



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 permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Hacker Tools, Techniques, and Incident Handling (Security 504)"
at <http://www.giac.org/registration/gcih>

Protecting Against Insider Attacks

GIAC (GCIH) Gold Certification

Author: Brad Ruppert, bradruppert@gmail.com

Advisor: Rick Wanner

Accepted: April 2nd 2009

Abstract

This paper will discuss the key factors in helping to enhance security to protect a company from internal attacks. Most companies focus their resources and defensive strategies on protecting the perimeter from outsider attacks but often the greatest damage can be done by someone already inside these defenses. System administrators can be a company's most trusted ally or their worst nightmare depending on their motivation or personal interest. It is imperative that companies implement internal controls to monitor, detect, and prevent access to sensitive resources to only those individuals that require it to perform their specific job function. The goal of this paper will be to identify high risk areas commonly neglected and to provide some best practice tips to enhance internal security controls.

1. Introduction

Companies large and small all have to deal with expansion and reduction in their employee workforce as the business and economy changes. These modifications to the number of employees will also affect the delegation or consolidation of duties. As duties change, permissions and access to specific assets should be changed as well to fit the current role of that employee. Lack of processes, to ensure that employee access is limited to systems or data that is required to do his or her job, is a major issue that most companies continue to struggle with. This problem is compounded when an employee is promoted from operations to management yet their permissions to systems are not updated to reflect their new role. Failure to remove access to sensitive assets for those employees that no longer have a legitimate business requirement increases an asset's exposure to unauthorized disclosure or alteration. This can be a common enabler of insider attacks which is often overlooked.

It is no secret that companies spend a majority of their security budget on protecting from external attacks but, *“one of the toughest and most insidious problems in information security, and indeed in security in general, is that of protecting against attacks from an insider.”* (Dimitrakos, Martinelli, Ryan, and Schneider, 2007) Typically an insider is an employee of the company that has greater access to sensitive information, a better understanding of internal processes, and knowledge of high-value targets and potential weaknesses in security. *“Consequently, an insider attack has the potential to cause significant, even catastrophic, damage to the targeted IT-infrastructure.”* (Dimitrakos, Martinelli, Ryan, and Schneider, 2007) While this problem is recognized in the security and law-enforcement communities, many companies still tend to rely on audit logs after the insider attack has occurred instead of focusing on developing tools and techniques for analyzing and solving the actual problem.

Many insider attacks are detectable if the proper logging mechanisms have been defined and are appropriately segregated and secured from the production systems. Some insider attacks are even preventable, but this may increase resource or manageability costs. It would be impossible to prevent all insider attacks; therefore a certain level of

trust must exist between a company and its employees. Despite this trust a company should employ some basic security standards such as limiting access to systems based on individual needs and segregating roles amongst team members. Applying these basic principles can go a long way towards protecting against insider attacks.

2. Scope

This paper will focus on identifying some high level areas of potential insider threats and how to employ some basic practices to protect against them. Although multiple areas of risk will be introduced which may include financial, legal, political, or economic, the emphasis will be around information security risks. Technologies mentioned in this paper may not be the best solution for every organization depending on the size, budget, and flavor of systems being supported. The degree of difficulty required to establish controls that protect against insider attacks will depend on the size of the company, number of employees, number of systems, locations of systems, and vendor types. The basic principles of this paper can be applied to any company looking to establish a minimum set of controls to protect assets from insider attacks.

3. Defining Insider Attacks

Understanding what an insider attack is and how it can happen will help to identify causes and how to best prepare defenses against them. What characterizes an insider is that they are usually a trusted employee, student, or contractor that is granted a higher level of trust than an outsider. (Stolfo, Bellovin, and Hershkop, 2008) This trust is usually established through some initial means of authentication followed by authorization to internal assets. Authentication is the process of establishing identity and ensuring the person is who they claim to be. This is usually accomplished through physically meeting the person and associating them with a name and a role (new employee or contractor). The authorization component is the provision of access to specific internal assets based on who they are. One employee may have access to floors 1 and 2 of the building and access to applications A and B. Another employee may have access to floors 3 and 4 and access to applications C and D. Limiting employee access to

areas needed to perform their job is a better means of providing security as opposed to granting all employees the same access to all locations and systems. Granting access based on roles limits exposure and strengthens accountability.

Knowing who is an insider is the first step to classifying internal attacks, and understanding what constitutes an insider attack will be the next step. Some common attacks made by employees, contractors, or students are:

- *“Making an unintentional mistake*
- *Trying to accomplish needed tasks – for example, in a cause in which the system does not support a particular action or the insider is blocked from accessing certain data, the insider may try workarounds to accomplish the same thing*
- *Trying to make the system do something for which it was not designed, as a form of innovation to make the system more useful or usable*
- *Checking the system for weaknesses, vulnerabilities or errors, with the intention of reporting the problems*
- *Acting with the intention of causing harm, for reasons such as fame, greed, capability, divided loyalty or delusion” (Stolfo, Bellovin, and Hershkop, 2008)*

4. First Step to Defending

The first step in protecting a company’s assets from internal attacks is to identify and classify what those assets are and what controls are currently in place to protect those assets. If a company’s most important asset is money, then it will be important to note its physical location, how it is accessed, how it is guarded, who currently protects it, how much of it exists, and how the amount is recorded and maintained safe from alteration. If the most important asset is data, it will be important to note what form is it stored in (electronic or physical), where it is stored (on a server, in a file cabinet), how it is accessed (over the network, physically opening a file cabinet), who has access to it (employees, managers), how changes are logged, and what controls are in place to secure it (usernames & passwords, lock & key). After identifying the assets and all the means of

accessing them, the company should determine who, within the company, has access to these assets. This list should be reviewed and re-evaluated against job roles to ensure that only those employees that actually need access to conduct their daily responsibilities continue to have access. For all other employees, regardless of rank or managerial influence, their access should be removed.

