus on internal threats as well as external threats. Social engineers exploit the human element to gain access to a network. More attacks are coming from inside a network today than previously. To prevent these attacks, an enterprise must harden its physical security by segmenting the network, addressing user privileges, implementing security patches, separating administrative accounts, monitoring and aggregating event logs, limiting stored credentials, implementing secure network addressing, and restricting remote desktop protocol (RDP) traffic. To secure today's network, employees must also be trained and communication made a priority.

Core Strengthening Project: Taking Security to Eleven in Your Microsoft Environment

Introduction

In today's world where collaboration, information sharing, and social networking are on the rise, the need to secure computer networks can no longer be measured on a scale of one to ten but instead must be ramped up to eleven. In the past, viruses, worms, and other external threats occupied the majority of an information technology (IT) department's security budget and effort. However, more recently the focus is shifting to internal and social engineering threats. In a May 2009 survey of IT executives, "52 percent of respondents said they are more concerned about the possibility of internal data leaks -- both accidental and malicious -- than they are about external threats" (Wilson, 2009). Most companies have invested in external protection for their networks but are only now beginning to realize how vulnerable their networks are to these types of internal threats. The same study reports, "59 percent of respondents said their organizations were either 'likely' or 'bound to' be infected in the next 12 months by malware that is unintentionally introduced by internal employees or business partners" (Wilson, 2009). As more and more vital and confidential information is stored and transferred over computer networks, addressing the increasing threat of internal and social engineering attacks becomes essential.

Companies must develop and implement a plan that will successfully limit the damage done by internal and social engineering attacks – even if removing the possibility of these attacks is not possible. The plan will need to address both physically securing the network and the creation and communication of policies and procedures to employees. This paper outlines the Steps that should be taken to secure a company's network against cyber attacks include segmenting the network, addressing user privileges, implementing security patches, separating administrative accounts, monitoring and aggregating event logs, limiting stored credentials,

implementing secure network addressing, and restricting remote desktop protocol (RDP) traffic. The plan should also establish a workgroup to take ownership of the process and to create policy and documentation. This group would also be responsible for the communication and education of employees on these policies. Creating and implementing a plan that addresses the technical fixes that can limit the ability of these attacks to damage the network and creates well-informed employees will give companies a head start on minimizing the number and strength of internal threats and social engineering attacks.

This paper outlines the steps an enterprise needs to take to increase its networks' security against internal threats. The first step is to analyze the threats to the network and identify the areas of concern about internal network security. Next, an enterprise must take measures that will mitigate the risks to their network or limit the damage done by these attacks. There are areas on which the enterprise must focus. Enterprises need to address social engineering attacks, internal threats, physical breaches, and cyber attacks by putting in place the needed security enhancements for their Microsoft environment that will allow them to move their security to the top of the chart and beyond.

Threats

Social Engineering

Social engineering is a human based attack to gain information an attacker can use to intrude into a network. These strategies are often most successful against uneducated employees. Once the attacker gains information such as username, password, IP, gateway, or personal information, they use it to gain access to your system. "Social engineering is the act of obtaining or attempting to obtain otherwise secure information or access by conning an individual into revealing secure information." (Gregg & Kim, 2005). Regardless of how much security you

have, when someone gains access to an administrative user account, it compromises the network. A comprehensive security plan needs to address this issue by teaching employees how to prevent this from happening. In addition, a network administrator must find ways to limit the impact of successful social engineering attacks.

Social engineering attacks come in a number of different forms. Four common types of attacks are authority abuse, e-mails, phone calls, and reverse social engineering. An example of authority abuse can be an attacker obtaining information by calling or meeting a person while pretending to be their superior or colleague. Often intruders launch e-mail attacks by pretending to be an important service the user has engaged. The email then requests information such as username, password, birth date, etc. in order to confirm the service or to fix a problem with the service. Another type of e-mail attack redirects the user to a fake website. When the user signs on, the attacker records whatever the user types. An example of this type of threat can be seen in Figure 1. In phone call attacks, the attacker may claim something is wrong and ask for a user's login information in order to fix the problem. In reverse social engineering, the attack persuades the user to get help from the attacker. The attacker damages a work application and modifies the error message so that it contains the attacker's contact information. The user may report the error and give out confidential information to fix it. "An alternative form of social engineering is as simple as guessing someone's password. Children's names, birthdays, and phone numbers are likely candidates to be guessed as passwords" (De Laet & Schauwers, 2004). A good security plan must address the different types of common social engineering intrusions to create the most secure network possible.

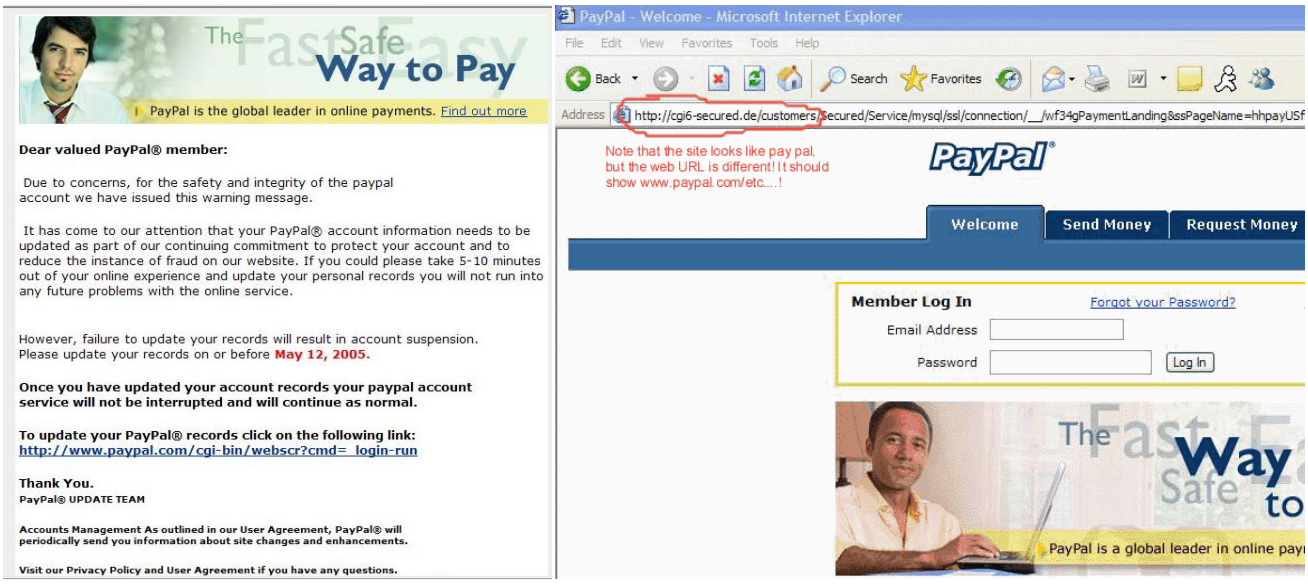

Figure 1. An example of a phishing email & website. Source: "Security & privacy tips", (2005).

Social engineering is very dangerous to network security no matter how many external protections are implemented. All of these protections can be defeated by obtaining a legitimate username and password. Untrained employees may give out vital information without checking to who they are giving it. According to Rich Mogull, a Gartner Senior Analyst, "[t]he single strongest defense for an enterprise against social engineering attacks is an educated employee…Creating a culture of security is the single most critical factor in building a security-aware enterprise and defending yourself from nearly every type of attack" (Chickowski, 2004, p. 5). A way to verify that you are dealing with a legitimate organization is to ask questions. When questioned, most social engineers will abandon their attack. However, this does not mean if the caller answers a few questions they are safe; always confirm callers and email before freely giving out information. Training employees is a critical step in limiting social engineering attacks, but it is also necessary to place physical security measures that can limit the damage done by a social engineer if they are successful in gaining access to the network.

Insider Threat

Much of the focus of network security changes and improvements over the last several years have focused on the need to secure an enterprise's network from outside attacks or social engineers. However, that is not the only threat to networks today. A 2003 survey by the FBI's Computer Security Institute and Ernst and Young showed that nearly 60 percent of all security threats come from internal sources (Gupta, 2003). While the statistics are open to debate on the number of insider vs. outsider threats, enterprises cannot deny that insider threats are a serious issue for network security.

Insider threats can take many different forms – some attacks are malicious and some are unintentional. According to Wilson, 2009, employees self-reported that they breeched security by changing their security setting to view websites, passed on corporate information that may have been sensitive, and even let other people use their work computer devices without supervision. A recent report from Ponemon Institute suggests that negligence may make up to 88 percent of insider security issues. Another study, according to Wilson, highlights the growing concern about malicious attacks. With layoffs a reality in today's economy, nearly 60 percent of employees have stored company data in anticipation of possible redundancy. Malicious attacks may occur by current or former employees. According to Contos & Kleinman (2006), 41% of attacks are by current employees and 59% by former employees. These inside threats may occur for reasons of revenge, greed, or simply for the thrill of being able to outwit the system. Figure 2 shows a graphic representation of an insider threat. Just as there are different motivations, the attacks vary significantly as well. The insider threat may be in the form of a destructive activity such as the release of a virus, the deletion of information, the editing of information, or the creation of data storms that can bring down the network. Another form the insider threat may

take is theft. Insider attacks may result in the loss to the enterprise of important knowledge capital, money, or even the identities of employees.

