

Insider Threat Study: Computer System Sabotage in Critical Infrastructure Sectors

Michelle Keeney, J.D., Ph.D.
Eileen Kowalski
National Threat Assessment Center
United States Secret Service
Washington, DC

Dawn Cappelli
Andrew Moore
Timothy Shimeall
Stephanie Rogers
CERT® Program
Software Engineering Institute
Carnegie Mellon University
Pittsburgh, PA

May 2005



Carnegie Mellon
Software Engineering Institute

Table of Contents

TABLE OF CONTENTS	2
SECTION 1: INTRODUCTION	3
PREVALENCE OF INCIDENTS OF INSIDER SABOTAGE	4
THE INSIDER THREAT STUDY	5
SECTION 2: CHARACTERISTICS OF INSIDER SABOTAGE ACROSS CRITICAL INFRASTRUCTURE SECTORS	11
INSIDER CHARACTERISTICS	11
ORGANIZATION CHARACTERISTICS	12
CONSEQUENCES TO INSIDERS	12
SECTION 3: KEY FINDINGS OF THE INSIDER THREAT STUDY OF SABOTAGE ACROSS CRITICAL INFRASTRUCTURE SECTORS	14
THE INSIDER’S MOTIVE	14
PRE-ATTACK BEHAVIOR AND PLANNING	15
ADVANCING THE ATTACK	16
DETECTING THE ATTACK	18
CONSEQUENCES FOR TARGETED ORGANIZATIONS	20
SECTION 4: IMPLICATIONS OF THE KEY FINDINGS FOR THE PREVENTION OF INSIDER SABOTAGE	22
THE INSIDER’S MOTIVE	22
PRE-ATTACK BEHAVIOR AND PLANNING	23
ADVANCING THE ATTACK	24
DETECTING THE ATTACK	30
CONSEQUENCES FOR TARGETED ORGANIZATIONS	32
SECTION 5: CONCLUSION: REFLECTIONS ON THE FINDINGS FOR THE PREVENTION OF INSIDER SABOTAGE	34
APPENDICES	35
APPENDIX A: INCIDENT DATE AND LOCATION	35
APPENDIX B: ORGANIZATIONAL SIZE	36
APPENDIX C: ADDITIONAL CASE EXAMPLES	37
APPENDIX D: GLOSSARY OF TECHNICAL TERMS	42
APPENDIX E: ACKNOWLEDGEMENTS	45

SECTION 1: INTRODUCTION

A system administrator, angered by his diminished role in a thriving defense manufacturing firm whose computer network he alone had developed and managed, centralized the software that supported the company's manufacturing processes on a single server, and then intimidated a coworker into giving him the only backup tapes for that software. Following the system administrator's termination for inappropriate and abusive treatment of his coworkers, a logic bomb previously planted by the insider detonated, deleting the only remaining copy of the critical software from the company's server. The company estimated the cost of damage in excess of \$10 million, which led to the layoff of some 80 employees.

An application developer, who lost his IT sector job as a result of company downsizing, expressed his displeasure at being laid off just prior to the Christmas holidays by launching a systematic attack on his former employer's computer network. Three weeks following his termination, the insider used the username and password of one of his former coworkers to gain remote access to the network and modify several of the company's web pages, changing text and inserting pornographic images. He also sent each of the company's customers an email message advising that the website had been hacked. Each email message also contained that customer's usernames and passwords for the website. An investigation was initiated, but it failed to identify the insider as the perpetrator. A month and a half later, he again remotely accessed the network, executed a script to reset all network passwords and changed 4,000 pricing records to reflect bogus information. This former employee ultimately was identified as the perpetrator and prosecuted. He was sentenced to serve five months in prison and two years on supervised probation, and ordered to pay \$48,600 restitution to his former employer.

A city government employee who was passed over for promotion to finance director retaliated by deleting files from his and a coworker's computers the day before the new finance director took office. An investigation identified the disgruntled employee as the perpetrator of the incident. City government officials disagreed with the primary police detective on the case as to whether all of the deleted files were recovered. No criminal charges were filed, and, under an agreement with city officials, the employee was allowed to resign.

These incidents of sabotage were all committed by "insiders:" individuals who were, or previously had been, authorized to use the information systems they eventually employed to perpetrate harm. Insiders pose a substantial threat by virtue of their knowledge of, and access to, employer systems and/or databases.