5. Assigning Owners and Custodians

The next step in the protection of internal assets is to assign information owners and information custodians. The owner is typically a senior ranking official that has a solid understanding of the high level business processes but is not involved in the daily routine of operations or maintenance. The owner is responsible for making decisions about the assets including who should have access to them, and for what purpose. The information custodian is responsible for the maintenance and administration of the assets. The custodian follows the directives of the information owner and provides the operational and security aspects of maintaining the asset. If the owner defines the “what and who”, the custodian provides the “how”. Under the guidance of the owner, the custodian must ensure the security (confidentiality, integrity, and availability) of the asset is maintained.

After assigning an owner and custodian, internal accessibility to the asset should be evaluated. Deciding who needs access to the asset and how it can be accessed are the next steps. Regardless of whether this is a new asset or existing one, questions to be answered are:

- Does this person truly need access to this asset to do their job, or can they rely on the output of others to obtain this information?
- What controls are in place to limit access to this asset to only those specifically granted permission?
- How can we detect if an unauthorized person tries to access this asset?

6. Areas of Vulnerability

Accessibility to information is a key deliverable within most companies and is often tied to performance. The ability to get a job done will rely on an employee's access to the correct information and their ability to obtain it in a timely and efficient manner. No one wants to be hindered by obstacles, or having to rely on others to get their own work done. Because of this, management will often focus their efforts on functionality and availability while ignoring security. Without addressing security in the design, many systems will unknowingly possess vulnerabilities, meaning they are *“open to unauthorized access, change, or disclosure of information and susceptible to interference or disruption of system services.”* (Wang, Cheung, and Liu, 2007) This is especially true of internal projects or applications developed solely for the use of employees. While it is important to extend a certain level of trust, there should be some controls in place to prevent an employee from taking advantage of a situation or resource for their own personal gain as opposed to that of the company.

Without interfering with an employee's ability to perform their job, certain business processes should be examined to ensure the appropriate level of security controls are in place to minimize the risk of disclosure, alteration, or destruction of information assets. Areas where large amounts of data can easily be shared or accessed, such as network file shares should be reviewed for appropriate permissions. If a company was concerned about potential data leakage, laptops and portable hard drives should be secured or have restricted use. Identifying potential issues or unauthorized changes requires logging or record keeping of all changes so as to be able to identify who made the change, when it happened, and the details of the changes.

6.1. Network File Shares

One of the most common vulnerabilities of companies is caused by their inherent desire to share everything internally. In today's world of copying multiple people on emails, posting messages on blogs, hosting SharePoint or intranet sites, many companies are focused on getting information out to everyone as quick and easy as possible. When members of a team want to communicate or share files with each other, they will create a folder on an internal file server, give it their team's name, and begin sharing files. Often

Author Name, email@address

these actions are focused on the benefits of sharing the information within the team but little to no thought is given to who also might have access to this information outside the team. Although we like to believe our employees are inherently good, it is not good practice to leave the bank vault completely unlocked. As with network file shares, if the Finance and Accounting team creates a folder that has employee or customer banking information in it, does this really need to be visible to everyone?

6.2. Legacy Permissions

Another issue that many companies face is legacy permissions to internal assets. No company starts out with 20,000 employees and builds their security infrastructure accordingly to this figure. Most companies start out small and grow over time. A small company may have one employee tasked with multiple jobs. As the company grows this employee will begin to delegate his responsibilities to new employees, thereby reducing his access requirements to specific assets. The trouble is, many companies focus their efforts on providing access to their employees and do not focus on removing access or ensuring alignment with actual job responsibilities. If an employee started out as a database developer and was promoted after three years to manager and then three years later to director of operations, it is likely that their access requirements would be significantly different today versus when they started. But there are many directors and vice presidents that still possess their same permissions that they had when they started with the company. This can pose a significant risk to a company if that VP or director becomes disgruntled or didn't get that raise they were expecting.

6.3. Data Portability

The Internet provides a backbone of communication for legitimate business use but also facilitates employees sending internal information outside the company. This can be accomplished by email, file transfer protocol, instant messaging, or even over the web via hypertext transfer protocol (HTTP). Along with relying on networks to send and receive data, employees can also take advantage of local data portability from their desktop or laptop via CD/DVD burners or even USB thumb drives. One study conducted by Promisec found that undocumented or unsecured USB devices was the largest of all internal security threats. (Cook, 2007) While the devices may simplify the transfer of

data between machines, their use also increases the risk of data theft. Employees with access to the company's intellectual property may rationalize the transfer from their work machines to their home systems to work at home. The problem is that once the data leaves a company computer, the company can no longer ensure the security or legitimate use of the data.

6.4. Change Control

Before Sarbanes-Oxley (Public Company Accounting Reform and Investor Protection Act of 2002) the software development life cycle (SDLC) was an unrefined process of taking a project from inception to production deployment in as short of time as possible. This may have contributed to a lack of standards, weak documentation, insufficient testing, scope creep, security holes, budget overruns, and potential project failure. Often times the developers were also responsible for testing as well as production deployment. The problem with the lack of controls over the SDLC is that developers can implement backdoors or changes directly to production systems. While this might make it easier or faster for a developer to make changes, it also increases the likelihood of introducing errors because the changes have not gone through testing. Having development access to production also increases the likelihood of data theft because it enables someone to create a program to access and extract valuable information from the company systems without anyone else detecting it.

