Common Sense Guide to Prevention and Detection of Insider Threats

2nd Edition – July 2006

Version 2.1

Carnegie Mellon University
CyLab

Authors

Dawn Cappelli
Andrew Moore
Timothy J. Shimeall
Randall Trzeciak

Table of Contents

INTRODUCTION

In 2005, the first version of the *Commonsense Guide to Prevention and Detection of Insider Threats* was published by Carnegie Mellon University's CyLab. The document was based on the insider threat research performed by CERT, primarily the *Insider Threat Study* conducted jointly with the U.S. Secret Service (USSS). Over the past year, CERT has continued analyzing insider threat cases with the USSS. CERT has also conducted additional insider threat research funded by Carnegie Mellon CyLab and the U.S. Department of Defense Personnel Security Research Center. Those projects have involved a new type of analysis of the insider threat problem focused on high-level patterns and trends observed in the cases. Specifically, the projects examine the problem in terms of the interaction of insider psychology, organizational culture, policies, practices, and technology over time.

CERT and the USSS have previously concentrated on analyzing insider threats according to critical infrastructure sector, specifically banking and finance, information technology (IT), and government. In addition, the research team published a report analyzing insider IT sabotage cases across all critical infrastructure sectors. CERT researchers believe it is now important to perform a different type of analysis – by type of malicious insider activity. Therefore, this version of the Commonsense Guide includes a new section that presents a high-level picture of different types of insider threats: fraud, theft of confidential or proprietary information, and sabotage. This section presents patterns and trends observed in each type of malicious activity.

In addition, this report includes one new practice–*Practice 1: Institute periodic enterprise-wide risk assessments*. It has become apparent that one of the overarching problems regarding insider threat is the absence in many organizations of an enterprise-wide, risk-based approach to security management. In addition, many simply do not recognize the risk posed to them by insiders. As a result, they are vulnerable to malicious insider activity; and, once attacked, recovery can be much more difficult.

Most of the practices in this guide reflect new insights from the past year's research at CERT.

Are insiders really a threat?

The threat of attack from insiders is real and substantial. The 2005 E-Crime Watch Survey™ conducted by the United States Secret Service, CERT® Coordination Center (CERT/CC), and CSO Magazine,[1] found that in cases where respondents could identify the perpetrator of an electronic crime, 20% were committed by insiders. The impact from insider attacks can be devastating. One complex case of financial fraud committed by an insider in a financial institution resulted in losses of almost $700 million. Another case involving a logic bomb written by a technical employee working for a defense contractor resulted in $10 million in losses and the layoff of eighty employees.

[1] http://www.cert.org/archive/pdf/ecrimesummary05.pdf

Over the past several years, Carnegie Mellon University has been conducting a variety of research projects on insider threat. One of the conclusions reached is that insider attacks have occurred across all organizational sectors, often causing significant damage to the affected organizations. These acts have ranged from "low-tech" attacks, such as fraud or theft of proprietary information, to technically sophisticated crimes that sabotage the organization's data, systems, or network. Damages are not only financial—widespread public reporting of the event can also severely damage the organization's reputation.

Insiders have a significant advantage over others who might want to harm an organization. Insiders can bypass physical and technical security measures designed to prevent unauthorized access. Mechanisms such as firewalls, intrusion detection systems, and electronic building access systems are implemented primarily to defend against external cyber threats. However, not only are insiders aware of the policies, procedures, and technology used in their organizations, but they are often also aware of their vulnerabilities, such as loosely enforced policies and procedures or exploitable technical flaws in networks or systems.

Partnering with the USSS, CERT has been conducting the *Insider Threat Study*, gathering extensive insider threat data from more than 150 case files of crimes involving most of the nation's critical infrastructure sectors. To date, the researchers have published two reports documenting the results of the study: *Insider Threat Study: Illicit Cyber Activity in the Banking and Finance Sector* and *Insider Threat Study: Computer System Sabotage in Critical Infrastructure Sectors*.[2] This study shows that use of widely accepted best practices for information security could have prevented many of the insider attacks examined. Part of our research of insider threat cases entailed an examination of how each organization could have prevented the attack or at the very least detected it earlier. Rather than requiring new practices or technologies for prevention of insider threats, the research instead identifies existing best practices critical to the mitigation of the risks from malicious insiders.

Based on our research to date, the practices outlined in this report are the most important for mitigating insider threats.

Who should read this report?

This guide is written for a diverse audience. Decision makers across an organization can benefit from reading it. Insider threats are influenced by a combination of technical, behavioral, and organizational issues, and must be addressed by policies, procedures, and technologies. Therefore, it is important that management, human resources, information technology, and security staff understand the overall scope of the problem and communicate it to all employees in the organization.

The guide outlines practices that should be implemented throughout organizations to prevent insider threats. It briefly describes each practice, explains why it should be

[2] See http://www.cert.org/insider_threat/ for all CERT insider threat publications.

implemented, and provides one or more case studies illustrating what could happen if it is not, as well as how the practice could have prevented an attack or facilitated early detection.

Much has been written about the implementation of these practices (a list of references on this topic is provided at the end of this guide). This report provides a synopsis of those practices, and is intended to convince the reader that someone in the organization should be given responsibility for reviewing existing organizational policies, processes, and technical controls and for recommending necessary additions or modifications.

Can insiders be stopped?

Insiders can be stopped, but stopping them is a complex problem. Insider attacks can only be prevented through a layered defense strategy consisting of policies, procedures, and technical controls. Therefore, management must pay close attention to many aspects of its organization, including its business policies and procedures, organizational culture, and technical environment. It must look beyond information technology to the organization's overall business processes and the interplay between those processes and the technologies used.

Patterns and Trends Observed by Type of Malicious Insider Activity

As part of the *Insider Threat Study* conducted jointly by the USSS and CERT, as well as subsequent research in CERT, the insider threat team collected and coded over 150 actual insider threat cases. One hundred sixteen of those cases were analyzed in detail for this report. Because the remaining cases did not fall into the critical infrastructure sectors analyzed in the *Insider Threat Study*, they have not been formally analyzed as yet.

This section of the document presents trends and patterns observed in those cases when analyzed by type of malicious insider activity: insider IT sabotage, fraud, and theft of confidential or proprietary information. Some cases fell into multiple categories. For example, some insiders committed acts of IT sabotage against their employers' systems, then attempted to extort money from them, offering to assist them in recovery efforts only in exchange for a sum of money. A case like that is categorized as both IT sabotage and fraud. Other insiders stole customer data and used it to commit credit card fraud; such cases fall under both theft of confidential information and fraud. Only one case involved all three: sabotage, fraud, and theft of IP. In this case the insider quit his job following an explosive argument with his coworkers. When no severance package was offered, he proceeded to make a copy of the software he had been developing for the company, deleted the software from its systems, and stole the backup tapes. He then offered to restore the software for fifty thousand dollars. Although he was convicted of the crime, the most recent version of the software was never recovered. The breakdown of the cases is shown in Figure 1.

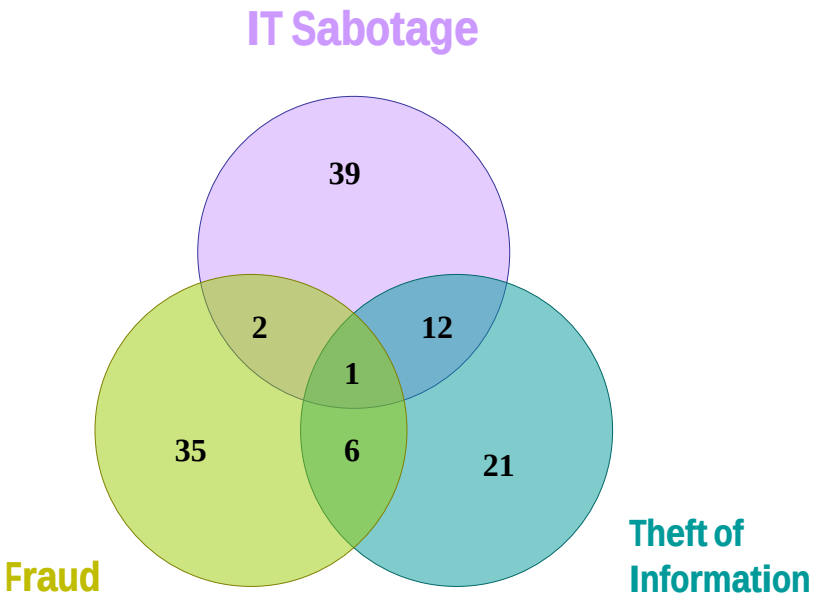

Figure 1: Breakdown of Cases

Insider IT Sabotage

In this report, insider IT sabotage cases are defined as follows: cases in which current or former employees or contractors intentionally exceeded or misused an authorized level of access to networks, systems, or data with the intention of harming a specific individual, the organization, or the organization's data, systems, and/or daily business operations.

CERT researchers analyzed 54 cases of IT sabotage. All occurred in critical infrastructure sectors in the United States. Thirty nine cases involved only sabotage, two also involved fraud, and twelve also involved theft of confidential or proprietary information. One case involved all three: sabotage, fraud, and theft of IP.

Who were the insiders?

The insiders who committed IT sabotage were primarily male and held highly technical positions, the majority hired with system administrator or privileged access. However, according to the U.S. Department of Labor Bureau of Labor Statistics, in 2004, 73% of all employees in computer and mathematical occupations were male.[3] Therefore, while it is useful to note that sabotage was typically committed by technically sophisticated employees, focusing attention only on male employees is probably not a logical conclusion. In addition, the majority of the insiders who committed IT sabotage were former employees.

Why did they do it?

Over half of the insiders were perceived as disgruntled, and most of them acted out of revenge for some negative event. Examples of negative events include termination, disputes with the employer, new supervisors, transfers or demotions, and dissatisfaction with salary increases or bonuses.

How did they attack?