Prevalence of Incidents of Insider Sabotage

Efforts to estimate how often companies face attacks from within are difficult to make. It has been suggested that insider attacks are under-reported to law enforcement and prosecutors.¹ Reasons for such under-reporting include an insufficient level of damage to warrant prosecution, a lack of evidence or insufficient information to prosecute, and concerns about negative publicity.²

Moreover, statistics vary regarding the prevalence of cases perpetrated by insiders compared to those perpetrated by individuals external to the target organizations.³ The E-Crime Watch Survey™⁴, carried out by the United States Secret Service (Secret Service), the CERT® Program of Carnegie Mellon University's Software Engineering Institute (CERT), and *CSO Magazine* in spring 2004, elicited responses from 500 security and law enforcement executives on issues related to electronic crimes. Among the 70 percent of respondents who were able to identify whether outsiders or insiders were responsible for an e-crime or intrusion committed in 2003, 71% reported that one or more attacks were known or suspected to have come from outsiders compared to 29% from insiders. Respondents identified current or former employees and contractors as the second greatest cyber security threat, preceded only by hackers.

Previous efforts to study insider incidents have focused on groups that could be conveniently sampled or represented more narrow areas of industry. These efforts have included

- workshops to develop a foundation of knowledge on insider threats⁵
- annual surveys of organizations on the number of insider incidents they have experienced in a given year⁶
- in-depth case studies of information technology insiders⁷
- workshops to develop and test a framework for detection of insider threats to U.S. national security⁸

¹ National Research Council, Computer Science and Telecommunications Board, *Summary of Discussions at a Panning Meeting on Cyber-Security and the Insider Threat to Classified Information*, November 2000.

² CSO Magazine, United States Secret Service and CERT® Coordination Center. (2004). 2004 E-Crime Watch Survey. Framingham, MA: CXO Media [Hereafter referred to as 2004 E-Crime Watch Survey].

³ Richardson, R. (2004). Ninth Annual CSI/FBI Computer Crime and Security Survey, Computer Security Institute [Hereafter referred to as CSI Survey]; E-Crime Watch Survey.

⁴ <http://www.cert.org/archive/pdf/2004eCrimeWatchSummary.pdf>

⁵ Anderson, R.H. (1999, August). Research and Development Initiatives Focused on Prevention, Detecting, and Responding to Insider Misuse of Critical Defense Information Systems. Santa Monica, CA: RAND (CF151); Department of Defense (2000). DoD Insider Threat Mitigation: Final Report of the Insider Threat Integrated Process Team. Washington, DC: Author.

⁶ 2004 E-Crime Watch Survey; CSI Survey.

⁷ Shaw, E., Post, J., and Ruby, K. (August 31, 1999). Final Report: Insider Threats to Critical Information Systems: Typology of Perpetrators, Security Vulnerabilities, Recommendations.

⁸ <http://www.mitre.org/news/events/tech04/8.html>.

Collectively, these initiatives have helped to inform the insider threat issue. However, existing gaps in the literature have made it difficult for organizations to develop a more comprehensive understanding of the insider threat and address the issue from an approach that draws upon human resources, corporate security, and information security perspectives. In particular, research to date has not examined the incidents from both behavioral and technical perspectives simultaneously.

The Insider Threat Study

Securing cyberspace has become a national priority. In *The National Strategy to Secure Cyberspace*⁹, the President's Critical Infrastructure Protection Board identified several critical infrastructure sectors¹⁰:

- banking and finance
- information and telecommunications
- transportation
- postal and shipping
- emergency services
- continuity of government
- public health
- food
- energy
- water
- chemical industry and hazardous materials
- agriculture
- defense industrial base

The National Strategy to Secure Cyberspace emphasizes the importance of public-private partnerships in securing these critical infrastructures and improving national cyber security. Similarly, one focus of the Department of Homeland Security is enhancing protection for critical infrastructure and networks by promoting working relationships between the government and private industry. The federal government has acknowledged that these relations are vital because most of America's critical infrastructure is privately held.

Since 2001, the Secret Service and CERT have collaborated on multiple efforts to identify, assess, and manage potential threats to, and vulnerabilities of, data and critical systems. The collaboration represents an effort to augment security and protective practices through two components:

⁹ The National Strategy to Secure Cyberspace. (February 2003). <http://www.whitehouse.gov/pcipb/>.