6.5. Logging & Monitoring

Changes to a company's data and systems can happen in milliseconds and over a billion times in a single day. *“One of the first complaints heard in most security shops is, ‘there is too much data to look at,’ and finding out what all the different security ‘widgets’ mean can be very confusing.”* (Babbin, Giuseppini, Kleiman, and Carter, 2006) Ensuring that only authorized visibility or authorized changes of the data is taking place, is a constant struggle of any company. Preventing unauthorized access or changes can be an extremely costly solution to implement with regard to both monetary resources and efficiency. Often a company will have to settle for detective controls which would mean logging or recording any access attempts or changes made to sensitive data and by whom. Capturing this type of information and securely storing it to prevent tampering may

sound like a cheaper alternative, but the major cost is associated with reviewing these logs. Many companies will log all types of transactions but fail to implement adequate log aggregation or alerting mechanisms. This defeats the purpose of logging if no one is monitoring or reviewing the changes that are taking place. If an employee were to be siphoning valuable company data for his or her own personal benefit and no one was reviewing the logs to detect this, the logs would, in a sense, be useless.

7. Protecting Yourself

Having introduced some major weaknesses that can lead to insider attacks, we'll now review some important principles and techniques to mitigate these risks. While there is no "Silver Bullet" to prevent against all risks, the key will be to implement some basic controls to help minimize exposure while still enabling the business to perform at its capacity. Some solutions will come at a higher cost than others and it will be up to management to decide if the risk of continuing operations as is outweighs the cost to implement preventive or detective controls.

Two very basic standards of security controls include the "Principle of Least Privilege" and "Segregation of Duties." The Principle of Least Privilege states that "*no entity within a system should be accorded privileges greater than those required to carry out its tasks.*" (Bidgoli, 2006) For example, a manager should be able to authorize an employee's access to a system but does not need the privileges to implement the actual change. For the actual implementation, a system administrator should have the privileges to make the change but only after having received authorization from a manager. (Bidgoli, 2006) Segregation of Duties is "*a fundamental principle of control that no individual should be able to process a transaction from initiation to completion. In electronic funds transfer systems, for example, two or more individuals are involved in the input and execution of a payment.*" (Wilding, 2006) Applying these basic guidelines of security will help enhance the security controls of any system.

7.1. Securing Network File Shares

As a best practice, only system and domain administrators should have full access to read, write, and create new file shares. All other users should only be granted read or

write access to specific folders on a case-by-case basis which coincides with business need. Standard users should never be allowed full access to create a new folder at the root level of a server nor should they be able to modify permissions. This only opens the doors for them to accidentally create a new share and grant “Everyone” read or write permissions. Although there may be some resistance to this if a company is not used to these controls, it will greatly reduce unnecessary internal exposure between departments and reduce the likelihood of data theft.

Establishing a policy whereby only system administrators have “Full Control” to create new network shares is a good practice going forward, but this does not address how to identify or cleanup issues with existing network file shares. In order to begin the identification process you’ll need to scan your existing servers for file shares and their associated permissions. This can be an arduous task depending on the number of servers and users in your environment. One means of tackling this challenge is by using a tool or script that scans all systems in a specific network or range of IP addresses and looks for shared folders with the “Everyone” group having read, write, or full control permissions.

There are multiple tools that can simplify the process of auditing shared folder permissions but some of the best ones include: Enterprise Security Reporter (www.scriptlogic.com), Nessus (www.nessus.org), and Hyena (www.systemtools.com). Enterprise Security Reporter (ESR) is a great commercial tool that performs a comprehensive discovery of all group membership and file security on Windows servers in a domain and stores this information in a SQL database which can be queried by the reporting engine. This product is very robust and built specifically for auditing file and folder security along with group permissions at both the local server and domain level. Individual queries can be run against the database after a discovery scan has been completed to look for all shared folders across all servers with the “Everyone” group having full control permissions. Below is an example of a folder permissions report generated by ESR.

All Permissions For Folder			4/10/2006 10:37:07 AM
Path	Permissions		
A/D	Account		
jon2003svr3.jondemo2.local			Discovery Date: 4/10/2006 10:05:45 AM
C:\SHARE\USERS\accountsmgr\accounts payroll.doc			
Allow	(I) Accounts Manager (accountsmgr@jondemo.local)	Full Control	This folder only
Allow	(I) NT AUTHORITY\SYSTEM	Full Control	This folder only
Allow	(I) \\jon2003svr3\Administrators	Full Control	This folder only
<u>\\jon2003svr3\Administrators</u> Administrator (JONDEMO2\Administrator) Domain Admins (JONDEMO2\Domain Admins) Enterprise Admins (JONDEMO2\Enterprise Admins) <u>Domain Admins (JONDEMO2\Domain Admins)</u> Administrator (JONDEMO2\Administrator) SLAdmin (SLAdmin@jondemo.local) <u>Enterprise Admins (JONDEMO2\Enterprise Admins)</u> Administrator (JONDEMO2\Administrator)			
C:\SHARE\USERS\accountsmgr\secret docs			
Allow	(I) Accounts Manager (accountsmgr@jondemo.local)	Full Control	This folder, subfolders and files
Allow	(I) NT AUTHORITY\SYSTEM	Full Control	This folder, subfolders and files
Allow	(I) CREATOR OWNER	Full Control	This folder, subfolders and files
Allow	(I) \\jon2003svr3\Administrators	Full Control	This folder, subfolders and files
<u>\\jon2003svr3\Administrators</u> Administrator (JONDEMO2\Administrator) Domain Admins (JONDEMO2\Domain Admins) Enterprise Admins (JONDEMO2\Enterprise Admins) <u>Domain Admins (JONDEMO2\Domain Admins)</u> Administrator (JONDEMO2\Administrator) SLAdmin (SLAdmin@jondemo.local) <u>Enterprise Admins (JONDEMO2\Enterprise Admins)</u> Administrator (JONDEMO2\Administrator)			