The majority of the insiders who committed IT sabotage did not have authorized access at the time of their attack. Only 31% used their own username and password; 56% of them compromised an account. Thirty three percent used another employee's username and password, and 17% used an unauthorized (backdoor) account they had created previously. They also used shared accounts, including some that had been overlooked in the termination process; 15% used system administrator or database administrator (DBA) accounts and 12% used company accounts.

Almost 36% used sophisticated technical means for carrying out their attacks. Commonly used technical methods included writing a script or program, such as a logic bomb, or creating a backdoor account for later use. Other technical mechanisms included planting a virus on customer computers, using password crackers, and installation of remote system administration tools.

Over 25% took technical preparatory actions prior to the attack, particularly in cases where they anticipated termination. For example, they wrote, tested, and planted logic

[3] http://www.bls.gov/cps/cpsaat9.pdf

bombs, sabotaged backups, and created backdoor accounts. Most logic bombs were designed to delete massive amounts of data; however, at least one was designed to disrupt business operations surreptitiously, six months following the insider's termination. Some backdoor accounts were fairly obvious and could have been detected easily in an account audit, while others were well concealed. Most insiders used remote access, and carried out their attack outside of normal working hours.

How was it detected?

Most of the attacks were detected manually due to system failure or irregularity. Non-security personnel, including customers in almost 25 % of the cases, often detected the attacks. Employees detecting the attacks included supervisors, coworkers, and security staff.

How was the insider identified?

In most cases, system logs were used to identify the insider, including remote access logs, file access logs, database logs, application logs, and email logs. Most of the insiders took steps to conceal their actions; some insiders, knowing that the logs would be used for identification, attempted to conceal their actions by modifying the logs. In some cases, they modified the logs to implicate someone else for their actions.

What were the impacts?

In 74% of the cases, the organization suffered some type of business impact, such as inability to conduct business due to the system or network being down, loss of customer records, or inability to produce products due to damaged or destroyed software or systems.

Other negative consequences resulted from

- negative media attention
- forwarding management email containing private information like strategic plans or plans of impending layoffs to customers, competitors, or employees
- exposure of personal information, like Social Security numbers
- web site defacements in which legitimate information was replaced with invalid or embarrassing content
- publication of confidential customer information on a public web site

In 32% of the cases an individual was harmed. Examples of harm to individuals include threats, modification of evidence to falsely implicate supervisors or coworkers, and exposure of personal or private information.

Fraud

In this report, insider fraud cases are defined as follows: cases in which current or former employees or contractors intentionally exceeded or misused an authorized level of access to networks, systems, or data with the intention of obtaining property or services from the organization unjustly through deception or trickery.

CERT researchers analyzed 44 cases of insider fraud. All occurred in critical infrastructure sectors in the United States. Thirty five cases involved only fraud, two also involved IT sabotage, and six also involved theft of confidential or proprietary information. One case involved all three: sabotage, fraud, and theft of IP.

Who were the insiders?

Only three of the insiders who committed fraud were former employees; all others were current employees when they committed their illicit activity. Half of the insiders were male and half were female. Only 16% held technical positions, four were managers, and the rest held some type of position in which they had legitimate access to modify the data in the systems they used to commit their fraud. For example, many of these insiders held data entry positions or were classified as clerks. Overall, organizations must be vigilant about all employees in attempting to prevent fraud.

Why did they do it?

Most of the insiders who committed fraud did not exhibit financial need and very few held a grudge against their organizations. Some insiders were approached by others outside their organizations and persuaded to modify data in return for payment. For example, some insiders were paid to modify credit histories. In some cases they were paid by people with poor credit histories, and in others by someone (like a car dealer) who would benefit from the beneficiaries' loan approvals. Other insiders were paid by external people to create false drivers licenses, to enter fake health care providers, and to generate false claims totaling significant amounts. Still others were paid to counterfeit federal identity documents.

Some insiders were motivated to provide additional income for their relatives, and a few insiders had large credit card debts or drug-related financial difficulties.

How did they attack?

Only two of the insiders did not have authorized access when they committed their fraud; all others were legitimate users. Two had system administrator access, over 50% had privileged access, and the rest, about 40%, had authorized, unprivileged access. Almost all of the insiders used only legitimate user commands to commit fraud. Only 16% of the fraud involved sophisticated technical techniques, like use of a script or program, creation of a backdoor account, or account compromise.

More than 75% of the insiders used their own usernames and passwords to commit their fraud. Almost 20% compromised someone else's account, two insiders used a company computer account, and only one insider created a backdoor account. One insider used

social engineering, and one used a computer left logged in and unattended by a coworker.[4]

Only two insiders took technical preparatory actions to set up their illicit activity. One insider enabled fraudulent medical care providers to be added to the database, and another disabled automatic notification of the security staff when a certain highly restricted function was used in the system.

Most of the fraud was committed during working hours, with only four insiders acting before normal working hours and three after working hours. Only 14% used remote access, with most insiders acting from within their employers' workplaces.

How was it detected?

Only two of the insiders were detected due to system failure. Half were detected due to system irregularity, and the majority of the cases were detected by non-technical means, such as notification of a problem by a customer, law enforcement officer, coworker, informant, an auditor, or other external person who became suspicious. Most of the fraud was eventually detected by multiple people. About 25% of the cases were detected by non-IT security personnel, 25% by other employees, 20% by customers, and 18% by the insider's supervisor. Only three cases were detected by the people responsible for the information or system and two cases by system administrators.

How was the insider identified?

In most cases system logs were used to identify the insider, including database logs in almost half of the cases, system file change logs, file access logs, and others.

What were the impacts?

The fraud cases analyzed for this report affected not only the insiders' organizations, but also other innocent victims. For example, a check fraud scheme resulted in innocent people receiving collection letters due to fraudulent checks written against their account. Other cases involved insiders committing credit card fraud by abusing their access to confidential customer data. Other insiders subverted the justice system by modifying court records. Some cases could have very serious consequences – cases in which insiders created false official identification documents or drivers licenses for illegal aliens or others who could not obtain them legally. Similarly, one insider accepted payment to modify a database to overturn decisions denying asylum to illegal aliens.

The insiders' organizations also suffered as a result of the fraud. Impacts included negative media attention as well as financial losses. One insider committed fraud against a state insurance fund for a total of almost $850,000, and another insider working for the same company was tied to almost $20 million in fraudulent or suspicious transactions. Another insider committed fraud against a federal agency for over $600,000. In a case involving both sabotage and fraud, an insider set himself up to benefit from the abrupt decline in his company's stock price when he deleted over 10 billion files on the company's servers, costing the organization close to $3 million in recovery costs.

[4] Some insiders used multiple accounts to commit fraud.

Theft of Confidential or Proprietary Information

In this report, cases involving theft of confidential or proprietary information are defined as follows: cases in which current or former employees or contractors intentionally exceeded or misused an authorized level of access to networks, systems, or data with the intention of stealing confidential or proprietary information from the organization.

CERT researchers analyzed forty cases of theft of confidential or proprietary information. All occurred in critical infrastructure sectors in the United States. Twenty one cases involved only information theft, twelve also involved IT sabotage, and six also involved fraud. One case involved all three: sabotage, fraud, and theft of IP.

Who were the insiders?

Eighty percent of the insiders who stole confidential or proprietary information were male and over half held technical positions. Twenty five percent were former employees; the other 75% were current employees when they committed their illicit activity. Interestingly, 45% of the insiders who were current employees at the time of their theft had already accepted positions with another company.

Why did they do it?

Some insiders were financially motivated, for example, stealing information to commit credit card fraud or selling information to their company's competitors. Others were about to start new jobs or form their own companies, and felt entitled to the information. Still others were disgruntled and chose to embarrass their employers by revealing private, sensitive, or confidential information. Some insiders did not realize what they were doing was wrong: for instance, stealing information to help a friend without recognizing the potential consequences, or taking information to a new job only for personal use.

How did they attack?

More than 75% of the insiders had authorized access when they committed their theft. Only one had system administrator access. One former employee was given authorized access to do some additional work; he used that access to commit his theft. The rest of the authorized users were fairly evenly split between privileged and unprivileged users.

More than 75% of the insiders used their own usernames and passwords to commit thefts. Thirty two percent used someone else's account, 14% used a shared account, and two insiders used a company computer account. No insiders created a backdoor account, although 45% of the insiders did compromise an account for use in their illicit activities.[5]

Less than 25% of the insiders took technical preparatory actions to set up their activities. One insider surreptitiously installed a modem for future access. Another insider requested installation of special software that he then used to copy a large amount of data onto multiple floppy disks. Another insider, lead developer for a production software system, deliberately created no backups of the source code and wrote no documentation, to amplify the impact after he deleted all of the source code from his organization's system.

[5] Some insiders used multiple accounts to commit theft of confidential or proprietary information.

The majority of the theft was committed during working hours, although three insiders acted before normal working hours, one on a weekend or holiday, and over 25% after working hours. Almost 75% of the insiders stole the information from within the employers' workplaces. Thirty percent used remote access, accessing their employers' networks from their homes or from another organization.[6]

How was it detected?

Only one theft was detected due to system failure. Half were detected due to system irregularity, and many were detected by non-technical means, such as notification by a customer or informant, detection by law enforcement investigating the consequences of the theft, rumors, audits, suspicions by coworkers, or detection by security staff reviewing log files. Most of the theft was eventually detected by multiple people. Twenty one percent of the cases were detected by system administrators or IT security personnel. Thirty four percent were detected by other employees, security personnel, or the person responsible for the information or system. Twenty one percent were detected by the insider's supervisor. Only four cases were detected by customers and one was discovered when reported by the company's competitor.

How was the insider identified?

In most cases, system logs were used to identify the insider, including file access logs in 37% of the cases and database logs in more than 25% of the cases. Email logs were used in 18% of the cases, and remote access logs in 18%.

What were the impacts?

Impacts on organizations as a result of their employees' theft of confidential or proprietary information included financial and other losses. Confidential information was publicly revealed on web sites, and at least one business was shut down as a result of its employee's acts. Some companies' trade secrets were stolen; in some instances the information was returned by ethical competitors, but in others it was compromised, resulting in legal action or loss of competitive advantage.