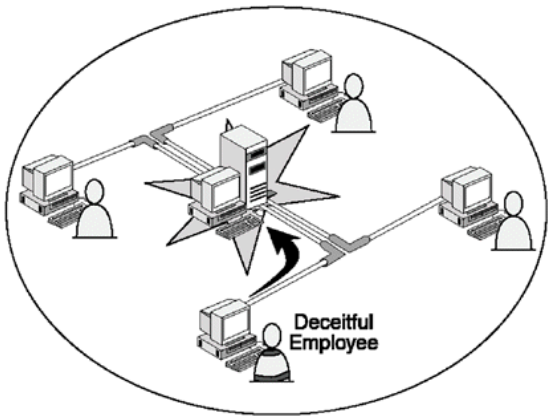

Figure 2. A generic representation of insider security threat. Source: Gupta, (2003).

While most of the fear about insider threats centers on malicious attacks, non-malicious or accidental attacks can often end with the same results. Employees who have higher-level access than they need to perform their jobs can delete data, unintentionally release viruses or other security threats, and can lose important knowledge capital simply because of an error. Insider threats in all their various forms are becoming an increasing concern for enterprises, and enterprises must take new steps to protect against this growing threat.

Physical Breaches

Physical security needs are an essential part of a security plan. Facility location and security measures to ensure limited access to vital components must be in place. In addition, in today's mobile world, the physical security of computers and laptops must be included. An enterprise must establish strong procedures and policies that address key card usage, sign in/sign out of restricted area, computer lockdown, and the disposal of physical assets. Another factor is limiting the physical access to only those who must have access. Develop classification levels for employees that restrict their access to only job necessary equipment. An enterprise needs to

invest in the proper locks, security systems, and other anti-intrusion and anti-theft devices. Monitoring and assessment of these procedures and devices must also be in place. In order to enforce these policies and procedures, there must be some plan in place to observe, verify, and assess the everyday implementation of physical security. Physical security breaches from both internal and external sources can be deterred if the proper devices, policies, and procedures are in place.

Cyber Attacks

To ensure a network's security against insider threats and social engineering attacks a company should implement the following steps.

Step 1: Segment the network. The creation of virtual local area networks (VLANs) can provide network segmentation for security and business process purposes.

Step 2: Address user privileges. Defining the individual user's access to include only the material and server access required to do their job increases the security of the network.

Step 3: Employ security-patching procedures. Revising and employing efficient procedures to implement security patches protects the network from known security vulnerabilities.

Step 4: Separate administrative accounts. Administrative accounts provide greater access to network resources and often to confidential information such as usernames and passwords. Separating these administrative accounts by department or division limits the area compromised by a successful attack.

Step 5: Install and apply event log monitoring and aggregation. Tracking system events and using an aggregator allows network administrators more awareness of unusual activity on

the network. This awareness is necessary to prevent and limit unauthorized access to the network.

Step 6: Limit stored credentials. Credentials such as username and password or network tickets can be stored in individual computer caches. Limiting a computer's ability to save this information helps plug a major security hole in many networks.

Step 7: Implement secure network addressing. Creating a stored directory of known machines on the network prevents unauthorized machines and virtual machines from adding themselves to the network.

Step 8: Restrict remote desktop protocol (RDP) traffic. RDP allows an outside user to gain control of another person's machine. Restricting this type of traffic to those in support positions that require this type of service limits the chances of an unauthorized user from using this service to gain access to network resources.

Segment the Network

Segmenting the network provides an enterprise with a number of security benefits. Segmenting the network is the act of separating the network into smaller sections for better management and maintenance. Creating virtual local area networks (VLANs) limits the access between different groups in the enterprise. For most enterprises, the employees and even the information technology staff, in Accounting have no need to access the Human Resources data or servers. "[W]hen users from mixed departments share a segment, undesirable information captures can occur. If someone from human resources or accounting sends sensitive data such as salaries, stock options, or health records on the shared network, anyone with a network monitoring package can decode the information" (Clark & Hamilton, 1999). The benefit of individual VLANs is to create additional internal firewalls. This is similar to creating real

firewalls in a building, so you can lockdown one segment of a building that is on fire and limit the damage that the fire creates. Another benefit is VLANs place administrative activities closer to the user and allows easier administration of users' access to needed resources without allowing them access to other people's resources. Segmenting the network ensures maintenance of the network's security with limited impact on productivity.

Address User Privileges

Another issue that the enterprise needs to address is user privilege levels. Before internal security became such a hot issue, the information technology (IT) department gave users higher levels of access than they may have needed. This allowed users to install software and updates but also created security vulnerability. By making administrative accounts available to a limited number of people in the enterprise and monitoring their actions while they are using the account is the way to protect administrative accounts. (Northcutt, Zeltser, Winters, Kent, & Ritchey, 2005). Forcing users to operate with the least amount of privilege to accomplish what they need to do is a better alternative. At a minimum, the enterprise must change users' privilege levels from administrators to user levels. By limiting the users to the user privilege level, the enterprise can combat both internal and external threats. If either an inside threat or a social engineer gains access to a computer, this limited privilege will not allow the attacker much leeway in what they can do. With only user privilege, attackers cannot install changes to the root kit, decrypt passwords, or other activities that would jeopardize a larger portion of the network than the individual computer to which the attacker has access.

Employ Security-Patching Procedures

The enterprise must ensure cohesive implementation of security patches across the network. Most attacks take advantage of new or known security holes in operating systems and

applications. According to Microsoft (2006), it is a recommended security procedure to create an automated patch management solution. The solution should detect, assess, acquire, test, and deploy patches in a systematic manner. By creating a definite plan to ensure that all machines receive the latest security patches in a timely and efficient manner, the enterprise greatly reduces the vulnerability of individual machines. By creating a uniform security profile for machines, the network administrators can better service and monitor each individual machine. This process ensure that not only does every machine on the network have the latest security patches but allows more information to be gathered regarding unusual network events. Security patching in a systematic and efficient manner is a major step in enhancing the security of the network.

Separate Administrative Accounts

Separating administrative accounts is another step on the path to a more secure network. One of the ways attackers exploit machines is to create a problem on a desktop. Then, when an IT worker arrives and uses his administrator password, the attacker logs the administrator password and now potentially has access to the entire network. By separating administrative accounts, the network segmentation is guaranteed even for administrative accounts. For example, the administrator who works with desktop applications would have limited privileges –the desktop admin password would only gain them access to their segment of the network protecting the other network segments. This type of administration separation is particularly important for administrators dealing with departments such as human resources or research and development where sensitive information is stored. If an attacker gained an IT technician's password, employee private information and proprietary knowledge capital would not be in jeopardy. "Grouping resources based on similarities in security-related attributes allows us to limit the attacker's area of influence if he gains access to a system inside the perimeter [of the network]"

(Northcutt, Zeltser, Winters, Kent, & Ritchey, 2005). Separating administrative accounts by function helps an enterprise protect its most sensitive data and increases the overall level of security on its network.

Install and Apply Event Log Monitoring and Aggregation

Real time event log monitoring and aggregation is an essential part of a solid network security plan. "Activity logging and monitoring will help assess policy compliance, identify intrusions and breaches, and support an effective response program" (Photopoulus, p. 70, 2008). It gives the enterprise the ability to track and monitor network events as they happen. These logs would note new account creation, group membership changes, failed logon attempts, service starts and stops, etc. and send the information to a centralized location. The aggregator then generates alerts and sends them to the appropriate party. This addresses the problem of event logs being local and the need to be physically at the individual computer to get a complete picture of the network activity. This also means technicians do not need to log into compromised machines with administrator credentials to examine the machine, which could compromise the administrator password. Implementing event log monitoring and aggregation provides an enterprise with an accurate and easily accessible view of all relevant events taking place on its network and increases the security level of the overall network.

Limit Stored Credentials

It is important for network security to limit the storage of credentials on individual machines. Credentials are stored using LAN Manager (LM) Hash, which is an easily unencrypted version of the password saved on the machine for backward capability. "LM hashes are more susceptible to brute force attacks, so your organization might want to disable storing this type of hash" (Lam, LeBlanc, & Smith, 2004). By refusing to allow computers to store

credentials or tickets issued by the domain controller, the enterprise can limit the reach of an attacker with access to a single machine. The attacker cannot use this store information on one computer to access other computers or servers. This is particularly important because of the prevalence of "pass the hash" attacks. Limiting the storing of credentials also prevents users from saving their password – an easy vulnerability to exploit, especially for inside threats. Limiting stored credentials on machines is an easy security vulnerability for enterprises to plug.

Implement Secure Network Addressing

In order to prevent one of the most common types of network security attacks from social engineers or insider threats it is imperative to use secure network addressing. One of the types of attacks currently is to establish a virtual machine that runs completely in the memory of an individual computer on the network. Since the virtual machine has no footprint on the hard drive, it can be difficult to detect. The virtual machine can then request and receive an internet protocol (IP) address for the network and the attacker is inside the network. Enterprises need to limit access to the network address space to computers registered in a directory. This ensures that only known machines can get an internet protocol (IP) address on the network. This prevents people from plugging in rogue machines into the network and receiving an IP address or using virtual machines to gain access. Implementing secure addressing can help an enterprise ensure that only authentic machines are gaining access to its network.

Restrict Remote Desktop Protocol (RDP) Traffic

A technology advancement of today's world is the ability to fix a computer issue without ever touching the computer. Using RDP, an IT technician could take control of a remote machine and make the necessary changes or adjustments to the machine to solve many technical problems. IT technicians can even monitor and maintain servers from home. The disadvantage is

that non-authorized people can use RDP to access not just a single computer but also an enterprise's entire network – without ever being physically present. "Providing remote access must be undertaken very cautiously, because, as soon as you allow employees to connect to your corporate network, you have to some degree, extended your network boundary to their workstations. This means that your network security is only as good as the security of the remote user's system or network. In many cases this borders on no security at all" (Seagren, p. 86, 2007). Restricting remote desktop protocol traffic will keep the wrong people from using RDP to go from machine to machine and attempting to access each server to determine the extent of their credentials. RDP allows enterprises to more efficiently maintain computers but can pose a serious security risk if not restricted to essential authorized personnel.