(www.scriptlogic.com)

The Nessus tool, which is a vulnerability scanner, can also be used to scan for open shares and the permissions associated with them. It works on both Windows and Unix environments and provides a plethora of built-in vulnerability scans above and beyond weaknesses in network folder share security. It is accessible from both the command line as well as a graphical user interface. Below is an example of a Nessus command that will scan an IP range to look for open network shares on the default Windows TCP ports 139 and 445.

```
# ./nessuscmd -U -p139,445 -V -i 10396 192.168.1.0/24
```

The flags used in this command perform the following functions:

nessuscmd Option	Description
-U	Disable safe checks
-p139,445	Limit the scan to TCP ports 139 and 445
-V	Force nessuscmd to print the full plugin output

Author Name, email@address

`-i 10396`Define the plugin ID (In this case plugin id 10396, "[SMB shares access](#)")

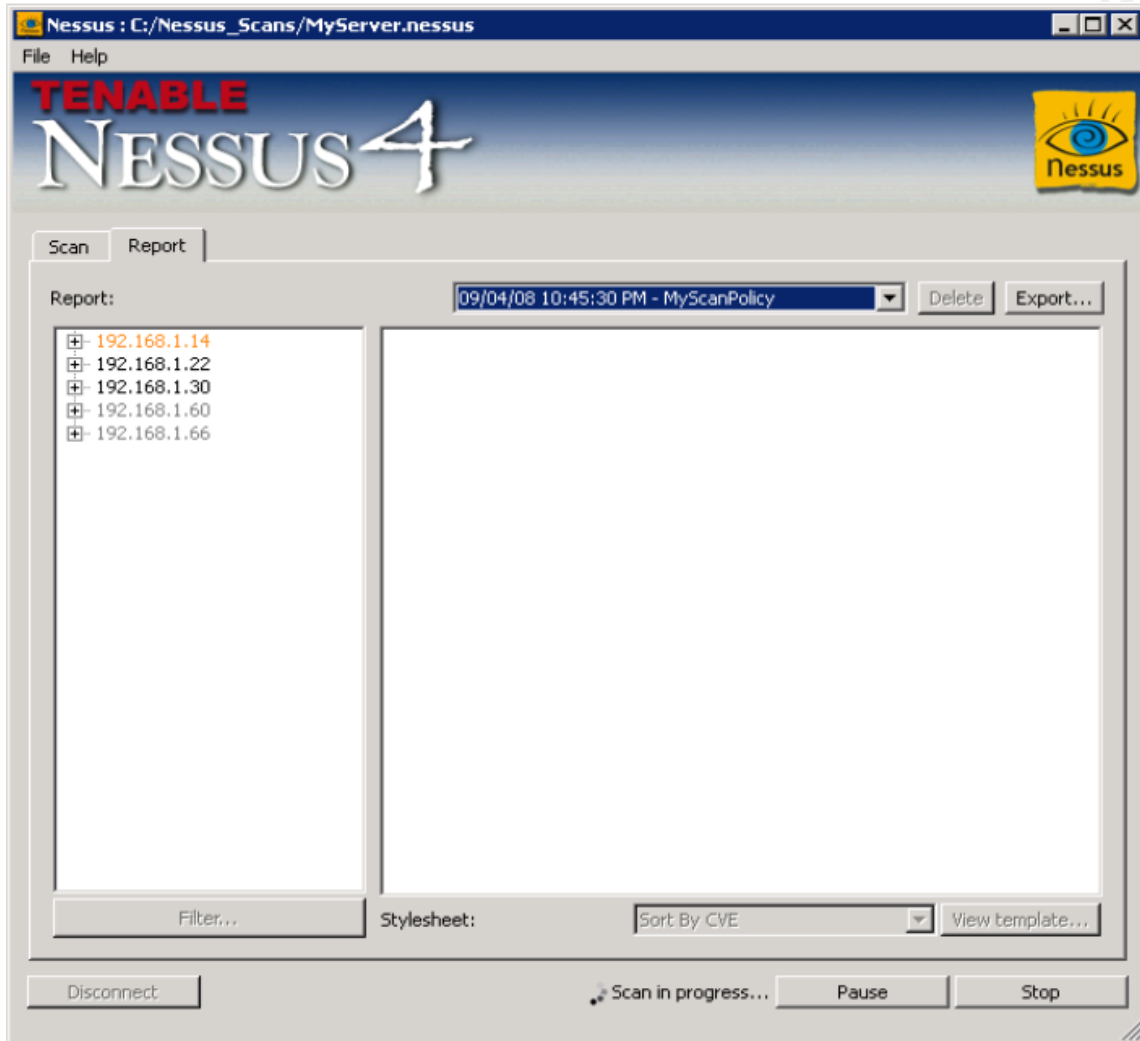
This will result in the following output:

```
+ Results found on 192.168.10.230 :
- Port netbios-ssn (139/tcp) is open
- Port microsoft-ds (445/tcp) is open
[!] Plugin ID 10396
|
| Synopsis :
|
| It is possible to access a network share.
|
| Description :
|
| The remote has one or many Windows shares that can be accessed
| through the network with the given credentials.
| Depending on the share rights, it may allow an attacker to
| read/write confidential data.
|
| Solution :
|
| To restrict access under Windows, open the explorer, do a right
| click on each shares, go to the 'sharing' tab, and click on
| 'permissions'.
|
| Risk factor :
|
| High / CVSS Base Score : 7.5
| (CVSS2#AV:N/AC:L/Au:N/C:P/I:P/A:P)
|
| Plugin output :
|
| The following shares can be accessed as
nessus6804946061421403042121321
| 621 :
|
| - backup - (readable,writable)
| + Content of this share :
| ..
| CreditApplication_Fax.pdf
| Payroll_2009.xls
| Invoice10001.doc
```