In some cases, there were other victims in addition to the organization (for instance, in five cases that involved credit card fraud). Some cases had extreme consequences. One insider used her system access in a sheriff's department to look up the address of her son's ex-wife in a restricted system. Her action inadvertently led to the death of the woman's fiancé when her son went to the house and stabbed him. In another case, an insider sabotaged his company's servers, then stole up to seventy-seven backup tapes, including those in off-site storage. As a result, the emergency services 911 address lookup system did not function for a large geographic area until the system could be restored.

Finally, the insiders themselves sometimes suffered unanticipated consequences. Some insiders were surprised that their actions were criminal in nature. In one case, the insider committed suicide before he could be brought to trial.

[6] Some insiders acted remotely and within the workplace, and some acted at various times of day.

Summary – Comparison of Insider Threat Types

Forty seven percent of the 116 cases analyzed for this report involved IT sabotage, 38% involved theft of confidential or proprietary information, and 34% involved fraud. Therefore, while IT sabotage was most prevalent, all three types of malicious activity were experienced in fairly comparable numbers. Therefore, organizations should consider whether each of these activities is a potential threat to them, and if so, consider the information regarding those types of crimes carefully.

The authors of this paper contend that IT sabotage is a threat to any organization that relies on an IT infrastructure for its business, regardless of the size or complexity of the configuration. Likewise, it is unlikely that many organizations can disregard theft of proprietary or confidential information as an insider threat. Therefore, all organizations should heed the practices detailed in the remainder of this report for prevention of sabotage and information theft. While fraud is not a risk for all organizations, none of the practices in this report apply only to fraud.

Who is the insider threat?

The dangers posed by disgruntled technical staff, both before and after termination or other negative work related events, need to be recognized as potential threats for insider IT sabotage. Data pertaining to fraud and information theft, on the other hand, suggests that organizations need to exercise some degree of caution with *all* employees. Current employees in practically any position have used legitimate system access to commit those types of crimes. Of special note, however, is the fact that almost half of the employees who stole information while still employed had already accepted other job offers. Therefore, extra caution should be exercised once the organization becomes aware of this type of information, either formally or via rumor.

A balance of trust and caution should factor into the organization's policies, practices, and technology.

How Can They be Stopped?

The methods of carrying out the malicious insider activity varied by type of crime. The IT sabotage cases tended to be more technically sophisticated, while the information theft and fraud cases tended to be less.

Once again, it is important that organizations carefully consider implementing the practices outlined in the remainder of this report to protect themselves from any of these malicious activities that pose a risk to them. Proactive measures need to be instituted and maintained at a constant level in order to prevent or detect technical preparatory actions by technical staff. Other measures are effective in preventing or detecting fraud and theft of information, and likewise, should be implemented consistently.

Too often organizations allow the quality of their practices to erode over time as they appear to be unnecessary since no malicious activity is detected. One of the vulnerabilities posed by insiders is their knowledge of exactly this: the quality of their organization's defenses.

What if an Insider Attack Succeeds?

One pattern in all of the cases is the importance of system logs in identifying the insider. Regardless of type of crime, system logs provide the evidence needed to take appropriate action. Since many technical insiders attempted to conceal their actions, sometimes by altering system logs, it is particularly important that organizations architect their systems to ensure the integrity of their logs.

The remainder of this document is structured around thirteen practices that could have been effective in preventing the insider threats analyzed for this report, or at the very least, would have enabled early detection of the malicious activity.

Best Practices for the Prevention and Detection of Insider Threats

Summary of practices for preventing insider attacks

Implementation of the following 13 practices for preventing insider attacks will provide an organization defensive measures that could prevent or facilitate early detection of many of the insider attacks other organizations have experienced.

PRACTICE 1: *Institute periodic enterprise-wide risk assessments.*
It is difficult for organizations to determine the proper balance between trusting their employees, providing them access to achieve the organization's mission, and protecting itself from those same employees. Insiders' access combined with knowledge of the organization's vulnerabilities in both technology and business processes gives them the ability to carry out malicious activity against their employer if properly motivated. An organization must protect itself from both insiders and outsiders using risk management principles. The organization must take an enterprise-wide view of information security, first determining its critical assets, then defining a risk management strategy for protecting those assets from both insiders and outsiders.

PRACTICE 2: *Institute periodic security awareness training for all employees.*
A culture of security awareness must be instilled in the organization so that all employees understand the need for policies, procedures, and technical controls. The first line of defense from insider threats is the employees themselves. All employees in an organization must understand that security policies and procedures exist, that there is a good reason why they exist, that they must be enforced, and that there can be serious consequences for infractions. Each employee needs to be aware of the organization's security policies and the process for reporting policy violations.

PRACTICE 3: *Enforce separation of duties and least privilege*.
If all employees are adequately trained in security awareness, and responsibility for critical functions is divided among employees, the possibility that one individual could commit fraud or sabotage without the cooperation of another individual within the organization is limited. Effective separation of duties requires the implementation of *least privilege*, that is, authorizing people only for the resources they need to do their jobs.

PRACTICE 4: *Implement strict password and account management policies and practices.*
No matter how vigilant employees are in trying to prevent insider attacks, if the organization's computer accounts can be compromised, insiders have an opportunity to circumvent both manual and automated mechanisms in place to prevent insider attacks.

PRACTICE 5: *Log, monitor, and audit employee online actions.*
If account and password policies and procedures are enforced, an organization can associate online actions with the employee who performed them. Logging, periodic monitoring, and auditing provide an organization the opportunity to discover and investigate suspicious insider actions before more serious consequences ensue.

PRACTICE 6: *Use extra caution with system administrators and privileged users*.
Typically, logging and monitoring is performed by a combination of system administrators and privileged users. Therefore, additional vigilance must be devoted to those users.

PRACTICE 7: *Actively defend against malicious code*.
System administrators or privileged users can deploy logic bombs or install other malicious code on the system or network. These types of attacks are stealthy and therefore difficult to detect ahead of time, but practices can be implemented for early detection.

PRACTICE 8: *Use layered defense against remote attacks*.
If employees are trained and vigilant, accounts are protected from compromise, and employees know that their actions are being logged and monitored, then disgruntled insiders will think twice about attacking systems or networks at work. Insiders tend to feel more confident and less inhibited when they have little fear of scrutiny by coworkers; therefore, remote access policies and procedures must be designed and implemented very carefully.

PRACTICE 9: *Monitor and respond to suspicious or disruptive behavior*.
In addition to monitoring online actions, organizations should closely monitor other suspicious or disruptive behavior by employees in the workplace. Policies and procedures should be in place for employees to report such behavior when they observe it in coworkers, with required follow-up by management.

PRACTICE 10: *Deactivate computer access following termination*.
When an employee terminates employment, whether the circumstances were favorable or not, it is important that the organization have in place a rigorous termination procedure that disables all of the employee's access points to the organization's physical locations, networks, systems, applications, and data.

PRACTICE 11: *Collect and save data for use in investigations*.
Should an insider attack, it is important that the organization have evidence in hand to identify the insider and follow up appropriately.

PRACTICE 12: *Implement secure backup and recovery processes*.
Despite all of the precautions implemented by an organization, it is still possible that an insider will attack. Therefore, it is important that organizations prepare for that possibility by implementing secure backup and recovery processes that are tested periodically.

PRACTICE 13: *Clearly document insider threat controls*.
As an organization acts to mitigate insider threat, clear documentation will help to ensure fewer gaps for attack, better understanding by employees, and fewer misconceptions that the organization is acting in a discriminatory manner.

Practice 1: Institute periodic enterprise-wide risk assessments.

Organizations need to develop a risk-based security strategy to protect its critical assets from both external and internal threats.

What to do?

It is not practical for most organizations to implement 100% protection against every threat to every organizational resource. Therefore, it is important to adequately protect critical information and other resources and not direct significant effort toward protecting relatively unimportant data and resources. A realistic and achievable security goal is to protect critical assets from both external and internal threats. Enterprise-wide risk assessments help organizations to identify critical assets and potential threats to those assets. Organizations should use the results of the assessment to develop or refine the overall strategy for securing their networked systems, striking the proper balance between countering the threat and accomplishing the organizational mission.[7]

Risk is generally understood to be the combination of threat, vulnerability, and mission impact. The lack of any one of these elements indicates the absence of risk.

The threat environment under which the system operates needs to be understood in order to accurately assess enterprise risk. Characterizing the threat environment can proceed in parallel with the evaluation of vulnerability and impact, however, the sooner the threat environment can be characterized the better. Unfortunately, many organizations focus on protecting information from access or sabotage by those external to the organization, and overlook insiders. Moreover, an information technology and security solution designed without consciously acknowledging and accounting for potential insider threats often leaves the role of protection in the hands of some of the potential threats—the insiders themselves. It is imperative that organizations recognize the potential danger posed by their employees' and contractors' knowledge and access, and specifically address that threat as part of an enterprise risk assessment.

Understanding the vulnerability of an organization to a threat is also important since if there is no vulnerability there can be no risk. But organizations often focus too much on low-level technical vulnerabilities, for example, using automated computer and network vulnerability scanners. While such techniques are important, our studies of insider threat have indicated that vulnerabilities in an organization's business processes are at least as important as vulnerabilities in the technologies used.

In addition, when an organization does not address vulnerabilities it can actually compromise its ability to achieve its mission. It will not reduce enterprise risk and will waste its time. Organizations need to manage the impact of threats rather than chase individual technical vulnerabilities.

The impact of insider threats may involve compromises to the integrity, availability, or confidentiality of information critical to an organization's mission. Insiders have affected the integrity of their organization's information in various ways, for example by

[7] See http://www.cert.org/nav/index_green.html for CERT research in Enterprise Security Management.

manipulating customer financial information or defacing their employer's web sites. They have also violated confidentiality of information by stealing trade secrets or customer information. Still others have disseminated confidential information to others, including private customer information as well as sensitive email messages between the organization's management. Finally, insiders have affected the availability of their organization's information by deleting data, sabotaging entire systems and networks, destroying backups, and other types of denial-of-service attacks.