Mitigations

Defending against insider threats and social engineering attacks is extremely difficult. "Insider attacks can only be prevented through a layered defense strategy consisting of policies, procedures, and technical controls" (Cappelli, Moore, Shimeall, & Trzeciak, 2006). The technical controls necessary to defend against these types of attacks must balance the security of a network with today's need for flexible, accessible network access. Network security cannot chain computer users to a single physical location from which they can complete their work. Instead, networks must accommodate for its employees accessing the network from changing locations and using changing technology. In addition, eliminating insider threats and social engineering attacks is nearly impossible. Creating a layered technical response that limits the opportunities and confines the damage done when an attacker achieves access is essential to a solid network security plan.

Social Engineering

Implementation of several precautions can help limit the impact of social engineering attacks. The first step is to create an updated, easily understandable, and enforceable security policy. The second is to educate employees on both the security policy and on social engineers and their tactics. "The best defense against social engineering tricks is training. Train employees in social engineering tactics and send regular notices of scams" (Whitaker & Newman, 2005). Enterprise must also implement identity management procedures. Implementing these three security measures can decrease the enterprise's vulnerability to social engineering attacks.

The enterprise must have a clear security policy in place. The security policy should be written so it is accessible to all employees. The policy should cover computer system usage, password procedures, data handling, personnel security, physical security, and security education and compliance (Gulati, 2003). The enterprise should require yearly review by employees and compliance sign-off at that time. New employees should also be required to review and sign-off on the policy. The policy needs to be available to employees and kept on the enterprise's intranet with links from frequently used pages. Creating, disseminating, and updating the security policy of the enterprise is the first step in combating social engineering attacks.

Education of employees can be used as a method of combating social engineering attacks. The enterprise should designate a cybersecurity specialist for each department. The role of the specialist should be to assist employees with questions, provide updates and increase awareness of social engineering tactics, and to monitor compliance of the security policies and procedures. In addition, the enterprise should develop both in-person and on-line training that shows real-life examples of social engineering tactics. For employees to understand the full impact of social

engineering attacks, they need to see what they look like. Real-life examples are a good way to help employees understand with what they may be dealing.

In addition, the enterprise needs to implement real testing. For example, they should send out phishing email and provide remedial training to anyone who clicks on the bogus links. The enterprise must develop new education tools and use the cybersecurity specialist to inform employees about changes or new developments in social engineering tactics. Employees should be trained to recognize proper domain structure in URLS. This will allow employees to ensure they are on www.Citibank.com/logon not www.citibank.com.logon.jsp.ru/logon. It is imperative that the enterprise create new education opportunities on a regular basis for the employees as new attacks are developed. Increasing employee awareness and education is essential to combating social engineering attacks.

The enterprise should implement identity management procedures. These procedures should include the automation of password distribution and updating. After automating the process, inform employees that they should never give out passwords to anyone. This eliminates technology personnel from the loop and gives social engineers one less person to imitate when trying to gain password information from employees. Secure password policies are also a part of identity management. Employees should be forced to use more secure passwords and encouraged to avoid using passwords that contain common information about themselves. Microsoft group policy objects can force users to create secure password that include a minimum of eight characters, capitalization and lower case letters, and numbers or special characters. Figure 3 shows the group policy options available to enforce strong passwords. The final part of identity management is to put in place ways for both technical and non-technical employees to ensure that that they are dealing with the people they think they are. Employees should use caller-id,

callbacks, and location identification to ensure they are talking to the help desk before providing sensitive information about their computer to someone. The help desk or technicians should use the same procedures before providing services (Gulati, 2003). Implementing biometric identification for sensitive data is another option an enterprise should consider. Improving identity management can help an enterprise combat social engineering plans.

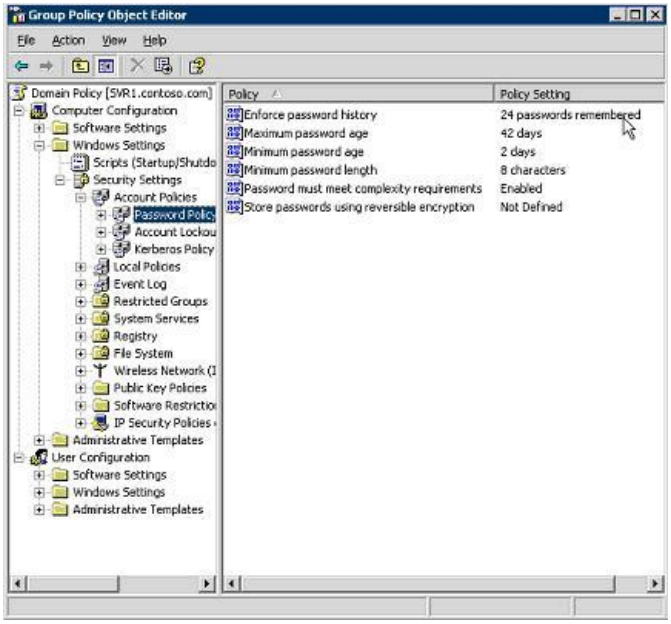

Figure 3. Group policy management console strong password options. Source: Lowe (2003).

Insider Threat

Enterprise must address insider threats to reach a higher security level. Employees must be aware of the security policies and procedures of the enterprise. Enterprises need to create employee awareness training. Activity monitoring and subsequent consequences must be in place. Identity management and access controls need to be implemented. Along with physical and electronic access controls, enterprises must implement employee awareness and monitoring processes to limit insider threats.

Employee awareness training can help prevent insider threats. “You are not educating insiders to change their behavior from bad to good, you are warning them not to change their

behavior from good to bad" (Stiennon, 2009). The enterprise acceptable use policy should be reviewed and employees should be made aware of both the policies and the consequences to them if they violate the policies. A large part of the training needs to focus on data security and confidentiality. Employees need to understand that the information and equipment that they use at work belong to the enterprise and that they need to protect them. Awareness training ought to be an on-going process and not a single training event. The enterprises goal with this training is to keep an employee's awareness at a heightened level. Employing training is one weapon in the war against insider threats.

Activity monitoring is another tool the enterprise can use to help prevent insider threats. "The very best counter to the insider threat is the cyber equivalent of the camera over your shoulder. Just as cameras can prevent shoplifting, visible activity monitoring can prevent data theft" (Stiennon, 2009). Enterprises with Microsoft networks have network activity monitoring software such as Microsoft Network Monitor (netmon) available to them. See Figure 4 for a screenshot of netmon. However, for activity monitoring to work, the software must be on and the contents reviewed. To install Netmon:

1. Open the Control Panel.
2. Click Add or Remove Programs.
3. Click Add/Remove Windows Components to open the Windows Components Wizard.
4. Select Management and Monitoring Tools. Click Details.
5. Check Network Monitor Tools, then click OK.
6. Click Next. If prompted for additional files, insert the installation CD.
7. At the end of the installation, click Finish.

In addition, the IT security group must designate filters and triggers in netmon to ensure that the data retrieved from the activity monitoring is usable. These filters and triggers will be determined based on the specific activities of the enterprise.

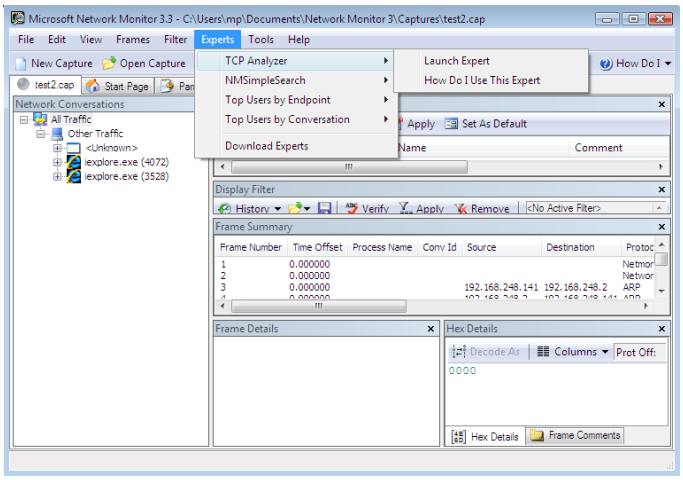

Figure 4. Netmon screenshot

Insider threats need to be addressed as a major consideration in an enterprise's security concern. Physical security and the securing of the network need to be addressed for insider threats as well as external threats. Employee awareness training can help warn employees about usage policy and confidentiality issues as well as consequences for violations. Activity monitoring of the network and employees also provides a way to track violations. Consequences for even minor violations need to occur. Protection from insider threats will help an enterprise crank up their security and could limit damages done by insiders.

Physical Breaches

An enterprise must secure their assets from the threat of a physical breach. All of their servers and network equipment should be stored in secured rooms. Access to these rooms should be limited to only employees who need to access the room to work on the computer and network resources. These rooms should be alarmed and access controlled with keycard or electronic combination locks with unique codes for each individual. This allows the enterprise to track who

enters the room, when they enter, and how long they are in the room. Implementation of this type of system also prevents a lost or copied key from providing physical access to the assets that are most important to the enterprise's network. Unique entry codes and individual card keys also makes it more efficient when dealing with accidental loss of keys or employee termination. The enterprise can simply invalidate these particular codes in the locking instead of requiring the enterprise to replace everyone's access method.

Ramping up the physical security of laptops and desktop computers is essential to increasing an enterprise's overall security. One of the easiest and cheapest ways to secure laptops is to provide employees with laptop cables. Along with providing the cables, an enterprise policy needs to be established on the use the cables. Laptop safes provide a higher level of security for computers that contain highly sensitive materials or for employees who travel frequently. The laptop safe also provides the added security of preventing theft of laptop cards and peripherals. The enterprise should also tag and label both desktop and laptop computers. While this seems like a simplistic solution, labels and tags can act as a deterrent to thieves since it requires additional effort to prepare the computers for resale (Ryder, 2001).