(Asadoorian, 2009)

If you prefer to use the graphical user interface, below is an example of Nessus performing a vulnerability scan on several machines listed by their IP address:

Author Name, email@address



(www.tenablesecurity.com)

After having corrected any issues with network file shares, you can implement some auditing rules to detect when users change or create new file shares on a server to ensure they adhere to the new policy. Below is an example of how to enable this auditing feature on a Windows server.

Problem

You want to audit the creation, modification, or deletion of file shares on your system.

Solution

It requires two steps. You must first enable object auditing through Group Policy. Then you must configure auditing for the specific object in the registry.

Author Name, email@address

Using Group Policy

The Group Policy setting shown below enables object audit events:

Configuring object auditing

Path	Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policy
Policy name	Audit object access
Value	Success and Failure

Using a graphical user interface

1. Open Registry Editor (regedit.exe).
2. Browse to the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\lanmanserver\Shares container.
3. Click Edit → Permissions.
4. Click Advanced.
5. Click the Auditing tab.
6. Click Add.
7. In the Enter the object name to select box, type Everyone and click OK.
8. In the Access list, check the Set Value checkboxes in the Successful and Failure columns (two checkboxes).
9. Click OK, OK, OK.

(Danseglio and Allen, 2005)

The examples provided above will help detect problems associated with network file shares but the prevention will ultimately rely upon policy, governance, and awareness. Administrators are the gate keepers of this, so it will be up to them to ensure enforcement. Limiting the ability to make changes to folder share permissions to a specific group of system administrators as opposed to the whole company is a step in the right direction. Establishing routine audit checks of network folder permissions is another means of verifying that policy is being followed.

7.2. Securing Legacy Permissions

All transitions within a company (new hires, promotions, terminations) are initiated by the Human Resources department. Therefore, for a company to ensure

Author Name, email@address

employee permissions on internal assets are aligned with current roles and responsibilities, a review process needs to be created that interfaces HR with IT. All transitions generated by HR must initiate a workflow that tracks employee cost center changes, employee manager changes, location changes, and roles & responsibility changes. Creating a change ticket within a company's internal change tracking system will help facilitate the transition and to log changes as they are made. Building in the steps to review current access permissions on internal systems with the new manager as well as consulting the old manager to remove access to those systems no longer needed by the new role will both be required components of this process. Having a checklist and documenting these access permission reviews is a must to ensure it is taking place in a timely and accurate manner.

Implementing this process requires a workflow system to track changes that occur as a result of an employee job transitions. Some of the important data that should be captured in the workflow system includes:

- the current role of the employee and current manager
- existing access or permissions to systems based on the current role
- the new role of the employee and new manager
- new access or permissions to systems based on new role
- the individual assigned to implement the access changes in each systems
- the results of the new access granted and removal from systems no longer required by the new role

Many companies will already have a partial workflow system that facilitates employee on-boarding but it may not be capable of capturing all required changes to systems or access. There are hundreds of workflow systems available today and finding the one that works best with your company will depend on your budget, resources, and infrastructure. A web-based solution may provide the greatest ease of portability and use so you may consider some of the following commercial examples: Numara FootPrints (www.numarasoftware.com), IssueTrak (www.issuetrak.com), Siebel (www.siebel.com), or even Bugzilla (bugzilla.mozilla.org) which is open source. Each of these were


Author Name, email@address

designed for tracking changes but some are better catered to help desk related changes, some are customer relationship management (CRM) oriented, and others are designed to track changes in product development. In practice, companies tend to acquire or build one type of change tracking application and attempt to make it work for all changes to their enterprise. The right solution(s) depends on the existing environment, infrastructure and budget of those implementing it.

Tracking new changes based on job transitions is important but it will also be necessary to ensure existing access privileges are properly aligned with current job roles. To facilitate this, all systems should be assigned a system owner and a system custodian¹ that will be responsible for validating user roles and permissions on each system. This should be a routine process that takes place annually or quarterly depending on the sensitivity of the system and classification of the data within it. The validation process should capture what was reviewed, who conducted the review, who provided validation of existing permissions, what the outcome or changes were, and when it took place.