In the types of insider incidents mentioned above, employees were able to compromise their organization's critical assets. It is important that protection strategies are designed focusing on those assets first: financial data, confidential or proprietary information, and other mission critical data. Once the critical assets have been determined, then potential impacts of policies, business processes, and technology on those assets can be assessed.

Case Studies: What could happen if I don't do it?

One organization failed to protect extremely critical systems and data from internal employees. It was responsible for running the 911 phone number to address lookup system for emergency services. An insider deleted the entire database and software from three servers in the organization's network operations center (NOC) by gaining physical access using a contractor's badge. The NOC, which was left unattended, was solely protected via physical security; all machines in the room were left logged in with system administrator access.

Although the NOC system administrators were immediately notified of the system failure via an automatic paging system, there were no automated failover mechanisms. The organization's recovery plan relied solely on backup tapes, which were also stored in the NOC. Unfortunately, the insider, realizing that the systems could be easily recovered, took all of the backup tapes with him when he left the facility. In addition, the same contractor's badge was authorized for access to the offsite backup storage facility, from which he next stole over fifty backup tapes.

Had a risk assessment been performed for this system prior to the incident, the organization would have consciously recognized the criticality of the systems, assessed the threats and vulnerabilities, and developed a risk mitigation strategy accordingly.

Another insider was the sole system administrator for his organization. One day, he quit with no prior notice. His organization refused to pay him for his last two days of work, and he subsequently refused to give them the passwords for the system administrator accounts for their systems. Over a period of three days, the insider modified the systems so that they could not be accessed by the employees, defaced the company web site, and deleted files. It is critical that organizations consider the risk they assume when they place all system administration power into the hands of a single employee.

Practice 2: Institute periodic security awareness training for all employees.

Without broad understanding and buy-in from the organization, technical or managerial controls will be short-lived.

What to do?

All employees need to understand that there is no "profile" of a malicious insider. Cases have involved both highly technical people and those with minimal understanding of the systems they exploited. Ages of perpetrators ranged from late teens to retirement. Both men and women have been malicious insiders, including introverted "loners," aggressive "get it done" people, and extroverted "star players." Positions have included low-wage data entry clerks, cashiers, programmers, artists, system and network administrators, salespersons, managers, and executives. They have been new hires, long-term employees, currently employed, recently terminated, contractors, and temporary employees.

Security awareness training should encourage identification of malicious insiders by behavior, not by stereotypical characteristics. Behaviors of concern include threats against the organization, bragging about the damage one could do to the organization, or discussing plans against the organization. Also of concern are attempts to gain employees' passwords or to obtain access through trickery or exploitation of a trusted relationship (often called "social engineering").

Training programs should create a culture of security appropriate for the organization, and include all personnel. For effectiveness and longevity, the measures used to secure an organization against insider threat need to be tied to the organization's mission, values, and critical assets, as determined by an enterprise-wide risk assessment. For example, if an organization places a high value on customer service quality, it may view customer information as its most critical asset and focus security on protection of that data. The organization could train its members to be vigilant against malicious employee actions, focusing on a number of key issues, including

- reducing risks to customer information by auditing access to customer records (see Practice 5)
- requiring separation of duties between employees who modify customer accounts and those who approve modifications or issue payments (see Practice 3)
- using secure backup and recovery methods to ensure availability of customer service data (see Practice 12)

Training on reducing risks to customer service processes would focus on

- protecting computer accounts used in these processes (see Practice 4)
- using malicious code detection tools (see Practice 7)
- detecting and reporting disruptive behavior by employees (see Practice 9)
- implementing proper system administration safeguards for critical servers (see practices 6, 8, and 10)

Training content should be based on documented policy (see Practice 13), including a confidential means of reporting security issues. Confidential reporting allows reporting of suspicious events without fear of repercussions, overcoming the cultural barrier of whistle blowing. Employees need to understand that the organization has policies and procedures and that they will respond to security issues in a fair and prompt manner.

Employees should be notified that system activity is monitored, especially system administration and privileged activity. All employees should be trained in their personal responsibility, such as protection of their own passwords and work products. Finally, the training should communicate Information Technology acceptable use policies.

Case Studies: What could happen if I don't do it?

The lead developer of a critical production application had extensive control over the application source code. The only copy of the source code was on his company-provided laptop; there were no backups performed, and very little documentation existed, even though management had repeatedly requested it. The insider told coworkers he had no intention of documenting the source code and any documentation he did write would be encrypted. He also stated that he thought poorly of his managers because they had not instructed him to make back-up copies of the source code.

A month after learning of a pending demotion, he erased the hard drive of his laptop, deleting the only copy of the source code the organization possessed, and quit his job. It took more than two months to recover the source code—after it was located by law enforcement in encrypted form at the insider's home. Another four months elapsed before the insider provided the password to decrypt the source code. During this time the organization had to rely on the executable version of the application, with no ability to make any modifications. If the insider's team members had been informed that the security and survivability of the system was their responsibility, and if they had been presented with a clear procedure for reporting concerning behavior, they might have notified management of the insider's statements and actions in time to prevent the attack.

Another insider case involved a much less technically sophisticated attack, but one that could have been avoided or successfully prosecuted if proper policies and training had been in place. Four executives left their firm to form a competing company. A few days before they left, one of them ordered a backup copy of the hard drive on his work computer, which contained customer lists and other sensitive information, from the external company that backed up the data. The company also alleged that its consulting services agreement and price list were sent by email from the insider's work computer to an external email account registered under his name. The insiders, two of whom had signed confidentiality agreements with the original employer, disagreed that the information they took was proprietary, saying that it had been published previously. Clear policies regarding definition of proprietary information and rules of use could have prevented the attack or provided a clearer avenue for prosecution.

Practice 3: Enforce separation of duties and least privilege.

While security awareness training is an excellent start, separation of duties and least privilege must be implemented to limit the damage that malicious insiders can inflict.

What to do?

Separation of duties requires dividing functions among people to limit the possibility that one individual could commit fraud or sabotage without the cooperation of another employee. One type of separation of duties, called *two-person rule*, is often used. It requires two people to participate in a task for it to be executed successfully. Examples include requiring two bank officials to sign large cashier's checks, or requiring verification and validation of source code before the code is released operationally. In general, employees are less likely to engage in malicious acts if they must collaborate with another employee.

Effective separation of duties requires implementation of *least privilege*, authorizing people only for the resources needed to do their job. Least privilege also reduces an organization's risk of theft of confidential or proprietary information by its employees.

Typically, organizations define roles that characterize the responsibilities of each job and the access to organizational resources needed to fulfill those responsibilities. Insider risk can be greatly mitigated by defining and separating roles responsible for key business processes and functions. For example

- requiring online management authorization for critical data entry transactions
- instituting code reviews for the software development and maintenance process
- using configuration management processes and technology to control software distributions and system modification
- designing auditing procedures to protect against collusion among auditors

Physical, administrative, and technical controls can be used to restrict employees' access to only those resources needed to accomplish their jobs. Access control gaps facilitated most incidents in the *Insider Threat Study*. For example, employees circumvented separation of duties enforced via policy rather than through technical controls. Ideally organizations should include separation of duties in the design of their business processes and enforce them via technical and non-technical means.

Access control based on separation of duties and least privilege is crucial to mitigating the risk of insider attack. These principles have implications in both the physical and the virtual worlds. In the physical world, organizations need to prevent employees from gaining physical access to resources not required by their work roles. Researchers need to have access to their laboratory space but do not need access to human resources file cabinets. Likewise, human resources personnel need access to personnel records but do not need access to laboratory facilities. There is a direct analogy in the virtual world in which organizations must prevent employees from gaining online access to information or services that are not required by their work roles. This kind of control is often called

role-based access control. Prohibiting access by personnel in one role from the functions permitted for another role limits the damage they can inflict if they become disgruntled or otherwise decide to exploit the organization for their own purposes.

Case Studies: What could happen if I don't do it?

In one case, a currency trader (who also happened to have a college minor in computer science) developed much of the software used by his organization to record, manage, confirm, and audit trades. He implemented obscure functionality in the software that enabled him to conceal his illegal trades. In this case, it was nearly impossible for auditors to detect his activities.

The insider, who consented to be interviewed for the *Insider Threat Study*, told the study researchers that problems can arise when "the fox is guarding the henhouse."[8] Specifically, the insider's supervisor managed both the insider and the auditing department responsible for ensuring his trades were legal or compliant. When auditing department personnel raised concern about the insider's activities, they were doing so to the insider's supervisor (who happened to be their supervisor as well). The supervisor directed auditing department personnel not to worry about the insider's activities and to cease raising concern, for fear the insider would become frustrated and quit.

This case illustrates two ways in which separation of duties can prevent an insider attack or detect it earlier:

- end users of an organization's critical systems should not be authorized to modify the system functionality or access the underlying data directly
- responsibility for maintaining critical data and responsibility for auditing that same data should never be assigned to the same person

In another case, a supervisor fraudulently altered U.S. immigration asylum decisions using his organization's computer system in return for payments of up to several thousand dollars per case, accumulating $50,000 over a two-year period. The insider would approve an asylum decision himself, request that one of his subordinates approve the decision, or overturn someone else's denial of an asylum application. Several foreign nationals either admitted in an interview or pleaded guilty in a court of law to lying on their asylum applications and bribing public officials to get approval of their applications. The organization had implemented separation of duties via role-based access control by limiting authorization for approving or modifying asylum decisions to supervisors' computer accounts. However, supervisors were able to alter any decisions in the entire database, not just those assigned to their subordinates. An additional layer of defense, least privilege, also could have been implemented to prevent supervisors from approving asylum applications or overturning asylum decisions with which they were not involved.