The last and possibly most important step the enterprise must take is to inform the employees that the security of their computer is their responsibility. This can help limit the sharing of enterprise equipment with unauthorized users, the careless leaving out of equipment, and the failure to properly lock and store equipment. Figure 5 shows examples of physical security equipment.

Figure 5. Examples of physical security equipment.

Another part of physical security is making sure that employees understand that should not permit any unauthorized personnel to be unsupervised in the secure areas. This includes repair people, cleaning staff, consultants, or any other personnel. Enterprises need to classify employees by their job function and their need to access equipment and server and network rooms. This will allow the enterprise more easily to limit access. The enterprise should also install alarms and security cameras. The presence of security cameras may act as a deterrent for internal threats as well as external threats.

Cyber Attacks

Segment the Network

Network segregation is the physical separation of computers along divisions that make sense based upon the enterprise. All desktops are placed on VLANs specific to their departments. Each VLAN has a firewall in place to prevent all unauthorized inbound traffic, with exceptions in place for the necessary infrastructure servers to have the access they need. Figure 6 contains an example of a VLAN separation. The plan separates the Servers into three classes. First are servers that require inbound access from the internet such as publicly available web servers. The

color green designates this server type. Second are the servers that are available across the enterprise. Orange designates this server type. Yellow designates servers dedicated to a specific division or group. These groups are definitive boundaries. Orange designated servers cannot be accessed through the enterprise firewall but can be access through the divisional firewalls surrounding the desktops. Yellow servers are accessible only from the specific division groups that need access. Green servers are accessed from the internet through specific ports opened in the enterprise firewall. To maintain the separation, Green servers cannot access desktops or Yellow designated servers. Creating a layered network design provides needed additional security layers to the network.

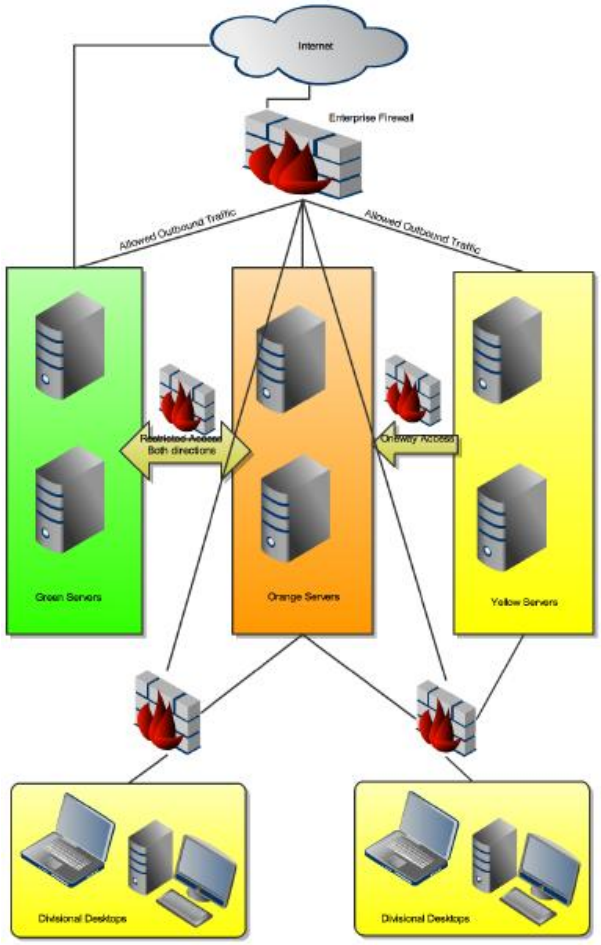

Figure 6. Example of color-coded VLAN architecture.

This separation is to enhance the security by firewalling the enterprise in to compartments that makes it easier to detect and contain compromised machines or intrusion. In an unsegmented network, if an infected machine starts scanning for other hosts to infect it would have access to all the other desktops and servers on the network. By segmenting the networks and firewalling the VLANs the intruder only has access to the machines on that VLAN. When the unauthorized entity starts to scan for machines on other subnets, it hits the firewall. This triggers alerts and notification of suspicious activity. With numerous segments and firewalls, the likelihood of an intruder actually making it out of that segment without detection is low. Administrators are also able to shut down a single segment by blocking all network traffic in and out of that VLAN. This eliminates the intruder's ability to control compromised machines or to compromise more machines. The advantage of this is that an attack will only result in the loss of the segment and not the entire network as it would if an intruder had to be cut off at the enterprise firewall until the compromised machines were found and eradicated. Segmenting the network is essential to limiting damage done by attacks.

Address User Privileges

Removing all standard users from the local administrators group and populating the group with a group of registered admin accounts limits the points of vulnerability. Removal of user admin privileges is a chore best done in stages. A wholesale removal can result in crippling the users. Selecting representative users from across the population, who interact with all the systems and processes, and establishing them as a pilot group will allow discovery of the failures and annoyances of the transition without bringing down the whole enterprise.

The first step is selecting the pilot group. One way to select pilot members is to ask for the managers of each department to select a few from each of their departments. However, managers

may select their least productive people to keep interruptions to a minimum and efficiencies at a maximum. The problem with this is that these people make poor test subjects since their lack of productivity usually means they touch fewer systems and functionality then a productive worker. Having the system administrator of the system choose people they know are power users of the system is often a more effective way to choose a pilot group. The system administrators usually know who uses the full system capabilities and who is likely to cooperate in troubleshooting issues that arise. A third option to transition users from being administrators of the local machine to non-privileged users is to divide up the user population in to groups based on applications and job responsibility. Once you have them grouped into similar roles, the administrator selects a small representative sample of users from each group. A pilot group can be selected using any of the three methods.

After establishing a pilot group, the next step is to build an Active Directory (AD) Security group for computers with no administrative privileges. Administrators then place the computer accounts of the pilot group into the new security group. A Group Policy Object is applied to the AD container that has the computer accounts within its structure. By default, this would be the container labeled “computers” under the domain container in Active Directory. Configure the GPO to control the administrators group by selecting the Group Policy Management Console (GPMC) from your administrative tools menu. Navigate in the left hand domain tree to the Organizational Unit (OU) that is appropriate to have the policy affect the correct computers. Right Click and select create and link policy here. Name the policy, select ID, and choose edit. In the policy edit window, browse to Computer Configuration-> Windows Settings-> Security Settings-> Restricted Groups. Right-click and choose "Add Group". Enter Administrators. See Figure 7 for a screenshot of the Group Policy Management Console. Whatever group entered

here will be the group that is restricted. Select the group and choose the allowed members, which should be Domain Admins, the desktop support security group appropriate to the computers this will be applied to, and the local admin account. The local admin account should be added even though it is disabled and the password randomized so that in the event that access the machine is required to attempt to repair the machine or recover data. Administrators can repeat the process to create a power user group as well. Creating AD groups to limit user privileges will increase the security of the network.

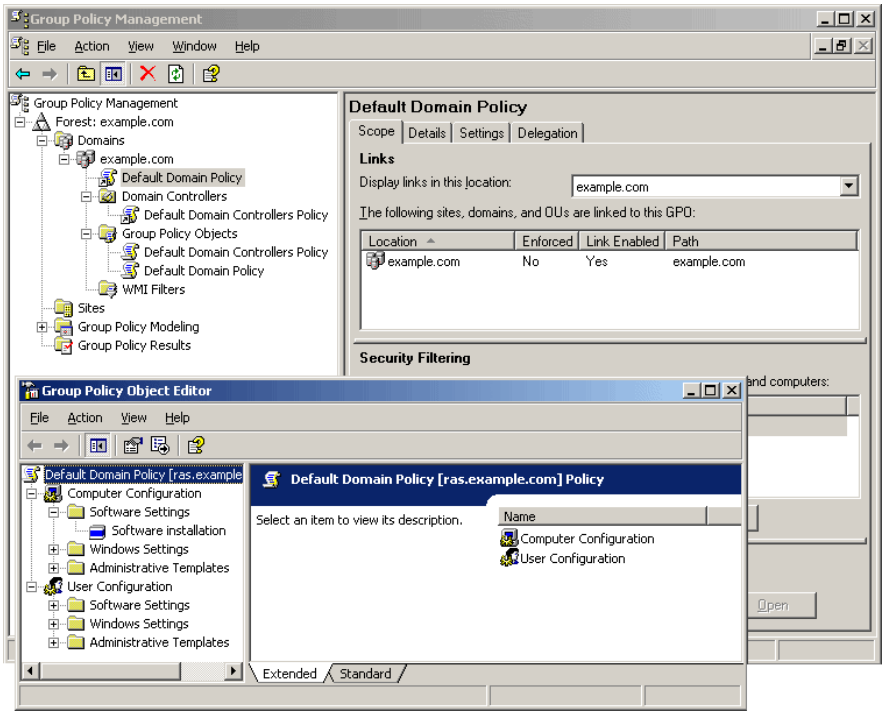

Figure 7. Group policy management console. Source: Lowe (2003).

Employ Security-Patching Procedures

In order to ensure consistent timely security patching, an enterprise should establishing a Windows Server Update Service server and define a group policy to have the desktops and servers obtain their Microsoft and OS patches from this server. The policy also groups the machines, defines automatic update timetables, and installs the updates on a daily basis. Update

Service machines checks for updates on a daily basis and reports to the patching administrator if the Windows Server Update Service (WSUS) server receives new updates. The patching administrator is responsible for evaluating the new patches and releasing the new patches on a regular schedule to a beta group. The beta group reviews, tests it, and troubleshoots compatibility and installation errors if there are problems. After the beta group has evaluated the updates, if there are no unresolved issues, the group releases the updates to the general population. The time from update/patch release until deployment is the most vulnerable time. Exploit code is developed for vulnerabilities by attackers utilizing the information provided in the Advisory or notification about updates/patches. In the cases where details are withheld, attackers use reverse engineering to discover the holes in the program. "We have reached the point where the patch is as revealing as an advisory" (Aitel, 2005).