If we wanted to conduct a review of permissions in a Windows environment, we may want to examine Windows Active Directory which is typically used to assign permissions to objects like users and computers. One particular area of importance is the “Domain Admins” group. This group has the highest level of permissions in a Windows domain which means the users of this group have full access to all systems and other users within a domain. Because of these permissions, the Domain Admins should be limited to a small set of trusted individuals and the members should be scrutinized regularly. If we wanted to know who these users were, we could use the Active Directory Users & Computers GUI which is accessible under the Control Panel -> Administrative Tools on all Windows servers. An example of this GUI is shown below:

¹ For additional details on system owners and system custodians, see Section 5



(www.wikipedia.org)


Identifying and validating who is currently part of the Domain Admins group is the first step to resolving privilege issues. The second and perhaps more important step is tracking when changes are made to this group. One way to achieve this is to configure the domain controllers to audit or log changes to the “Domain Admins” group. Logging these changes is good, but it is better to have some alerts set to notify when changes like this occur. Another means of accomplishing this is to write a script to query the Domain Admins from Active Directory and check this against a known list of them. Fortunately Windows comes with a basic scripting language (.vbs files) that will allow us to run such a query. Appendix A is an example of a “.vbs” file that will gather the Domain Admins from Active Directory and check these against a predefined group that you can set. The script can be placed on any Windows server in your domain and run as often as you like using the Windows Scheduler. This should be scheduled to run at least daily if not more often to detect when changes to this group occur. When the script runs, it will store the results of the gathered Domain Admins and write them to a text file along with the date

Author Name, email@address

and time. It also has an email function to notify the information security group if the Domain Admins found in Active Directory have either increased or decreased from the defined set. The script can be modified to accommodate your specific domain name and email addresses. See Appendix A for details. (**Note:** *The author assumes no responsibility for the use of the script. Please work with a system administrator before deploying in a production environment*)

7.3. Securing Data Portability

Preventing data leakage is a difficult task for any company considering the various ways data can be transferred to other media or over the Internet. While it will be near impossible to block every vector without interfering with routine business there are some mechanisms a company can employ to reduce this threat. Below is an example of some of the various means data can leak outside a company using a peripheral device.



(www.deviceclock.com)

Fortunately there are tools and applications that companies can use to restrict this access such as DeviceLock (www.deviceclock.com), Sanctuary Device Control (www.lumension.com), and USB Blocker (www.netwrix.com). These are applications managed at the enterprise level that can be used to monitor and prevent the installation or use of USB media or CD/DVD drives on workstations. This can prevent employees from exporting mass amounts of intellectual property onto a thumb drive and taking it home. There are also ways of manipulating Windows Group Policy Objects (GPO) to restrict USB access, but this has been difficult to achieve in practice. If an employee is suspected of removing intellectual property on a thumb drive, this can be validated in the registry by examining the registry using “regedit.exe” and looking up the following key:

Author Name, email@address

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Enum\USBSTOR to find out what devices have been plugged into the system. (Davis, Philipp, and Cowen, 2004)

Almost all companies provide employees with access to the Internet but controls should be put in place around this vector prevent or detect data leakage as well. One means of preventing data leakage over the web is to implement a proxy or filter that limits employee browsing to a small subset of websites required for employees to do their job. Some examples of these products include: Websense (www.websense.com), Barracuda (www.barracudanetworks.com), and Blue Coat (www.bluecoat.com). Each of these commercial tools will provide the protection needed but deciding upon which one to use will depend upon your environment and your budget. Some of these web-filtering products also offer a managed solution if you would prefer not to manage the system internally. One key area of the Internet to consider is the use of external web email such as Hotmail or Yahoo mail. If not required as part of standard business use, web based email sites and other file transfer sites should be blocked.

Content filtering applications can also be deployed to detect outbound traffic which can capture or block any company or customer sensitive data that may be leaving the internal network. An example of this is Vontu (www.vontu.com) which can be configured to detect character strings such as credit card numbers or social security numbers. The only drawback to content filters is that they cannot see encrypted traffic. So if someone inside your company encrypts the data in a password-protected compressed file such as Winzip or Winrar, you will not be able to detect it. Firewalls and Intrusion Prevention Systems can also be configured to block specific file transfer ports, or common applications used to transmit data outside the company's network.


7.4. Securing Change Control

Segregating development, quality assurance, and production environments is the first step to implementing proper change control. Along with separating environments, job roles should be clearly defined and segregated as well. It will be important to distinguish between developers and production support and to ensure that neither have access rights to the other's environment. This will prevent development changes from accidentally or intentionally being deployed to production and thereby reducing errors or the likelihood

Author Name, email@address

of malicious code going unnoticed. Implementing a configuration management system is also an important backbone of change control. This facilitates the workflow and transition between teams and provides a record of the hand-off as well as version control. Best practice is for development teams to “check-in” code they have completed into a source code repository and to have quality assurance teams “check-out” this code to test in their own environment. If the code passes the tests, the production support team can “check-out” the certified code and deploy to production. If the code does not pass the quality assurance testing, it is sent back to the developers to create a new version and the process is repeated.

Change control should be integrated into the System/Software Development Life Cycle used within the company to develop all products or applications. Typically a System Development Life Cycle will include the following steps:



(www.sharpertutorials.com)

Several products can be used to help guide the system development life cycle including Rational (www.rational.com), CA Software Delivery (www.ca.com), or NetDirector (www.emusoftware.com). Each will have their own pluses and minuses but ensuring that the process is followed will depend on policy, governance, awareness, and ultimately the separation of duties and environments between developers, testers, and production support. If you wish to detect or prevent unauthorized modifications to production environments, you can use a product such as Tripwire (www.tripwire.com) to set notifications or implement rollback features.