¹⁰ Please see an updated description of the critical infrastructure sectors described in Homeland Security Presidential Directive 7.

1. Finding ways to identify, assess, and mitigate cyber security threats to data and critical systems that impact physical security or threaten the mission of the organization.
2. Finding ways to identify, assess, and manage individuals who may pose a threat to those data or critical systems

The overall goal of the collaborative effort is to develop information and tools that can help private industry, government, and law enforcement identify cyber security issues that can impact physical or operational security and to assess potential threats to, and vulnerabilities in, data and critical systems.

Consistent with that goal, this report was written for a diverse audience that includes:

- corporate and government managers
- technical staff
- human resources personnel
- security officers
- law enforcement
- policy makers

The Insider Threat Study (ITS), being conducted by the Secret Service National Threat Assessment Center (NTAC) and CERT, is a central component of this multi-year collaboration. This effort was made possible, in part, through funding by the Department of Homeland Security, Office of Science and Technology, which provided financial support for the study in fiscal years 2003 and 2004.

The ITS focuses on the *people* who have access to such information systems and have perpetrated harm using them, and examines each incident from a behavioral and a technical perspective. The project combines the Secret Service's expertise in behavioral and incident analysis with CERT's technical expertise in network systems survivability and security.

The ITS is an extension of earlier studies conducted by both organizations. Previous Secret Service studies have focused on identifying information that is operationally relevant and could help prevent future violent or disruptive incidents. The goal of this earlier research was to find information that could help enhance threat assessment efforts – efforts to identify, assess, and manage the risk of harm an individual may pose, before the individual has an opportunity to engage in violent behavior.

Previous CERT research, sponsored by the Department of Defense, focused on cyber insider threats in the military services and defense agencies. The work is part of an ongoing partnership between CERT and the Defense Personnel Security Research Center (PERSEREC) in response to recommendations in the 2000 *DoD Insider Threat Mitigation* report.¹¹ The focus of that partnership is to identify characteristics of the environment surrounding insider cyber events evaluated for criminal prosecution by

¹¹ www.defenselink.mil/c3i/org/sio/iptreport4_26dbl.doc .

DoD investigative services. The primary use of this information will be to guide future operating, security, and personnel procedures to reduce the threat to critical information systems in the DoD and its contractor community.

The ITS consists of several components:

- an aggregated case-study analysis that provides an in-depth look at insider incidents that have occurred in critical infrastructure sectors between 1996 and 2002 (this report presents the second series of findings from this analysis)
- a review of the prevalence of insider activity across critical infrastructure sectors over a 10-year time frame
- a survey of recent insider activity experienced by a sample of public- and private-sector organizations¹²

The first report from the aggregated case study analysis, *Insider Threat Study: Illicit Cyber Activity in the Banking and Finance Sector*, was published in August 2004.¹³ That report reviewed 23 incidents of insider threat in the banking and finance sector. This report examines insider incidents across critical infrastructure sectors in which the insider's primary goal was to sabotage some aspect of the organization (for example, business operations, information/data files, system/network, and/or reputation) or direct specific harm toward an individual.

Incidents in which the insider's primary motivation was financial gain or theft of information or property are not included in this report. Those cases were included in the banking and finance sector report if the affected organization fell within that sector. Similarly, the next report, which will examine insider activity within the information and telecommunications sector, will include those cases if the organization affected by the incident fell within that sector.

Methodology

The Study Sample

The cases examined in the Insider Threat Study are incidents perpetrated by insiders (current or former employees or contractors) who intentionally exceeded or misused an authorized level of network, system, or data access in a manner that affected the security of the organizations' data, systems, or daily business operations. Incidents included any compromise, manipulation of, unauthorized access to, exceeding authorized access to, tampering with, or disabling of any information system, network, or data. The cases examined also included any in which there was an unauthorized or illegal attempt to view, disclose, retrieve, delete, change, or add information. Only those cases meeting these inclusion criteria, occurring in the United States, and in which the affected organization fell within a critical infrastructure sector were included in the study.

¹² 2004 E-Crime Watch Survey.