Waiting a week or months to deploy security updates gives the attackers time to develop and distribute the code to exploit the known vulnerabilities,. Ideally, the updates should be released to the beta group as soon as they are available. After testing, the updates should be deployed to the general population. The testing needs to ensure that all the applications and services are still functioning and available to maintain continuity of day-to-day operations. All applications should be tested along with the connectivity to essential services such as email, intranet websites and file and print resources. The beta group should be representative of the various operating systems and applications that exist across the network environment. Using virtual machines to test updates for compatibility is recommended but it is time consuming and cost prohibitive to develop a separate virtual environment specifically to test updates. Selecting a range of development servers and a representative sample of desktops that are utilized daily gives a much more characteristic beta group. This allows real-world testing of updates for both

deployment and compatibility issues. Selecting power users who perform various job duties across the different divisions or departments and who are willing to cooperate in identifying and notifying administrators of potential problems with updates is key to a successful beta testing program and speeds the deployment of updates to the general population. Testing in a virtual environment requires an administrator to dedicate his time to verify and operate the many different representative virtual machines that have to be maintained to keep current with the actual physical machines. By simply deploying the updates to a select group of representative machines, the users do the actual testing in a real work environment while this process still limits the impact of a bad update on the network.

After deployment of the patches and successful beta group testing for a day or two, it is prudent to access the WSUS console and release the updates to the rest of the machines. It is best to have a standard patching pattern. The recommendation is to release beta patches for deployment on Tuesday as this is the standard release day for Microsoft updates. Release of the test patches to the general population on Thursday. The update group policy object (GPO) sets all of the network machines to install updates at 11:59 p.m. everyday. Therefore, if the administrator downloads the patch from Microsoft on Tuesday and releases patches to the beta group Tuesday evening, the beta group machines install the patches overnight without interrupting the workers. The beta group then tests on Wednesday and Thursday and the administrator authorizes release of the patch to the rest of the machines on Thursday evening. Friday morning all the machines are up to date and the impact on the workers is minimal. Figure 8 shows the workflow of the security patching process.

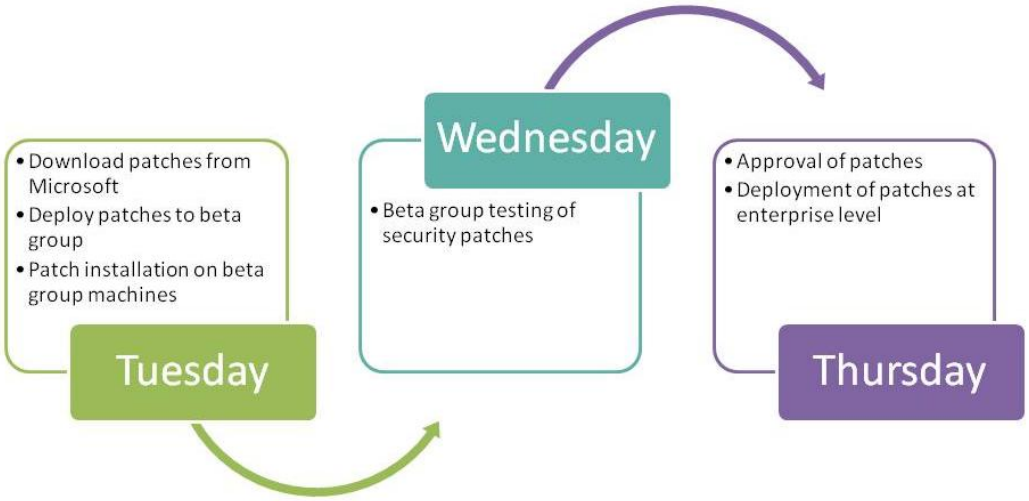

Figure 8. Security patch distribution timeline.

Exceptions to the process, such as a zero day exploit, must be accommodated. These are exceptionally dangerous as the exploit often predates the update to secure the vulnerability. In these cases, depending on the severity and risk, it is often prudent to release the update to all the machines as soon as it is available. This is the reason why the GPO is set to install updates every night and not just on Thursdays. The WSUS admin can release these types of updates as soon as they are available in the WSUS console and they will install overnight. This limits the window of exposure to the existing exploit code. If the GPO reflected the actual patch schedule and a zero day vulnerability update was released on a Friday, it would not be deployed to all of the machine until the following Thursday night update leaving the majority of your computers vulnerable to the exploit for an entire week.

The exact time of patch installation can be staggered by utilizing multiple GPOs targeting specific groups of desktops or servers across your environment to avoid having all the machines in your enterprise reboot simultaneously. A strategy for scheduling needs to be thought out and planned carefully but is easy to accomplish with Microsoft's Group Policy Management console. Considerations are availability of services, work schedules and backup windows. Availability of

services involves maintaining integrity of services provided from multiple servers. For example, a web farm of six redundant servers should be set to patch three servers at 10:30 p.m. every night with the other three set to go at 11:59 p.m. This allows the machines to provide uninterrupted services during the update process. The same goes for your Active Directory Domain Controllers (DC). In order to maintain authentication and connectivity for the domain, the patch time of the domain controllers will need to be staggered to ensure that one DC is available at all times.

A good alternative to WSUS is Microsoft's System Center Configuration Manager (SCCM). SCCM allows updates to be sent out and targeted to specific machines or collections of machines. The updates are downloaded from the Windows update site the same as with WSUS. Machines are scanned and determined to need updates using the WSUS engine. The real difference is in the exact control and detailed reporting achievable with SCCM. SCCM logs and records the installation process and the results, allowing detailed reporting and analysis of the update process and quicker troubleshooting of problems. However, there are several drawbacks to the SCCM update model. First, it is very time consuming. The updates need to be packaged for SCCM deployment and the packages may be operating system specific. This could result in the need to build three or more packages for each security update. Secondly, unlike WSUS, SCCM requires installation of software on all the client machines. SCCM client software is required for the SCCM server to install updates, track installation progress, and to determine installed software. Lastly, SCCM, unlike WSUS, is not free. SCCM requires its own license, a Microsoft SQL server and license, and each client machine requires a client license. These costs add up quickly and make the use of SCCM only cost-effective if you are utilizing SCCM for its numerous other feature and not simply as an update tool.

One of the numerous benefits to utilizing the SCCM Server is the ability to do software base lining. SCCM runs a software inventory on a regular schedule, allowing the administrators to know exactly what software is installed on which server and desktop across the network environment. SCCM allows establishment of numerous control machines that are each configured as all the machine of this type should be. These control machines have all the software and updates that are considered necessary and appropriate. SCCM creates baselines using these control machines. Client machines are assigned to baseline groups and reports of nonconformity are generated based on the last software inventory the SCCM server has. These reports can be generated on a regular schedule allowing the constant monitoring of not only compliance to released updates but also adherence to allowed software installations.

It is not only operating systems that require security updates but also many third-party applications. Many of these applications can be ubiquitous across your enterprise. For example, Adobe Reader is a common PDF viewer installed on almost all workstations. SCCM gives administrators the ability to patch or update third party software through SCCM's software deployment feature. This feature allows administrators to package and deploy both new software programs and updates to existing program. Once packaged, the updates can be deployed to all the SCCM-managed workstations and servers. This extends the capability of SCCM to not only monitor patch state and software compliance but to also deploy security and program updates to the SCCM clients. Another advantage of SCCM is the ability to allow updates to be deployed to collections of machines. Based on the information collected during the hardware or software inventories, SCCM dynamically builds creates collections of client machines. For example, SCCM can build a collection of machines that have installations of Adobe PDF viewer with a version less than 9.0.1. Administrators can then deploy a software package to this collection that

installs the current PDF viewer 9.0.1 during the next update window and sets the program to be mandatory. Once the new viewer is installed, during the next inventory cycle these client machines will move out of the deployment collection, as they no longer have a version of the PDF viewer within the collection parameters. Leaving this collection in place has the added benefit that any SCCM client that tries to revert or install an older version returns to the collection and goes through the update process again. Using SCCM collections enforces conformity with software minimum version requirements and avoids client machines installing or rolling back to vulnerable versions of software.

SCCM leverages WSUS to obtain security patches and determine patching requirements for the clients on the network. SCCM offers advantage in its ability to schedule updates individually and with the concept of maintenance windows that can really pay for the licensing cost by automating the patching. The SCCM maintenance window allows administrators to assign time slices to collections of machines with in the SCCM console. Time slices are specific times and dates when SCCM deploys software or reboots the machines. Use of time slices allows administrators to send an update to a large group of machines without having all the machines deploy and reboot simultaneously. The standard practice for patching is to have predetermined outage time for patching and system maintenance. These times are static and communicated to the customers and administrators whose systems are affected by the possible downtime. Administrators can build collections based on the day of the week and time when there are scheduled maintenance outages. Using time slices helps ensure that patches and updates install during the normal maintenance times and limits non-productive work time for the client machines.

Separate Administrative Accounts

One of the major avenues of exploitation is to compromise an account that has access to a large number of machines in the network environment. If all the network computers share the same local admin password or compromise of a single computer can lead to a loss of all the computers, the network is too vulnerable. Once the attackers have compromised a computer, they can download the Security Account Manger (SAM) database file and work on cracking the passwords by any number of means. The local admin account for a windows machine is easily identified regardless of its name because the Security Identifier (SID) follows a pattern. The Administrator SID always starts with "S-1-5" and ends "-500". Once the attackers have broken the administrator password from the compromised computer, they can then try the same password against the other machines on the network. The best way to prevent this is to eliminate or reduce the number of accounts that can access multiple machines with administrator level access.