7.5. Adequate Logging & Monitoring

All successful and failed access attempts of sensitive company resources should be logged and monitored, or reviewed on a regular basis. Along with this, all changes to production data and systems should also be logged to include what change was made, when it took place, and by who. It would be impossible to log and monitor everything that takes place in a company, so it will be important to identify the most critical assets and provide logging and monitoring mechanisms around these assets first. *“New regulations such as Gramm-Leach-Bliley, Sarbanes-Oxley, and California’s SB1386 and AB1950 all require system controls, protections, and the ability to verify or prove the existence of said controls. Under these regulations, the security professional is required to prove that adequate security controls are in place.”* (Maier, 2006) Ensuring that your logging requirements meet or exceed current regulations is a good starting point. The monitoring piece is just as important if not more important than the logging, so it is recommended that the logs be tuned or filtered to only report the most relevant information. Log aggregation solutions should be deployed to correlate logs from multiple systems, ensure synchronization of system times, and to alert an administrator if an anomalous pattern is matched. An example of a Windows tool that will help facilitate this first part is DAD (<http://sourceforge.net/projects/lassie>). This enables the aggregation of logs from thousands of systems and does not require an agent to be installed on the server. Although this is a big step in the process, the pattern matching is the most important piece because the amount of data being logged will make it near impossible for someone to visually detect an anomaly just by reviewing them. Alerts should be set to trigger based on uncommon behavior as well as unauthorized access.

Author Name, email@address

The uncommon behavior will require a baseline to be established for what is considered common behavior to compare against. This is a timely process but a necessity for any company that deals with sensitive data.

Establishing a log aggregation solution that correlates logs from multiple devices and prevents modification or changes to the logs is of utmost importance. If system administrators can manipulate the logs or delete recorded events, then it will be impossible to ensure their accuracy or completeness. System administrators should have access to administer individual business systems while log aggregation administrators should only have access to the log aggregation system. This separation of duties will help to reduce the likelihood of log file manipulation.

System logs will build up very rapidly so it will be important to establish filters or alerts to notify security teams in the event of a critical change or incident. Due to the overwhelming amount of data that can be logged, these filters should be tuned to ignore standard business operations but highlight anomalous activity. This can be a very difficult task that will require some highly trained individuals and will require a serious investment of time. Having a tool may help with some of this “noise reduction” but it will still be a challenging task regardless because every business and infrastructure is slightly different. One such application, ArcSight (www.arcsight.com) is very effective for managing just this type of information known as Security Information and Event Management (SIEM). Other SIEM vendors include RSA(EMC), Cisco, IBM, Symantec, NetIQ, SenSage, and Q1 Labs. These systems can be expensive to purchase and even more expensive to manage but without them there is little to no visibility over what’s happening on the network.

8. Conclusion

The risk of insider attacks are inherent to any business but can be adequately reduced given the proper preparation and forethought. Most insider attacks are made possible because of management’s focus on availability and functionality with little regard to internal security. While all companies must impart trust of their employees to act in a sensible and responsible manner, the company should not turn a blind eye to all

Author Name, email@address

internal actions. Admitting that the risk exists is the step to protecting against it. This is even a greater issue considering the “*likelihood of malicious insider attacks increases with each day of economic bad news. According to a study released by Cyber-Ark Software in December, nearly 60 percent of U.S. workers say they have already downloaded sensitive corporate data in anticipation of a future layoff. Approximately the same percentage of terminated employees do, indeed, take that data with them when they leave, according to another survey published last month by Ponemon Institute.*” (Wilson, 2009) Following some simple security principles like implementing a principle of least privilege and ensuring the segregation of duties will go a long way to providing overall security of company assets. Having change control mechanisms in place and tracking those changes will also ensure only authorized transactions are happening. At the end of the day, greater security comes at the cost of less availability. It will be up to executive management to decide how much risk they are willing to assume to keep business operating as usual.

9. References

- Asadoorian, Paul (2009, April, 1). Tenable Network Security: nessuscmd Tip: Finding Open SMB File Shares. Retrieved June 6, 2009, from Tenable Network Security Web site: <http://blog.tenablesecurity.com/2009/04/nessuscmd-tip-finding-open-smb-file-shares.html>
- Babbin, Jacob, Giuseppini, Gabriele, Kleiman, Dave, & Carter, Everett F. (2006). *Security Log Management*. Rockland, MA: Syngress.
- Bellovin, Steven M., Hershkop, Shlomo, & Stolfo, Salvatore (2008). *Insider Attack and Cyber Security*. New York, NY: Springer.
- Bidgoli, Hossein (2006). *Handbook of Information Security*. Hoboken, NJ: John Wiley and Sons.
- Cook, Rick (2007, May 7). The 10 Most Common Internal Security Threats. Retrieved May 18, 2009, from CSO Web site: http://www.cso.com.au/article/205087/10_most_common_internal_security_threats

Author Name, email@address

- Danseglio, Mike, & Allen, Robbie (2005). *Windows Server 2003 Security Cookbook*. Sebastopol, CA: O'Reilly Media Inc.
- Davis, Chris, Philipp, Aaron, & Cowen, David (2004). *Hacking Exposed Computer Forensics*. New York, NY: McGraw-Hill Professional.
- Dimitrakos, Theo, Martinelli, Fabio, Ryan, Peter, & Schneider, Steve (2007). *Formal Aspects in Security and Trust*. Hamilton, Ontario: Springer.
- Maier, Phillip Q. (2006). *Audit and Trace Log Management*. Boca Raton, FL: CRC Publishing
- Wang, Yuping, Cheung, Yiu-ming, & Liu, Hailin (2007). *Computational Intelligence and Security*. New York, NY: Springer.
- Wilding, Edward (2006). *Information Risk and Security*. Farnham, U.K: Gower Publishing.
- Wilson, Tim (2009, March 9). Reports: Security Pros Shift Attention from External Hacks to Internal Threats. Retrieved May 18, 2009, from Dark Reading Web site: <http://www.darkreading.com/insiderthreat/security/vulnerabilities/showArticle.jhtml?articleID=215801195>