¹³ Available on-line at <http://www.cert.org/archive/pdf/bankfin040820.pdf> and http://www.secretservice.gov/ntac_its.shtml

Cases were identified through public reporting or as a computer fraud case investigated by the Secret Service.¹⁴ Public reporting included references in various media outlets (found through searches on Lexis-Nexis news databases) and criminal justice databases (found through searches on Lexis court databases).

The cases studied here may or may not be representative of cases not mentioned in media, court, or Secret Service databases. As noted, organizations may be reluctant to expose these incidents, even to law enforcement. This report and others from the study will articulate only what we found among these known cases. This limits the ability to generalize the study findings and underscores the difficulty other researchers have faced in trying to better understand the insider threat. This limitation does not, however, diminish the value in analyzing these incidents. This study provides insight into actual criminal acts committed by insiders. This insight may be useful to those in the sectors charged with protecting their critical assets as they begin to examine ways of improving their defense against insider attacks.

Research and Analysis

The ITS adapted methods used in previous research performed by the Secret Service and CERT to conduct in-depth examinations of network, system, and data compromises and other insider activity. Researchers focused primarily on tracing insider incidents from the initial harm backward in time to when the idea of committing the incident first occurred to the insider. In tracing the incidents backward, researchers tried to identify the behaviors and communications in which the insiders engaged – both online and offline – prior to and including the insiders' harmful activities.

For each case examined in the study, researchers from the Secret Service and from CERT reviewed primary source material on the case, including investigative reports, court records, and other materials, as well as secondary source material from news articles.¹⁵ Researchers also conducted supplemental interviews with case investigators and organization representatives.¹⁶ Researchers used the information gleaned from these sources to complete several hundred questions about the insider and the behavioral and technical aspects of the incident. The questions were organized around the following major topic areas:

1. components of the incident
2. detection of the incident and identification of the insider
3. pre-incident planning and communication
4. nature of harm to the organization
5. law enforcement and organizational response
6. characteristics of the insider and the organization

¹⁴ Examples of computer fraud cases include cases where an individual(s) fraudulently obtains a credit card issuer's records via a computer; places a virus, Trojan horse, or worm on, or conducts a denial-of-service attack against a computer; or obtains unauthorized access to a computer system by using a password.

¹⁵ Appendix E provides a list of Secret Service and CERT personnel who reviewed cases for the study.

¹⁶ For this report, researchers interviewed representatives from 25 companies and 39 law enforcement or prosecutorial agencies, as well as one of the insiders whose incident was reviewed for the study.

7. insider background and history
8. insider technical expertise and interests

Overview of the ITS Study Findings

The ITS' examination of incidents of insider sabotage across critical infrastructure sectors found that most of the insiders who committed acts of sabotage were former employees who had held a technical positions with the targeted organizations. The majority of the incidents examined under the ITS were perpetrated against private sector organizations.

Insider activities caused organizations financial losses, negative impacts to business operations and damage to reputation. As a result of their involvement in the incidents reviewed for this study, almost all of the insiders were charged with criminal offenses. The majority of these charges were based on violations of federal law.

Among the key findings of the ITS study of insider sabotage across critical infrastructure sectors are the following:

- A negative work-related event triggered most insiders' actions.
- Most of the insiders had acted out in a concerning manner in the workplace.
- The majority of insiders planned their activities in advance.
- When hired, the majority of insiders were granted system administrator or privileged access, but less than half of all of the insiders had authorized access at the time of the incident.
- Insiders used unsophisticated methods for exploiting systemic vulnerabilities in applications, processes, and/or procedures, but relatively sophisticated attack tools were also employed.
- The majority of insiders compromised computer accounts, created unauthorized backdoor accounts, or used shared accounts in their attacks.
- Remote access was used to carry out the majority of the attacks.
- The majority of the insider attacks were only detected once there was a noticeable irregularity in the information system or a system became unavailable.
- Insider activities caused organizations financial losses, negative impacts to their business operations and damage to their reputations.

Organization of the Report

The remainder of this report is organized into four sections. Section 2, "Characteristics of Insider Sabotage Across Critical Infrastructure Sectors," provides information on the individuals who carried out the insider attacks and the organizations that were the target of those attacks. This section also describes the consequences of the insider attack on the insider.