The best way to defeat this attack is to segregate the administrative accounts and eliminate local accounts with the same password. To accomplish this, the first objective is to separate the administrative functions into logical divisions based on roles. The first division is into those dealing with workstations, servers, and infrastructure servers. For each of these groups, designate a specific purpose administrative account. Then prefix the account with a standard identifier such as DA for Desktop Admin, SA for Server Admin, etc. followed by initials or first name of the administrator. Then determine which of the administrative personnel need which accounts. Administrators who need to work on desktops will need a DA account. Administrators who work on servers will need an SA account. If administrators work on both servers and desktops, they will require one of each account in order to maintain the separation of administrative functions.

When a desktop is compromised or infected and an administrator logs in, he is not providing credentials that the attackers can use against your servers or your Active Directory infrastructure.

After segmenting the administrative functions by equipment, further subdivide the workstation administrators and server administrators based on the sensitivity level of the data contained on the machines they access. The machines with the highest risk are usually the Human Resources (HR) and Financial Departments. These department machines contain Personally Identifiable Information (PII) both locally and on the network because of handling resumes and expense reports as part of their work routine. Create security groups for each of the enterprise's divisions and add only the appropriate administrator accounts for those managing these desktop or servers to the group. Manage the security groups on the machines through Group Policy's Restricted Groups settings. This setting allows the administrator to enforce the membership of the local groups on the computers. Set the designated security group as a member of the local administrators group along with Domain Admins, System, and Administrator. Because this group has been set as a restricted group through group policy, the membership is enforced. If a new administrator were added to the group locally, they would automatically be removed at the next group policy refresh. These refreshes happen at regular intervals of 4 hours by default but the enterprise can change the refresh rate as appropriate for its individual situation. The refresh of policies also occurs at each reboot and logon.

The enterprise must handle server administrators in a different manner. Figure 9 shows an example of the division needed for administrators. Each server has a security group designated for that server. The local administrator group for each server contains Domain Administrators, System Administrator and ServerName Administrators. The server administrators group for each server is located in an active directory stored in a location with restricted permissions. This

ensures that only a domain administrator can alter membership and prevents unauthorized people from adding new accounts to the server security groups. Maintaining these groups is crucial to maintaining the security of the servers. It is often necessary to add or remove people as server administrators when consultants or application developers are working on building servers applications. Network administrator need to ensure these people do not maintain the ability to add administrators or elevate users on your server.

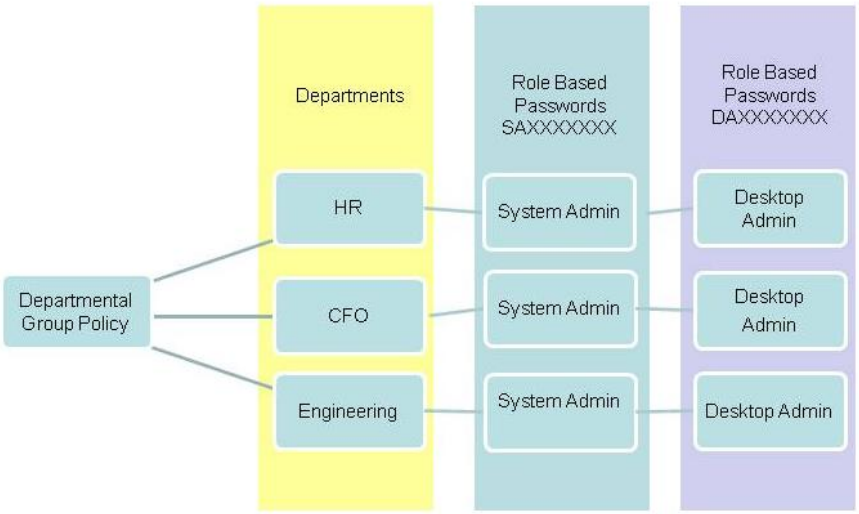

Figure 9. Example of segmenting administrator accounts.

The final step the enterprise must take to ensure segmentation of administrator accounts is to write and enforce a policy. This policy must regulate the usage of administrator passwords and limit the use of passwords to the appropriate situation and equipment. It should also define the roles an employee must have to receive a password. The policy should set the group policy refresh rate and determine the process to issue and remove outside entities' credentials. Creating and maintaining an enterprise-level policy on segmentation of administrator accounts helps ensure the success of segmenting administrator accounts and the network's security.

Install and Apply Event Log Monitoring and Aggregation

In order to maintain and understand the network environment of an enterprise, the enterprise must be aware of what is happening on its network. The best record of what is happening is in the event logs of the computers on the network. When the computers audit and record events to their system logs, they create record of all the information needed. The problem is how to read these individual logs in a timely manner. Log aggregation allows the enterprise to consolidate the logs from all the systems into a central repository. This eliminates the need to go to each machine and view the log files looking for security events, system failures or other irregularities that may indicate a security or hardware problem. The logs are easier to parse for these events if they are held in a central location and a central repository eliminates the need to repeat the search at each individual machine. This system also makes it easier to spot trends in the information. The aggregation of logs shows results for all machines with the errors and this allows a quick view of how many machines are affected by specific events. Implementing software such as Splunk is an excellent way to achieve event log aggregation. Figure 10 shows the Splunk log aggregation. In addition to the ability to function as a syslog server, it has a built-in search engine. The Splunk search engine is similar to the Google search engine. It allows searching of the data in a user-friendly manner and even provides suggestions on what to search for based on the data in the repository. Searches in Splunk can be saved and then scheduled to run at designated times. This allows network administrators to generate reports automatically and to establish monitoring for specific situations. For example, the network administrators can write a report that finds all the failed logon attempts. This report can then be set to alert the appropriate authorities if the number of failed logon attempts for a certain account surpasses five. This

parameter is a good indicator that an attacker is using that account to brute force the password for computers or resources.

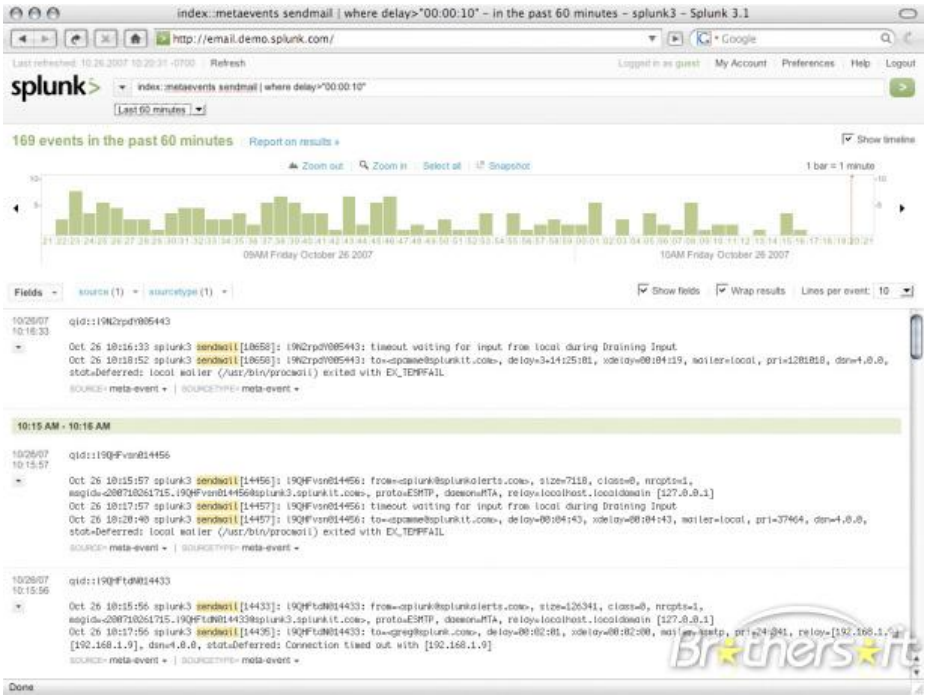

Figure 10. Splunk log aggregation. Source: brothersoft.com, (2009).

In order to get a computer's Windows event log data to the Splunk server, each computer needs an agent installed. Intersect Alliance makes an open source client for syslog servers called the Snare Client for Windows. The Snare client integrates with the event log and auditing processes on Windows machines and allows forwarding of specific or all log entries to a syslog or Splunk server. Figure 11 shows the Snare agent for Windows. The Snare client is configurable to comb through the event logs as they are created and only forwards log entries that meet the specified criteria. This allows network administrators to limit the amount of traffic generated. There may only be specific event IDs from individual computers that are monitored while all the events from mission critical servers are monitored.

Figure 11. Snare agent for windows. Source: static.rbytes.net, (2009).

Another advantage of the Snare client installation is its ability to be scripted and automated which allows deployment through Group Policy to all servers and desktops. This type of deployment occurs in an efficient hands-off manner. The registry key stores the filtering parameters and group policy deploys and maintains them to the necessary equipment. This ensures the administrators can control the Snare client filtering on any level from individual filtering requirements to enterprise-wide standards.

The combination of Snare clients and Splunk server allow monitoring of all types of system log events and the creation of alerts against stored parameters. Aggregating the system log files provides a timely way to review all the important events on the network. Creating alerts allows administrators to be more aware of what is happening on the network. Staying on top of the network activity allows the enterprise to react to threats and problems as they occur and not just deal with cleaning up the mess afterwards.