[8] *Insider Threat Study: Illicit Cyber Activity in the Banking and Finance Sector*. http://www.cert.org/archive/pdf/bankfin040820.pdf.

Practice 4: Implement strict password and account management policies and practices.

If the organization's computer accounts can be compromised, insiders can circumvent manual and automated control mechanisms.

What to do?

No matter how vigilant employees are about insider threats, if the organization's computer accounts can be compromised, insiders have an opportunity to circumvent mechanisms that are in place to prevent insider attacks. Therefore, computer account and password management policies and practices are critical to impede an insider's ability to use the organization's systems for illicit purposes. Fine-grained access control combined with proper computer account management will ensure that access to all of the organization's critical electronic assets

- is controlled so that unauthorized access is not possible
- is logged and monitored so that suspicious access can be detected and investigated
- can be traced from the computer account to the individual associated with that account

Some methods used by malicious insiders to compromise accounts included using password crackers, obtaining passwords through social engineering or because employees openly shared passwords, and using unattended computers left logged in. Password policies and procedures should ensure that all passwords are strong,[9] employees do not share their passwords with anyone, employees change their passwords periodically, and all computers execute password-protected screen savers. As a result, all activity from any account should be attributable to its owner. Employees should also report all attempts at account compromises rather than permit a compromise to happen due to ignorance of potential consequences or lack of a reporting mechanism.

Some insiders created backdoor accounts that provided them with system administrator or privileged access following termination. Other insiders found that shared accounts, like system administrator, DBA, and training accounts, were overlooked in the termination process and were still available to them. Periodic account audits combined with technical controls enable identification of

- backdoor accounts that could be used later for malicious actions by an insider, whether those accounts were specifically set up by the insider or were left over from a previous employee
- shared accounts whose password was known by the insider and not changed after that person's termination

The need for every account should be re-evaluated periodically. Limiting accounts to those that are absolutely necessary, with strict procedures and technical controls that enable auditors or investigators to trace all online activity on those accounts to an

[9] See *Choosing and Protecting Passwords*: http://www.us-cert.gov/cas/tips/ST04-002.html.

individual user, diminishes an insider's ability to conduct malicious activity without being identified. Account management policies that include strict documentation of all access privileges for all users enable a straightforward termination procedure that reduces the risk of attack by terminated employees.

Case Studies: What could happen if I don't do it?

A disgruntled software developer downloaded the password file from his organization's UNIX server to his desktop. Next, he downloaded a password cracker from the Internet and proceeded to "break" approximately forty passwords, including the root password. Fortunately, he did no damage, but he did access parts of the organization's network for which he was not authorized. The insider was discovered when he bragged to the system administrator that he knew the root password. As a result, his organization modified its policies and procedures to implement countermeasures to prevent such attacks in the future. System administrators were permitted to run password crackers and notify users with weak passwords, and it improved security training for employees on how and why to choose strong passwords.

A second case also illustrates the importance of employee awareness of password security. Two temporary data entry clerks and one permanent employee were able to embezzle almost $70,000 from their company by fraudulently using other employees' computer accounts. The employees within their group openly shared their passwords to enhance productivity. The system's role-based access provided the other employees' accounts with access to privileged system functions. The clerks used those accounts without authorization to subvert the business process governing vendor payment. First, they entered valid data into the database using their own accounts. Then they used the other, unauthorized accounts to modify the vendor's name and address to that of a friend or relative, issued the check from the system, and then modified the data back to the original, valid vendor information. The fraud was discovered after almost five months when an accountant in the general ledger department noticed that the number of checks issued was larger than normal, and further investigation revealed the irregularities in the handling of the checks.

Practice 5: Log, monitor, and audit employee online actions.

Logging, monitoring, and auditing can lead to early discovery and investigation of suspicious insider actions.

What to do?

If account and password policies and procedures are in place and enforced, an organization has a good chance of clearly associating online actions with the employee who performed them. Logging, monitoring, and auditing provide an organization with the opportunity to discover and investigate suspicious insider actions before more serious consequences ensue.

Auditing in the financial community refers to examination and verification of financial information. In the technical security domain it refers to examination and verification of various network, system, and application logs or data. To prevent or detect insider threats, it is important that auditing involve the review and verification of *all* of the organization's critical assets.[10] Furthermore, auditing must examine and verify the integrity as well as the legitimacy of logged access.

Automated integrity checking should be considered for flagging suspicious transactions that do not adhere to predefined business rules for manual review. Insider threats are most often detected by a combination of automated logging and manual monitoring or auditing. For example, integrity checking of computer account creation logs involves automated logging combined with manual verification that every new account has been associated with a legitimate system user and that the user is aware of the account's existence. Likewise, data audits typically involve manual processes, such as comparing electronic data modification history to paper records or examining electronic records for suspicious discrepancies.

Auditing should be both ongoing and random. If employees are aware that monitoring and auditing is a regular, ongoing process and that it is a high priority for the individuals who are responsible for it, it can serve as a deterrent to insider threats. For example, if a disgruntled system administrator is aware that all new computer accounts are reviewed frequently, then it is less likely that he or she will create backdoor accounts for later malicious use.

On the other hand, it probably is not practical to institute daily monitoring of every financial transaction in a financial institution. Monthly and quarterly auditing provides one layer of defense against insiders, but it also provides a predictable cycle on which insiders could design a fraud scheme that could go undetected over a long period of time. Random auditing of all transactions for a given employee, for example, could add just enough unpredictability to the process to deter an insider from launching a contemplated attack.

---

[10] Many risk management methodologies are based on protection of critical assets. For example, see the OCTAVE® (Operationally Critical Threat, Asset, and Vulnerability Evaluation℠) risk-based strategic assessment and planning technique for security: http://www.cert.org/octave/.

Case Studies: What could happen if I don't do it?

A large international company, while performing remote access monitoring, noticed that a former consultant had obtained unauthorized access to its network and created an administrator account. This prompted an investigation of the former insider's previous online activity, revealing he had run several different password-cracking programs on the company's network five different times over a ten-month period. Initially, he stored the cracked passwords in a file on the company's server. Later he installed a more sophisticated password-cracking program on the company's system. This program enabled him to automatically transfer all accounts and passwords that could be cracked to a remote computer on a periodic basis. Five thousand passwords for company employees were successfully transferred. This case illustrates the importance of logging and proactive monitoring. Because of those practices, this insider's actions were detected before any malicious activity was committed using the accounts and passwords or the backdoor account.

Another insider attack provides a contrasting example—one in which lack of auditing permitted the insider to conduct an attack that was less technically sophisticated but that enabled him to steal almost $260,000 from his employer over a two-year period. The insider was the manager of a warehouse. The attack proceeded as follows:

- The insider convinced his supervisor that he needed privileged access to the entire purchasing system for the warehouse.
- Next, he added a fake vendor to the list of authorized suppliers for the warehouse.
- Over the next two years, he entered 78 purchase orders for the fake vendor, and, although no supplies were ever received, he also authorized payment to the vendor.

The insider was aware of approval procedures, and all of his fraudulent purchases fell beneath the threshold for independent approval. The bank account for the vendor happened to be owned by the insider's wife. The fraud was accidentally detected by a finance clerk who noticed irregularities in the paperwork accompanying one of the purchase orders. This fraud could have been detected earlier by closer monitoring of online activities by privileged users, particularly since this particular user possessed unusually extensive privileged access. In addition, normal auditing procedures could have validated the new vendor, and automated integrity checking could have detected discrepancies between the warehouse inventory and purchasing records.

Practice 6: Use extra caution with system administrators and privileged users.

System administrators and privileged users have the technical ability, access, and oversight responsibility to commit and conceal malicious activity.

What to do?

Recall that the majority of the insiders who committed sabotage, and over half of those that stole confidential or proprietary information, held technical positions. Technically sophisticated methods of carrying out and concealing malicious activity included writing or use of scripts or programs (including logic bombs), creation of backdoor accounts, installation of remote system administration tools, modification of system logs, planting of viruses, and use of password crackers.

System administrators and privileged users[11] by definition have a higher system, network, or application access level than other users. This higher access level comes with higher risk due to the following:

- They have the technical ability and access to perform actions that ordinary users cannot.
- They can usually conceal their actions, since their privileged access typically provides them the ability to log in as other users, to modify system log files, or to falsify audit logs and monitoring reports.

Techniques that promote non-repudiation of action ensure that online actions taken by users, including system administrators and privileged users, can be attributed to the person that performed them. Therefore, should malicious insider activity occur, non-repudiation techniques allow that activity to be attributed to a single employee. Policies, practices, and technologies exist for configuring systems and networks to facilitate non-repudiation. However, keep in mind that system administrators and other privileged users will be the ones responsible for designing, creating, and implementing those policies, practices, and technologies. Therefore, separation of duties is also very important: network, system, and application security designs should be created, implemented, and enforced by multiple privileged users.

Even if online actions can be traced to the person who engaged in the action, it is unreasonable to expect that all user actions can be monitored proactively. Therefore, while the practices discussed above ensure identification of users following detection of suspicious activity, additional steps must be taken by organizations to defend against malicious actions before they occur. For instance, system administrators and privileged users have access to all information within their domains. Technologies such as

[11] The term "privileged users" refers to users who have an elevated level of access to a network, computer system, or application that is short of full system administrator access. For example, database administrators (DBAs) and webmasters are privileged users as they have the ability to create new user accounts and control the access rights of users within their domains.

encryption can be implemented to prevent such users from reading or modifying sensitive files to which they should not have access.

Policies, procedures, and technical controls should enforce separation of duties and require actions by multiple users for all modifications to critical systems, networks, applications, and data. In other words, no single user should be permitted or be technically able to release changes to the production environment without online action by a second user. These controls would prevent an insider from releasing a logic bomb without detection by another employee.

Finally, many of the insiders in the *Insider Threat Study* were former employees; therefore, organizations must be particularly careful in disabling access, particularly for former system administrators and privileged users. Thoroughly documented procedures for disabling access can help ensure that stray access points are not overlooked. In addition, the two-person rule should be considered for the critical functions performed by these users to reduce the risk of extortion after they leave the organization.