Appendix A

```
##### Domain Admin Script #####
### (Note: The author assumes no responsibility for the use of the script. Please work with a system administrator
### before deploying in a production environment)
```

```
On Error Resume Next
DIM fso, myFile, counter, newAdmins, infoSecText, scriptServerLocation, scriptNetworkLocation
dim standardCount, msgText, infoSecEmail, fromEmailAddress
##### Set the existing domain admin count and location of this script #####
dim adminArray(3)
standardCount = 3
scriptNetworkLocation = "\\serverName\Reports\Domain_Admins"
scriptServerLocation = "C:\Reports\Domain_Admins\domain_admins"
infoSecEmail = "infoSec@myCompany.com"
fromEmailAddress = "admin@myCompany.com"
#####

#### Below are the existing Domain Admins #####
adminArray(0) = "CN=Bill Duke"
adminArray(1) = "CN=Peter Gozenya"
adminArray(2) = "CN=William Frazier"
```

Author Name, email@address

```

##### call the getAdminList Function to gather those in AD and compare to the defined set above #####
getAdminList

Function verifyAdmin(str)
    dim notFound
    notFound = true
    for i = 0 to ubound(adminArray)
        if str = adminArray(i) then
            notFound = false
        end if
    next
    if notFound then
        newAdmins = newAdmins & str & vbCrLf
    end if
End Function

Function getAdminList()
##### change YourDomainName in the next line to your domain #####
Set objGroup = GetObject _
    ("LDAP://CN=Domain Admins,OU=IT-Active Directory Services,DC=YourDomainName,DC=com")
objGroup.GetInfo
arrMemberOf = objGroup.GetEx("member")
Set fso = CreateObject("Scripting.FileSystemObject")
Set myFile = fso.CreateTextFile(scriptServerLocation & dateToday & ".txt", True)

##### Gather the Domain Admins and write to a File along with the Date #####
myFile.WriteLine ("Members of Domain Admin Group: " & "    Date Today: " & now)
counter = 1
For Each strMember in arrMemberOf
    if counter < 10 then
        myFile.WriteLine(counter & ". " & left( strMember,instr(1,strMember,",")-1))
    else
        myFile.WriteLine(counter & ". " & left( strMember,instr(1,strMember,",")-1))
    end if
    counter = counter + 1
    verifyAdmin(left( strMember,instr(1,strMember,",")-1))
Next
myFile.Close

counter = counter - 1
##### If the number of domain admins has changed, then send an email with the details #####
If counter <> standardCount then
    msgText = "Old Admin count was: " & standardCount & " and new count is: " & counter & vbCrLf & vbCrLf & _
        "Please update the script to reflect the new changes. Click here to view latest report" & vbCrLf & _
        scriptNetworkLocation
    sendEmail infoSecEmail ,fromEmailAddress, "Domain Admin Count has Changed", msgText
end if

If len(newAdmins) > 0 then
    msgText = "Default Admin List has changed. The following users were added to original list:" & vbCrLf & _
        newAdmins & vbCrLf & vbCrLf
    infoSecText = "Please update the script to reflect the new changes. Click here to view latest report" & vbCrLf & _
        scriptNetworkLocation

```

Author Name, email@address

```

msgText = msgText & infoSecText
sendEmail infoSecEmail ,fromEmailAddress, "Domain Admin List has Changed", msgText
end if

If Err.Number <> 0 Then
'WScript.Echo "You caused Error " & Cstr(Err.Number) & " " & Err.Description
Err.Clear
'set WshShell = WScript.CreateObject("WScript.Shell")
WScript.Sleep 5000
getAdminList
End If
End Function

'##### Format Today's Date #####
Function dateToday()
    dim today, mon, dy, yr
    today = now()
    mon = Month(today)
    dy = day(today)
    yr = year(today)
    if len(mon) < 2 then
        mon = "0" & mon
    end if
    if len(dy) < 2 then
        dy = "0" & dy
    end if
    dateToday = yr & "-" & mon & "-" & dy
end function

'#### Standard SendEmail Function #####
FUNCTION SendEmail(sTo, sFrom, sSubject, sBody)
    ' send by connecting to port 25 of the SMTP server
    Dim iMsg, iConf, Flds, strHTML, strSmartHost

    Const cdoSendUsingPort = 2
    StrSmartHost = "smarthost.newcentury.com"
    set iMsg = CreateObject("CDO.Message")
    set iConf = CreateObject("CDO.Configuration")
    Set Flds = iConf.Fields

    ' set the CDOSYS configuration fields to use port 25 on the SMTP server
    With Flds
        .Item("http://schemas.microsoft.com/cdo/configuration/sendusing") = cdoSendUsingPort
        .Item("http://schemas.microsoft.com/cdo/configuration/smtpserver") = strSmartHost
        .Item("http://schemas.microsoft.com/cdo/configuration/smtpconnectiontimeout") = 10
    .Update
    End With

    ' apply the settings to the message
    With iMsg
        Set .Configuration = iConf
        .To = sTo
        .From = sFrom
        .Subject = sSubject

```

Author Name, email@address

```
.TextBody = sBody  
.Send  
End With  
  
' cleanup of variables  
Set iMsg = Nothing  
Set iConf = Nothing  
Set Flds = Nothing  
END FUNCTION
```