Limit Stored Credentials

Removal of stored credentials is important in eliminating easily obtained credentials from the machines across the enterprise. These credentials are a liability that will be exploited when one of your machines is exploited or compromised. This is particularly dangerous for portable computers that are often taken off-site and are at greater risk of theft. Theft of a computer with the credentials of not only the user but also the last ten logons to that machine can pose a huge threat to the enterprise. These credentials can be used to bypass the network firewall and gain unlimited access to the internal network resulting possible theft and destruction of data.

Eliminating the stored credentials can all be accomplished through Group Policy Objects. The first step is to eliminate is the LAN Manager Hash or LM hash. Credentials stored on local computers are stored in LM Hash.

> "This type of hash has two substantial weaknesses. The first is that the password length is limited to 14 characters, and these 14 characters are broken up into to independent 7-character chunks. To make matters worse, the password is case-insensitive, reducing the possible key space to 68 characters. At a rate of 10 million hashes per second, this can be tested in less than a week—far faster than most domains require password changes" (Lam, LeBlanc & Smith, 2004).

Administrators need to remove the user's ability to save passwords inside the desktop (checkbox) – or server to avoid this issue. This change will take effect when the user next changes their password. When the password is changed, the hash in the SAM database is replaced with an invalid hash. Implementation of this change can be done via group policy. See Figure 12 for a screenshot of the LM Hash security policy setting. The Group policy is configured by setting two options under Computer Configuration ® Policies ®Windows Settings ®Security Settings ® Local Policies ® Security Options section of the policy. Setting the Network security: Do not store LAN Manager Hash value on next password change to "Enabled"

and the Network security: LAN Manager Authentication level to “Send NTLMv2 response only. Refuse LM”

These restrictions will then force the machines to use NTLM v.2, which is a more secure process to validate authentication between computers and to replace the LM hash at the next password change with a bogus hash. This would be a good time to force the users in the domain to change their passwords, thus wiping out the hashes currently stored in the SAM database. “NTLM hashes do not break up passwords into chunks, are case-sensitive, and can support very long passwords—up to 128 or 256 characters on Windows 2000 and later systems” (Lam, LeBlanc & Smith, 2004). Additionally scripting the registry change to set the “no hash” and “use NTLMv2 “ further ensures that even if the machines are off the enterprise network the changes will not revert to their original setting. Both settings are additions to the key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa. Set the value of the LM Compatibility Level to four and add a key entitled “NoLMHash”. This script can be set as a start-up script for this group policy.

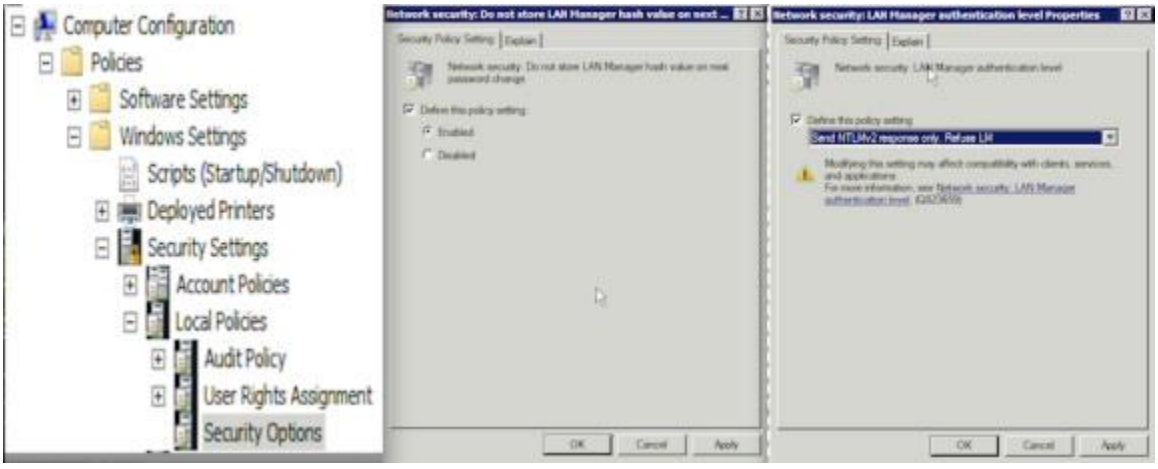

Figure 12. Disabling LM Hash with group policy. Source: Lam, LeBlanc, & Smith, (2004).

The removal of LM Hash does cripple backward compatibility. Legacy systems such as MS-DOS and Windows Workgroup cannot communicate with servers running NTLM v.2. In

addition, Windows 9.x versions will need configuration changes to use NTLM v.2. However, this is actually a security enhancement since machines that cannot run NTLM v.2 are by definition not secure. Allowing non-secure machines to communicate on the network lowers the overall security level of the entire network.

Implement Secure Network Addressing

The first step to securing a network is knowing what should be on the network. Only by knowing what should be on the network can administrators determine what should not be on the network. By establishing a central repository for host registration, the process of identifying network hosts can begin. There are many products available commercially to create a repository. However, an Integrated Host Warehouse (IHW), which is based on several open source programs and servers including MySQL and Django, can also be built. IHW generates the Domain Name System (DNS) zone tables and the Dynamic Host Configuration Protocol (DHCP) reservations. All of the hosts requiring network connectivity must be registered in IHW. Failure to do so will result in the machine not receiving a DNS entry or a DHCP reservation. Entering data on all the hosts on a network is a daunting task for many enterprises. The process can be made easier if the IHW has the ability to upload large CSV files containing the data. Once uploaded the data should then be accessible through a web interface that allows searching and editing of the host data online. Adding new hosts going forward is completed by using a standardized web form, which integrates not only host registration but DNS and network addressing requests. Using an open framework for IHW allows the data stored in it to be leveraged by other applications. A robust IHW can also be used to maintain the integrity of the networks. One method is monitoring the Address Resolution Protocol (ARP) tables on the network switches. Any new entries in the ARP tables are checked against IHW to make sure the

host is registered and on the right network. This process occurs on changes to the ARP cache which limits the networks overhead of having to constantly check on existing machines. Only new or changed ARP entries are scrutinized and validated against IHW. If the MAC address in the ARP table is not registered in IHW, the machine is determined to be a rogue machine. The network port that the rogue host is connected to is shut off and a notification is generated. This prevents rogue machines from attaching to an open port or malware from generating a virtual machine that obtains it own IP address for malicious scanning of your network.

Implementation of this process is easier if the network already has another system in place to scan for new hosts on the network. This is a standard security task that enterprises often already have in place. One such system is Scavenger; an open source software product developed by Argonne National Lab. Figure 13 shows a screenshot of the Scavenger software. Scavenger scans for new hosts on the network and is already polling the ARP tables to find new or changed entries. Scavenger forwards on the new ARP entries to be validated by IHW and if the validation is successful continues with its security scans. If it fails the IHW check or the security scans, a notification is generated and the network port is shut off. If a rogue host tries to spoof an existing MAC address, the system recognizes which system was already on and shuts down the port of the newer entry in the ARP table. Implementing measures to ensure that only known computers are requesting and receiving addresses on the network limits the ability of attackers to gain access to the network through various spoofing techniques.

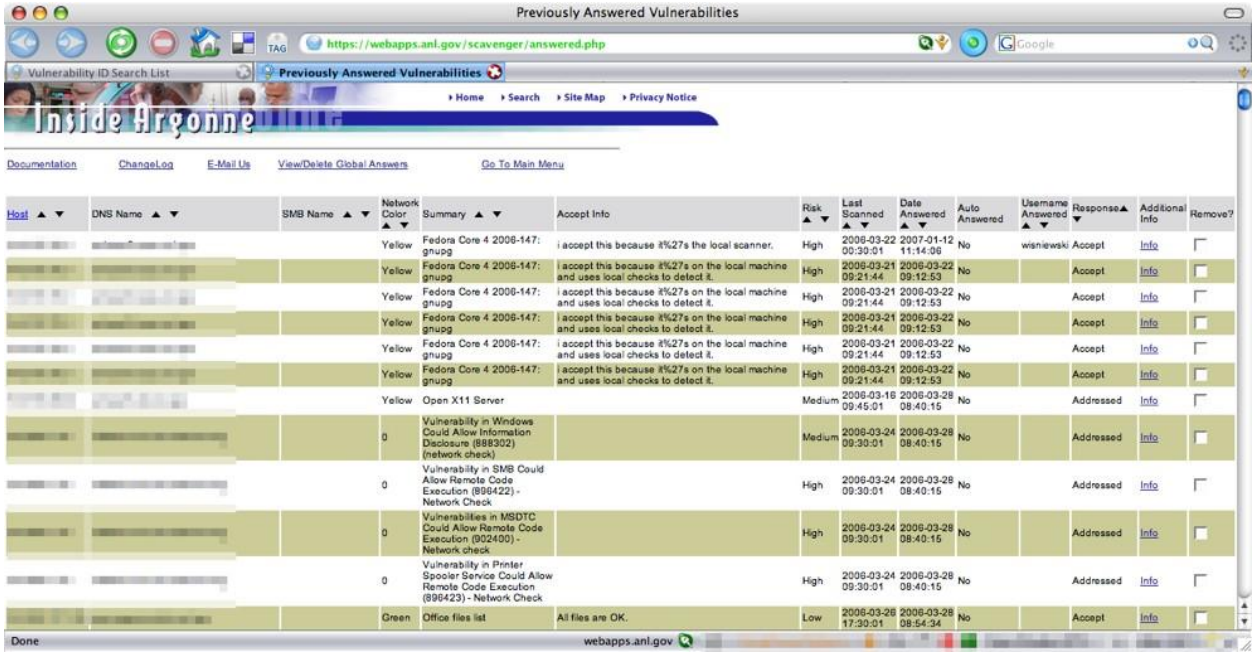

Figure 13. Scavenger software screenshot. Source: Argonne National Lab (2007).