Case Studies: What could happen if I don't do it?

A system administrator at an international financial organization heard rumors that the annual bonuses were going to be lower than expected. He began constructing a logic bomb at home and used authorized remote access to move the logic bomb to the company's servers as part of the typical server upgrade procedure over a period of two and a half months. When he was informed by his supervisor that his bonus would be significantly lower than he had expected, he terminated his employment immediately. Less than two weeks later, the logic bomb went off at 9:30 a.m., deleting 10 billion files on approximately 1,000 servers throughout the United States. The victim organization estimated that it would cost more than $3 million to repair its network, and the loss affected 1.24 billion shares of its stock.

In another case, an insider was promoted from one position to another within the same organization. Both positions used the same application for entering, approving, and authorizing payments for medical and disability claims. The application used role-based access to enforce separation of duties for each system function. However, when this particular insider was promoted, she was authorized for her new access level, but administrators neglected to rescind her prior access level (separation of duties was inadequately enforced). As a result, she ended up having full access to the application, with no one else required to authorize transactions (payments) from the system. She entered and approved claims and authorized monthly payments for her fiancé, resulting in payments of over $615,000 over almost two years.

Practice 7: Actively defend against malicious code.

While insiders frequently use simple user commands to do their damage, logic bombs and other malicious code are used frequently enough to be of concern.

What to do?

Many organizations defend against malicious code using antivirus software and host or network firewalls. While these defenses are useful against external infections, their value is limited in preventing attacks by malicious insiders in two important respects: they do not work against new or novel malicious code (including logic bombs planted by insiders) and they are concerned primarily with material spread through networking interfaces rather than installed directly on a machine. To deal with these limitations, a more systematic and active approach is needed.

First, organizations should identify baseline software and hardware configurations. An organization may have several baseline configurations, given the different computing and information needs of different users (accountant, manager, programmer, receptionist). But as configurations are identified, the organization should characterize the hardware and software that makes up those configurations.

Characterization can be a basic catalog of information, tracking information like versions of installed software, hardware devices, and disk utilization. However, such basic characterizations can be easily defeated, so more comprehensive characterizations are often required. These characterizations include

- cryptographic checksums (using SHA-1 or MD5, for example)
- interface characterization (such as memory mappings, device options, and serial numbers)
- recorded configuration files

Once this information is captured, computers implementing each configuration can be validated by comparing it against the baseline copy. Discrepancies can then be investigated to determine whether they are benign or malicious. Using these techniques, changes to system files or the addition of malicious code will be flagged for investigation. There are tools called *file integrity checkers* that partially automate this process and provide for scheduled sweeps through computer systems.[12]

Computer configurations do not remain unchanged for long. Therefore, characterization and validation should be part of an organization's configuration management process. For protection against malicious insiders, part of the configuration management process should be separation of duties. For example, validation of a configuration should be done by a person other than the one who made changes so that there is an opportunity to detect and correct malicious changes (including planting of logic bombs).

[12] See http://www.sans.org/resources/idfaq/integrity_checker.php for a discussion of file integrity checkers.

Case Studies: What could happen if I don't do it?

A manufacturing firm's system administrator began employment as a machinist. Because of his technical ability he, over a ten-year period, created the company's network and had sole authority for system administration. The company eventually expanded, opening additional offices and plants nationally and internationally. The insider

- began to feel disgruntled at his diminishing importance to the company
- launched verbal and physical assaults on coworkers
- sabotaged projects of which he was not in charge
- loaded faulty programs to make coworkers look bad

He received a verbal warning, two written reprimands, was demoted, and finally fired as a result of his actions. A few weeks later, a logic bomb executed on the company's network, deleting one thousand critical manufacturing programs from the company's servers. The estimated cost of the damage exceeded $10 million, leading to the layoff of approximately eighty employees. The investigation revealed that the insider had actually tested the logic bomb three times on the company's network prior to his termination.

Practices for detection of malicious code would have detected that a new program had been released with timed execution. Configuration control procedures with a two-person rule for release of system-level programs, and characterization procedures, could have detected the release of a new system file that was not part of the original system baseline.

Another organization built automated monitoring into its software that sent automatic notification to the security officer any time a highly restricted screen was used to modify information stored in the database. Role-based access control restricted access to this screen to a few privileged users; the automated notification provided a second layer of defense against illegal data modification using that function. However, a developer of the application who happened to have access to that function modified the code so that the automated notification was no longer sent. He then proceeded to use the function to steal a large sum of money from his employer.

Interestingly, the organization had a comprehensive configuration management system in place for software changes. When a program was compiled, a report was produced listing which files were compiled, by which computer account, and when. It also listed modules added, modified, or deleted. Unfortunately, this report was not monitored, and therefore the application changes were not detected during the year and a half over which the fraud was committed. Had it been monitored, or had the configuration control system enforced a two-person rule for releasing new versions of software, the removal of the security notification would have been detected and the insider could not have committed the fraud.

Practice 8: Use layered defense against remote attacks.

Remote access provides a tempting opportunity for insiders to attack with less risk.

What to do?

Insiders often attack organizations remotely using legitimate access provided by the organization, or following termination. While remote access can greatly enhance employee productivity, caution is advised when remote access is provided to critical data, processes, or information systems. Insiders have admitted that it is easier to conduct malicious activities from home because it eliminates the concern that someone could be physically observing the malicious acts.

The vulnerabilities inherent in allowing remote access suggest that multiple layers of defense should be built against remote attack. Organizations may provide remote access to email and non-critical data but should strongly consider limiting remote access to the most critical data and functions. Access to data or functions that could inflict major damage to the company should be limited to employees physically located inside the workplace as much as possible. Remote system administrator access should be limited to the smallest group practicable, if not prohibited altogether.

When remote access to critical data, processes, and information systems is deemed necessary, the organization should offset the added risk with closer logging and frequent auditing of remote transactions. Information such as login account, date/time connected and disconnected, and IP address should be logged for all remote logins. It also is useful to monitor failed remote logins, including the reason the login failed. If authorization for remote access to critical data is kept to a minimum, monitoring can become more manageable and effective.

Disabling remote access is an often overlooked but critical part of the employee termination process. It is critical that employee termination procedures include

- disabling remote access accounts (such as VPN and dial-in accounts)
- disabling firewall access
- changing the passwords of all shared accounts (including system administrator, database administrator [DBA], and other privileged shared accounts)
- closing all open connections

A combination of remote access logs, source IP addresses, and phone records usually helps to identify insiders who launch remote attacks. Identification can be straightforward because the user name of the intruder points directly to the insider. Of course, corroboration of this information is required, because the intruders might have been trying to frame other users, cast attention away from their own misdeeds by using other users' accounts, or otherwise manipulate the monitoring process.

Case Studies: What could happen if I don't do it?

For a period of five years, a foreign currency trader with an investment bank "fixed" the bank's records to make his trading losses look like major gains for the bank. His actions made it appear that he was one of the bank's star producers, resulting in lucrative bonuses for his perceived high performance. In actuality, the bank lost hundreds of millions of dollars and drew a large amount of negative media attention as a result of his actions. While initially most of the insider's fraud occurred at work, he increasingly found it easier to conduct his illicit activities from home in the middle of the night because he did not have to worry about anyone in the office or at home looking over his shoulder. Therefore, the risk that other traders would find out about his fraudulent activities was reduced significantly.

In an interview for the *Insider Threat Study*, the insider said that group trading (trading by a team of traders), rather than individual trading, can help mitigate an organization's risks, because it is easier to detect illegal or suspicious trading practices when there are multiple team members trading from the same account.[13] In this case isolated trading, along with the anonymous nature of remote access, emboldened the insider to continue a fraud in which he otherwise might not have engaged.

In another case, a government organization notified one of its contract programmers that his access to a system under development was being eliminated and that his further responsibilities would be limited to testing activities. After his protests were denied, the programmer quit the organization. Then, three times over a two-week period, the insider used a backdoor into the system with administrator privilege (which he presumably installed before leaving) to download source code and password files from the developmental system. The unusually large size of the remote downloads raised red flags in the organization, which resulted in an investigation that traced the downloads to the insider's residence and led to his arrest, prosecution, and imprisonment. This case demonstrates the value of vigilant monitoring of remote access logs and reaction to suspicious behavior in limiting damage to the organization's interests.

[13] *Insider Threat Study: Illicit Cyber Activity in the Banking and Finance Sector*. http://www.cert.org/archive/pdf/bankfin040820.pdf.

Practice 9: Monitor and respond to suspicious or disruptive behavior.

One method of reducing the threat of malicious insiders is to proactively deal with difficult employees.

What to do?

An organization's methods of dealing with difficult individuals should start in the hiring process. A consistent practice of performing background checks and evaluating individuals based on the information received can reduce insider threats. Background checks should investigate previous criminal convictions, verify credentials and past employment, and include discussions with prior employers regarding the individual's competence and approach to dealing with workplace issues. While this information may not be the dominant component in the hiring process (and, arguing fairness, should not be), the information may help in dealing proactively with the individual. The *Insider Threat Study* revealed a surprisingly high number of malicious insiders with prior criminal convictions when hired.[14] These proactive measures should not be punitive in nature; rather, the individual should be indoctrinated into the organization with appropriate care.

Organizations should invest time and resources in training supervisors to recognize and respond to inappropriate or concerning behavior. Many times, less serious but inappropriate behavior is noticed in the workplace but not necessarily acted on because it does not rise to the level of a policy violation. However, failure to define or enforce security policies may embolden employees to commit repeated violations that escalate in severity, eventually resulting in insider IT sabotage. It is important that organizations consistently investigate and respond to all rule violations committed by employees.

Given that financial gain is a primary motive for much insider fraud, organizations should monitor indications by employees of possible financial problems or unexplained financial gain. Sudden changes in an employee's financial situation, including increasing debt or expensive purchases, may be indicators of potential insider threat.

Once employed, if an employee's behavior becomes suspicious, the organization must act with due care in dealing with it. Policies and procedures must exist for employees to report their concerns or to report disruptive behavior by others, and reports should always be investigated. (Some checks and balances must exist to limit frivolous reporting.) Disruptive employees should not be allowed to migrate from one position to another within the enterprise, evading documentation of disruptive or concerning activity. Threats, malicious boasting ("You wouldn't believe how easily I could trash this net!") and other negative sentiments should also be treated as concerning behavior. Many employees will have concerns and grievances from time to time in an organization, and

[14] *Insider Threat Study: Illicit Cyber Activity in the Banking and Finance Sector*. http://www.cert.org/archive/pdf/bankfin040820.pdf and *Insider Threat Study: Computer System Sabotage in Critical Infrastructure Sectors*. http://www.cert.org/archive/pdf/insidercross051105.pdf.

the provision of a formal and accountable process for addressing those grievances may act to satisfy those who might otherwise resort to malicious activity. In general, any employee experiencing difficulties in the workplace, who also has access to critical information assets, should be aided in the resolution of those difficulties.

Once concerning behavior is identified, several steps may aid an organization in managing risks of malicious activity. First, the employee's access to critical information assets should be evaluated. His or her level of network access should also be considered. While this is done, the organization should provide options to the individual for coping with the behavior, including access to a confidential employee assistance program.

Case Studies: What could happen if I don't do it?

A system administrator was hired to run the engineering department for an organization and three months later was named as the lead for a major new project. He then began to bully his coworkers, and was taken off the project a month after it started. Less than two months after that, he was terminated for poor performance and conduct. Customers had complained that he was rude, and coworkers said that he thought he was better than everyone else. His superiors realized that he was not as good technically as they had originally believed and suspected that he was attempting to hide that fact by criticizing others. The company did provide counseling, but he resented it.

Almost two months after his termination, the insider obtained a system administrator account password from a female employee, still with the company, with whom he'd had a relationship. Using this password, the insider was able to hide the project folder on the server that was needed the next day for an important customer demonstration. Although the company did employ standard recommendations in handling this insider, he still managed to sabotage the company's system. This case highlights the fact that companies should consider social relationships that terminated insiders have with employees still working for the company.

One insider, working as a vice president for engineering and responsible for oversight of all software development in the company, was engaged in a long-running dispute with higher management. This dispute was characterized by verbal attacks by the insider and statements to colleagues about the degree of upset he had caused to management. The insider engaged in personal attacks once or twice a week and on one occasion in a restaurant screamed personal attacks at the CEO of the company. A final explosive disagreement led the insider to quit.

When no severance package was offered, he copied a portion of a product under development to removable media, deleted it from the company's server, and removed the recent backup tapes. He then offered to restore the software in exchange for $50,000. He was charged and convicted of extortion, misappropriation of trade secrets, and grand theft. However, the most recent version of the software was never recovered. If the organization had paid attention to earlier disruptive behavior and acted to secure assets against his access, substantial losses could have been avoided.

Practice 10: Deactivate computer access following termination.

It is important that organizations follow rigorous termination procedures that disable all open access points to the networks, systems, applications, and data.

What to do?

While employed, insiders have legitimate, authorized access to the organization's network, system, applications, and data. Once employment is terminated, it is important that the organization have in place and execute rigorous termination procedures that disable all open access points. Otherwise, the organization's network is vulnerable to access by a now-illegitimate, unauthorized user. Some organizations choose to permit continued access by former employees for some time period under favorable termination circumstances; it is important that organizations have a formal policy in place for these circumstances and carefully consider the potential consequences.

If formal termination policies and procedures are not in place, the termination process tends to be ad hoc, posing significant risk that one or more access points will be overlooked. The *Insider Threat Study* shows that insiders can be quite resourceful in exploiting obscure access mechanisms neglected in the termination process. If a formal process exists, it must be strictly followed. It is also critical that organizations remain alert to new insider threat research and periodically review and update these processes.

Some aspects of the termination process are quite obvious, such as disabling the terminated employee's computer account. However, organizations that have been victims of insider attacks were often vulnerable because of poor, non-existent, or non-comprehensive account management procedures. Many employees have access to multiple accounts; *all* account creations should be tracked and periodically reviewed to ensure that all access can be quickly eliminated when an employee is terminated.

Accounts sometimes overlooked in the termination process are shared accounts, such as system administrator accounts and database administrator (DBA) accounts. In addition, some applications require administrative accounts that are frequently shared among multiple users. It is important that the organization meticulously maintain a record of every shared account and every user authorized to have the password to each.

Remote access is frequently exploited by former insiders. Remote access or virtual private network (VPN) accounts must be disabled, as well as firewall access, in order to prevent future remote access by the terminated employee. In addition, any remote connections already open by that employee must be closed immediately.

In summary, a layered defensive model that accounts for all access methods should be implemented. Remote access should be disabled, but if an obscure remote access method is overlooked, the next layer of defense is accounts. All accounts should be disabled for use by the former employee, so that even if remote access is established, the insider is prevented from proceeding further. Therefore, it is important that intranet accounts, application-specific accounts, and all other accounts for which the user was authorized be

disabled. Also, keep in mind that if the terminated insider was responsible for establishing accounts for others, such as employees, customers, or external web site users, then those accounts could also be accessible to the terminated insider.

Finally, termination procedures must include steps to prevent physical access. Insiders have exploited physical access to gain access to their former employer's systems.

Case Studies: What could happen if I don't do it?

A credit union's system administrator was terminated suddenly with no notice that his employer was dissatisfied with his work. That night he suspected that his replacement, who he felt was technically inferior, had not disabled his access. He attempted to access the system from home and found that his replacement had failed to disable his access through the company firewall. Although his account had been disabled, she had failed to change the password of the system administrator account. The insider used that account to shut down the organization's primary server, one that had been having problems and had in fact crashed the previous weekend (and had taken him an entire weekend to bring up again). It took the credit union three days to bring the server back into service; during that time none of its customers were able to access any of their accounts in any way. This case illustrates the necessity of thoroughly disabling access, as well as the consequences when an organization has no competent backup for a single system administrator.

In another case, a system administrator logged in one morning and was notified by her custom-written login software that her last login was one hour earlier. This set off immediate alarms, as she had in fact not logged in for several days. She had previously taken steps to redirect logging of actions by her account to a unique file rather than the standard shell history file. Therefore, she was able to trace the intruder's steps and saw that the intruder had read another employee's email using her account, then deleted the standard history file for her account so that there would be no log of his actions.

The login was traced to a computer at a subsidiary of the company. Further investigation showed that the same computer had logged into the company's system periodically for the past month. Monitoring revealed that a former employee had accessed up to sixteen of his former employer's systems on a daily basis during working hours. The insider

- gained access to at least 24 user accounts
- read electronic mail
- reviewed source code for his previous project
- deleted two software modification notices for the project

The former employee had been terminated for non-performance and then went to work for the subsidiary. This case illustrates the importance of terminating access completely for former employees, careful monitoring for post-termination access, and paying particular attention to terminated technical employees.

Practice 11: Collect and save data for use in investigations.

Collecting and saving usable evidence preserves response options, including legal options.

What to do?

The first questions that often follow any computer incident, whether malicious or not, are "what happened?" and "who is responsible?" In cases where malicious insiders are suspected, these questions are particularly urgent. Answering these questions in an actionable manner requires a detailed record of system and network actions. However, malicious insiders may act to corrupt, falsify, or delete such a record, impacting options for corrective and responsive actions.

System logs were used to identify the insider in most cases. Furthermore, system logs served as evidence in many of the insiders' trials. To best protect critical information and equipment, multiple sources of information should be maintained, particularly sources that may support one another. This includes logging

- data access (reading, modifying, or deleting data)
- application usage (when applications were started and exited and by which user)
- system commands and file change logs
- method of connection (console, local-area networked, dial-in, Internet)
- the source and destination of connections

Phone system and physical access records should also be maintained. As this information is collected, it should also be placed on backup media for archival storage.

As difficult as collecting all this information is, analysis of it is often harder. The signs of malicious insider activity can be subtle, such as an abnormal pattern or rate of data modification, or an off-hours download of information the insider is authorized to read. Log files need to be monitored periodically to try to identify such situations.

Organizations may need to involve a forensics specialist, both to design a routine analysis procedure for identifying malicious insiders and for more specialized analysis once the insider is identified. There have been insider threat cases in which inappropriate handling of system logs has rendered them unacceptable for prosecution. In the event of a suspected security incident, involve an expert in the investigation of electronic crimes.

Case Studies: What could happen if I don't do it?

An employee of a subcontractor for a government agency was nearing completion of his contract. Ordinarily, under these circumstances, the government agency would offer the employee a permanent position if his or her performance had been satisfactory. The insider initiated this hiring process and was required to take a drug test. The drug test results came back positive for cocaine, so his employment possibilities for the agency were forfeited. He remained employed with his current

employer for a few days until that organization was notified of his drug test results and terminated his employment immediately. His physical access cards were confiscated, he was escorted from the building, his personal computer account was disabled, and the password was changed on the system administrator account to which he had access.

The following Monday morning, the subcontractor's system was down. The logs showed that the system had been shut down via commands from an account that was not associated with any legitimate user. Remote access logs showed that attempts to log in began Friday evening and continued until successful login early Saturday. Once authenticated, the user deleted a number of printer drivers in the system, altered and changed certain user passwords, and finally entered the command to shut down the system. The logs on the remote access server stored the phone number of the incoming connections, and it was traced to the home address of the terminated insider. These logs were key in successfully prosecuting the insider.