Restrict Remote Desktop Protocol (RDP) Traffic

Remote desktop protocol can present security issues if an enterprise does not enact proper precautions. If a computer is compromised, the intruder can then use those credentials to RDP into the network servers. While this may not get them administrative control of the machine, they can use this as a place from which to stage attacks on other machines or look for vulnerabilities on the server that are only available from an internal desktop. To remedy this situation all RDP traffic into and between the server networks must be restricted. All RDP traffic should be funneled through a bastion host. The bastion host is established as a terminal server that is not in the AD domain. It is on a separate VLAN and is only accessible from designated administrative networks. The administrators can establish a Virtual Private Network (VPN) session to the administrative network in order to work remotely. Each user has a separate user account local to the bastion host with which they log on to the Bastion and initiate their RDP sessions to the servers they need to access. This separation means that even if the administrator's credentials are

compromised at his workstation, they cannot be used to RDP into a server. Likewise, if the credentials on the bastion host are compromised, they are not valid on any of the servers. Figure 14 shows an example of RDP restrictions and bastion host placement. This separation allows you to continue using RDP for ease of administration with a greatly reduced risk to the enterprise while only adding minor overhead for the administrators.

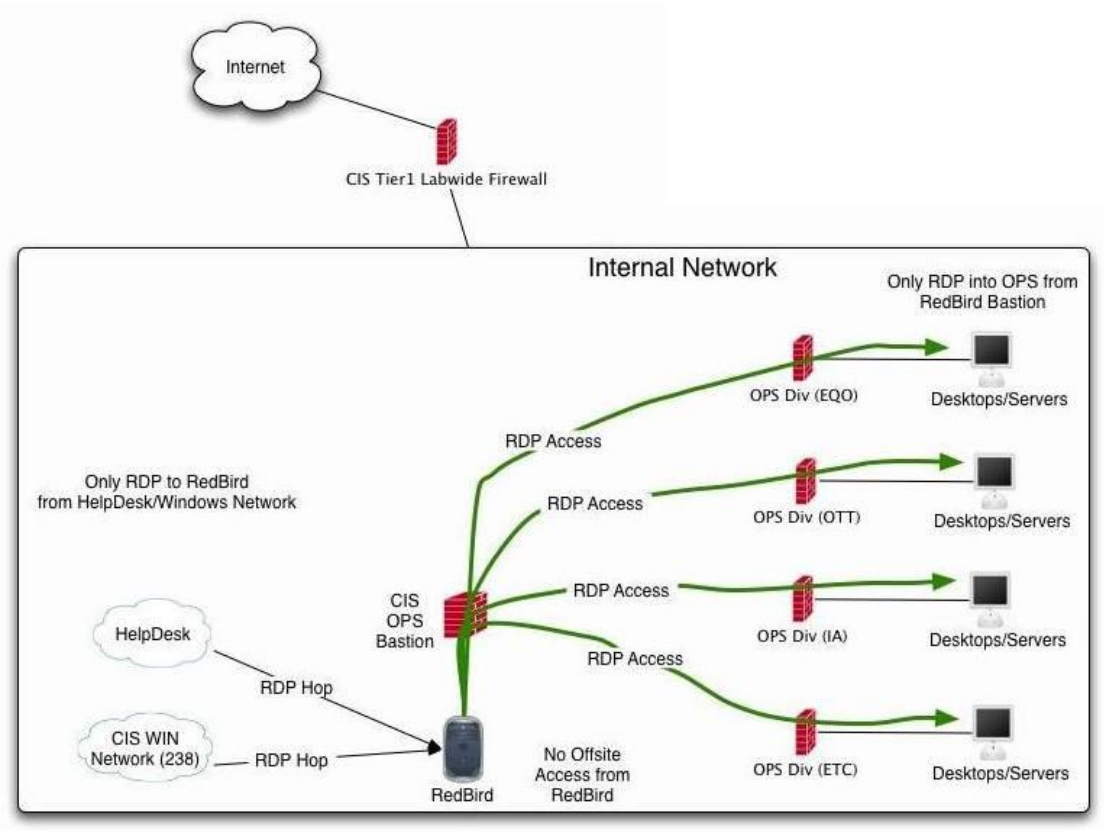

Figure 14. Example of RDP restrictions.

Conclusion

Networks today are more vulnerable than ever as enterprises take advantage of the new technologies available and the increased prevalence of the Internet. While enterprises focused most of their security efforts over the last several decades on keeping outsiders from accessing the network, today's focus has shifted to how to detect and minimize damage done by insider threats and social engineering attacks. The first steps that an enterprise has to make are to harden

the physical security of the network. Enterprises need to segment their networks to ensure that employees do not have access to information that is not necessary for their job. Implementing virtual local area networks is one way to accomplish this task. Another step that needs to be taken is to ensure that users do not have a higher level of privileges than is necessary. As legacy systems developed, users were often given administrative privileges that are not needed in an enterprise that uses managed software deployment. Creating and implementing an effective security patch deployment program is another essential to network security. A network is most vulnerable from the time a security vulnerability is published until the patch is implemented. Administrative accounts should also be separated based on need. Group policies establish departmental allegiances and passwords are distributed based on roles. Event log monitoring and aggregation allows an enterprise to respond more quickly and notice trends that may indicate security breaches. Stored credentials provide attackers with an easy way to attack a network and removing LM Hash storage is an easy fix. Secure network addressing is essential to ensure that rogue machines do not access an enterprise's network using spoofing attacks. Finally, restricting remote desktop traffic using bastion servers can limit the vulnerability that is created by this helpful but dangerous activity. Creating a more physically secure network is important to protecting the enterprise against attack and devastating information loss.

Once the network is more physically secure, an enterprise must address the issue of communication and training of employees. With the increase in social engineering attacks and the growing concern over insider threats, an enterprise must have established policies and documentation to help reinforce its security. Establishing a workgroup is one good way to address these issues. The workgroup would be responsible for the process of creating policies and documentation that provide employees with guidance on security. Establishing solid

guidelines on security provides employees with known parameters in which they can operate. The workgroup would also address the need to communicate to employees. The number one way to combat social engineering attacks is informed employees. The workgroup would be responsible for establishing training for employees and ensuring the training was completed. It would also put together a plan for regular security updates to be provided to employees about phishing or other frauds that might affect the employees. Policies, procedures, and training are an essential part of an enterprise's security program and establishing a workgroup to take ownership of this is essential.

References

Aitel, D. (2005, July 1). Reverse engineering patches making disclosure a moot choice?. *Security Focus*. Retrieved from http:// www.securityfocus.com/news/11235

Cappelli, D., Moore, A., Shimeall, T.J., & Trzeciak, R. (2006). *Common sense guide to prevention and detection of insider threats (version 2.1).* Carnegie Mellon CyLab, at www.cylab.cmu.edu/pdfs/CommonSenseInsiderThreatsV2.1-1-070118-1.pdf.

Chickowski, E. (2004). Trust no one: Low-tech hackers can wreak damage with social engineering. *Processor*, (26)7, 5.

Clark, K., & Hamilton, K. (1999). *Cisco LAN switching* [Electronic version]. Indianapolis, IN: Cisco Press.

Contos, B. & Kleinman, D. (2006). *Enemy at the water cooler: Real-life stories of insider threats and enterprise security management countermeasures*. Rockland, MA: Syngress Publishing, Inc.

De Laet, G., & Schauwers, G. (2004). *Network fundamentals* [Electronic version]. Indianapolis, IN: Cisco Press.

Gregg, M., & Kim, D. (2005). *Inside network security assessment: Guarding your IT investment* [Electronic version]. USA: Sams Publishing.

Gullati, R. (2003). *The threat of social engineering and your defense against it*. Sans Reading Institute at www. sans.org/reading_room/whitepapers/engineering/ the_threat_of_social_engineering_and_your_defense_against_it_1232.

Gupta, M. (2003). *Building a virtual private network* [Electronic version]. Portland, OR: Premier Press.

Lam, K., LeBlanc, D., & Smith, B. (2004). *Assessing network security* [Electronic version]. Redmond, WA: Microsoft Press.

Lowe, S. (2003). Streamline tasks with Group Policy Management Console in Windows Server 2003. *Tech Republic*. Retrieved from http:// articles.techrepublic.com.com/5100-10878_11-5089188.html

Microsoft (2006). *How to implement patch management*. Retrieved from http://msdn.microsoft.com/en-us/library/aa302364.aspx

Northcutt, S., Zeltser, L., Winters, S., Kent, K. & Ritchey, R. (2005). *Inside network perimeter security, second edition* [Electronic version]. USA: Sams Publishing.

Photopoulus, C. (2008). *Managing catastrophic loss of sensitive data.* Rockland, MA: Syngress Publishing, Inc.

Ryder, J. (2001). *Laptop Security, Part One: Preventing Laptop Theft.* Retrieved from http://www.securityfocus.com/infocus/1186

Seagren, E. (2007). *Secure your network for free: Using Nmap, Wireshark, Snort, Nessus, and MRTG.* Rockland, MA: Syngress Publishing, Inc.

Security and privacy tips. (2005). Retrieved from http://www.justtext.com/credit-card-fraud/pay-pal-scam/paypalscam.html

Stiennon, R. (2009, August 30). Countering the insider threat. Message posted to http://threatchaos.com/

Whitaker, A., & Newman, D. (2005). *Penetration testing and network defense* [Electronic version]. Indianapolis, IN: Cisco Press.

Wilson, T. (2009). *Hacks to internal threats: Majority of IT and security execs say insider vulnerabilities worry them most*. Dark Reading Security. Retrieved August 2, 2009, from

http://www.darkreading.com/insiderthreat/security/vulnerabilities/showArticle.jhtml?articl
eID=215801195