In contrast to the above case, in which logs were stored appropriately and used to identify the user for prosecution, the following case illustrates the opposite: A contractor for a large company was responsible for handling customer service calls. A fraud scheme conducted by four employees over a period of almost three years resulted in losses for the company of $500,000. However, once the fraud was suspected by the company's fraud investigator, it was discovered that, since the company "recycled" its computer logs, they only provided specific activity by login name and computer terminal as far back as one month. Fortunately, one of the employees involved testified as to the history and duration of the fraud. This case illustrates the importance of securely backing up system logs for long time periods in case they are needed for investigations or prosecution.

Practice 12: Implement secure backup and recovery processes.

Despite all of the precautions implemented by an organization, it is still possible that an insider will attack. Therefore, it is important that organizations prepare for that possibility by implementing secure backup and recovery processes that are tested periodically.

What to do?

Prevention of insider attacks is the first line of defense. However, experience has taught that there will always be avenues for determined insiders to successfully compromise a system. Effective backup and recovery processes need to be in place and operational so that if compromises do occur business operations can be sustained with minimal interruption. Our research has shown that effective backup and recovery mechanisms can make the difference between

- several hours of downtime to restore systems from backups
- weeks of manual data entry when backups are not available
- months or years to reconstruct information for which no backup copies existed

When possible, multiple copies of backups should exist, with redundant copies stored offsite in a secure facility. Different people should be responsible for the safekeeping of each copy so that it would require the cooperation of multiple individuals to compromise the means to recovery.

System administrators should ensure that the physical media on which backups are stored are also protected from insider corruption or destruction. Insider cases in our research have involved attackers who

- deleted backups
- stole backup media (including offsite backups in one case)
- performed actions that could not be undone due to faulty backup systems

Some system administrators neglected to perform backups in the first place, while others sabotaged established backup mechanisms. Such actions can amplify the negative impact of an attack on an organization by eliminating the only means of recovery. To guard against insider attack, organizations must

- perform and periodically test backups
- protect media and content from modification, theft, or destruction
- apply separation of duties and configuration management procedures to backup systems just as they do for other system modifications

Unfortunately, some attacks against networks may interfere with common methods of communication, thereby increasing uncertainty and disruption in organizational activities, including recovery from the attack. This is especially true of insider attacks, since insiders are quite familiar with organizational communication methods and, during attack, may interfere with communications essential to the organization's data backup process. Organizations can mitigate this effect by maintaining trusted communication

paths outside of the network with sufficient capacity to ensure critical operations in the event of a network outage. This kind of protection would have two benefits: the cost of strikes against the network would be mitigated, and insiders would be less likely to strike against connectivity because of the reduced impact.

Case Studies: What could happen if I don't do it?

Centralization of critical assets and sabotage of backups has enabled some insiders to amplify the impact of their attacks by eliminating redundant copies and avenues for recovery. One insider, the sole system administrator, centralized the only copy of all of the company's critical production programs on a single server and convinced management to institute policies mandating this practice. That server was later the target of a logic bomb written by the same insider. No other current copy of the software was available to recover from the attack, since he had also requested and received, through intimidation, the only backup tape, violating company policy. The logic bomb, which deleted all of the company's programs, cost the company millions of dollars and caused company-wide layoffs. While centralization can contribute to the efficiency of an organization, care must be taken that backups are performed regularly and are protected to ensure business continuity in the event of damage to or loss of centralized data.

In another case, an insider was terminated because of his employer's reorganization. The company followed proper procedure by escorting the insider to his office to collect his belongings and then out of the building. The IT staff also followed the company's security policy by disabling the insider's remote access and changing passwords. However, they overlooked one password that was known to three people in the organization; the terminated insider used that account to gain access to the system that night and to delete the programs he had created while working there. Some of these programs supported the company's critical applications.

Restoration of the deleted files from backup failed. While the insider had been responsible for backups, company personnel believe that the backups were not maliciously corrupted. The backups had simply not been tested to ensure that they were properly recording the critical data. As a result, the organization's operations in North and South America were shut down for two days, causing more than $80,000 in losses. This case illustrates the delay that can be caused in recovery following an insider attack if backups are not tested periodically.

Practice 13: Clearly document insider threat controls.

To ensure consistent handling and to protect against accusations of discrimination, procedures for dealing with malicious insiders must be clearly documented.

What to do?

Cases involving malicious insiders are difficult to handle. Relationships between management and employees may be strained, with individuals taking sides with the organization or with the employee. A clearly written set of policies and procedures, developed with protection of the rights of everyone involved in mind, may help to defuse this situation. All of the organization's efforts to control damage by malicious insiders should be identified, together with circumstances under which these efforts are appropriate. As individuals join the organization, they should receive a copy of this description that clearly lays out what is expected of them, together with the consequences of violations. Evidence that each individual has read and agreed to the organization's policies, such as the individual's signature, should be maintained.

This description should also form the basis of ongoing training as described in Practice 2. If the organization experiences damage due to a malicious insider or if other risks evolve, such as new forms of internal or external attack, the description and training should be updated. The training should be given periodically to all employees, to help individuals act properly in the organization.

Case Studies: What could happen if I don't do it?

An insider accepted a promotion, leaving a system administrator position in one department for a position as a systems analyst in another department of the same organization. In his new position, he was responsible for information sharing and collaboration between his old department and the new one. The following events ensued:

- The original department terminated his system administrator account and issued him an ordinary user account to support the access required in his new position.
- Shortly thereafter, the system security manager at the original department noticed that the former employee's new account had been granted unauthorized system administration rights.
- The security manager reset the account back to ordinary access rights, but a day later found that administrative rights had been granted to it once again.
- The security manager closed the account, but over the next few weeks other accounts exhibited unauthorized access and usage patterns.

An investigation of these events led to charges brought against the analyst for misuse of computing systems. These charges were eventually unsuccessful, in part because there was no clear policy regarding account sharing or exploitation of vulnerabilities to elevate account privileges. This case illustrates the importance of clearly established policies that are consistent across departments, groups, and subsidiaries of the organization.

References/Sources of Best Practices

Alberts, Christopher; Dorofee, Audrey; Killcrece, Georgia; Ruefle, Robin; & Zajicek, Mark. *Defining Incident Management Processes for CSIRTs: A Work in Progress* (CMU/SEI-2004-TR-015). Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, 2004. http://www.sei.cmu.edu/publications/documents/04.reports/04tr015.html.

British Standards Institute. *IT Service Management Part 1: Specification for service management* (BS 15000-1:2001), 2005. http://www.bsonline.bsi-global.com/server/index.jsp.

CERT. Survivability and Information Assurance Curriculum (SIA), 2006. http://www.cert.org/sia (2006).

CERT. Virtual Training Environment (VTE), 2005. https://www.vte.cert.org/.

Corporate Information Security Working Group (CISWG). Adam H. Putnam, Chairman; Subcommittee on Technology, Information Policy, Intergovernmental Relations & the Census Government Reform Committee, U.S. House of Representatives. "Report of the Best Practices and Metrics Teams." November 17, 2004; updated January 10, 2005. http://www.educause.edu/LibraryDetailPage/666&ID=CSD3661.

Department of Homeland Security, National Cyber Security Division and Carnegie Mellon University. *BuildSecurityIn*, 2006. https://buildsecurityin.us-cert.gov/daisy/bsi/home.html.

Federal Financial Institutions Examination Council. *IT Examination Handbook: Information Security*. December 2002. http://www.ffiec.gov/ffiecinfobase/html_pages/it_01.html.

Information Security Forum. "Standard of Good Practice for Information Security," Version 4.1, January 2005. http://www.isfsecuritystandard.com/index_ie.htm.

Information Systems Audit and Control Association. *COBIT 4.0*. 2006. http://www.isaca.org.

International Standards Organization. "Information technology -- Security techniques -- Code of practice for information security management," (SO/IEC 17799:2005), 2005. http://www.iso.org/iso/en/CatalogueDetailPage.CatalogueDetail?CSNUMBER=39612&ICS1=35&ICS2=40&ICS3=.

International Standards Organization. "Information technology -- Security techniques -- Information security management systems – Requirements," (ISO/IEC 27001:2005), 2005. http://www.iso.ch/iso/en/CatalogueDetailPage.CatalogueDetail?CSNUMBER=4210.

MasterCard International. “Payment Card Industry Data Security Standard.” January 2005. https://sdp.mastercardintl.com/pdf/pcd_manual.pdf.

National Institute of Standards and Technology. “Generally Accepted Practices and Principles for Securing Information Systems” (NIST 800-14), September 1996. http://csrc.nist.gov/publications/nistpubs/800-14/800-14.pdf.

National Institute of Standards and Technology. “Engineering Principles for Information Technology Security (A Baseline for Achieving Security), Revision A” (NIST 800-27 Rev. A), June 2004. http://csrc.nist.gov/publications/nistpubs/800-27A/SP800-27-RevA.pdf.

National Institute of Standards and Technology. “Recommended Security Controls for Federal Information Systems,” (NIST 800-53), February 2005. http://csrc.nist.gov/publications/nistpubs/800-53/SP800-53.pdf.

National Institute of Standards and Technology. “Draft Special Publication 800-100: Information Security Handbook: A Guide for Managers” (NIST 800-100), June 7, 2006. http://csrc.nist.gov/publications/drafts/Draft-SP800-100_Handbook06-07-06.zip.

National Institute of Standards and Technology, Computer Security Resource Center. “Standards for Security Categorization of Federal Information and Information Systems” (FIPS PUB 199), February 2004. http://csrc.nist.gov/publications/fips/fips199/FIPS-PUB-199-final.pdf.

National Institute of Standards and Technology, Computer Security Resource Center. “Minimum Security Requirements for Federal Information and Information Systems” (FIPS PUB 200), March 2006. http://csrc.nist.gov/publications/fips/fips200/FIPS-200-final-march.pdf.

United Kingdom Office of Government Commerce, Information Technology Infrastructure Library. http://www.ogc.gov.uk/index.asp?docid=1000368.

Visa U.S.A., Inc. “Payment Card Industry Data Security Standard,” Version 1.0, December 15, 2004. http://usa.visa.com/download/business/accepting_visa/ops_risk_management/cisp_PCI_Data_Security_Standard.pdf.