Section 3, "Key Findings of the Insider Threat Study of Sabotage Across Critical Infrastructure Sectors," presents the key findings of the study. In Section 4, "Implications of the Key Findings for the Prevention of Insider Sabotage," the report discusses the key findings that appear to have implications for the development of strategies to prevent incidents of insider sabotage. Section 5, "Conclusion: Reflections on the

Findings for the Prevention of Insider Sabotage,” offers some concluding thoughts on the study’s examination of insider sabotage across the critical infrastructure sectors.

Finally, this report includes four appendices. Appendix A provides tables containing information on the date and location of the incidents examined by this study. Appendix B describes the size of the targeted organizations. Appendix C provides additional case examples. Appendix D contains a glossary of technical terms used in this report. The first use of each term included in the glossary is indicated in the text of the report with an underscore. Appendix E identifies and acknowledges the efforts of individuals who worked on this project.

SECTION 2: CHARACTERISTICS OF INSIDER SABOTAGE ACROSS CRITICAL INFRASTRUCTURE SECTORS

This report of the ITS examines cases in which the primary goal of the insider was to sabotage some aspect of an organization or direct specific harm toward an individual(s). Forty-nine incidents that occurred across the critical infrastructure sectors between 1996 and 2002 were studied.

Information on each research question that was investigated for the study was not available for all cases. Thus, percentages in this report are based on the total number of cases for which information was available for a given research question. For example, if information on a particular research question was available for only 46 of the 49 cases, then that particular statistic is based on a total of 46 cases. Instances in which information for a particular research question was not available for more than 10% of the cases are indicated by a footnote.

Analysis of the study findings identified the following general characteristics of the insiders, the organizations they targeted, and the consequences that befell the insiders:

Insider Characteristics

The majority of the insiders were former employees.

- At the time of the incident, 59% of the insiders were former employees or contractors of the affected organizations and 41% were current employees or contractors.
- The former employees or contractors left their positions for a variety of reasons. These included the insiders being fired (48%), resigning (38%), and being laid off (7%).

Most insiders were either previously or currently employed full-time in a technical position within the organization.

- Most of the insiders (77%) were full-time employees of the affected organizations, either before or during the incidents. Eight percent of the insiders worked part-time, and an additional 8% had been hired as contractors or consultants. Two (4%) of the insiders worked as temporary employees, and one (2%) was hired as a subcontractor.
- Eighty-six percent of the insiders were employed in technical positions, which included system administrators (38%), programmers (21%), engineers (14%), and IT specialists (14%). Of the insiders not holding technical positions, 10% were employed in a professional position, which included, among others, insiders employed as editors, managers, and auditors. An additional two insiders (4%) worked in service positions, both of whom worked as customer service representatives.

Insiders were demographically varied with regard to age, racial and ethnic background, gender, and marital status.

- The insiders ranged in age from 17 to 60 years (mean age = 32 years)¹⁷ and represented a variety of racial and ethnic backgrounds.
- Ninety-six percent of the insiders were male.
- Forty-nine percent of the insiders were married at the time of the incident, while 45% were single, having never married, and 4% were divorced.

Just under one-third of the insiders had an arrest history.

- Thirty percent of the insiders had been arrested previously, including arrests for violent offenses (18%), alcohol or drug related offenses (11%), and non-financial/fraud related theft offenses (11%).

Organization Characteristics

The incidents affected organizations in the following critical infrastructure sectors:

- banking and finance (8%)
- continuity of government (16%)
- defense industrial base (2%)
- food (4%)
- information and telecommunications (63%)
- postal and shipping (2%)
- public health (4%)

In all, 82% of the affected organizations were in private industry, while 16% were government entities. Sixty-three percent of the organizations engaged in domestic activity only, 2% engaged in international activity only, and 35% engaged in activity both domestically and internationally. Appendix B provides a table on the affected organizations' sizes by number of employees.

Consequences to Insiders

Almost all insiders were charged criminally, the majority of which were based on federal law.

- Ninety-percent of the insiders faced formal criminal charges.
- Of those insiders criminally charged, 61% faced federal charges, 36% faced state charges, and 2% faced federal and state charges.
- The most common charges were violations of 18 USC 1030 (65%), a computer related state criminal statute (19%) and/or a non-computer related state criminal statute (5%).

The majority of insiders were found guilty of violating a criminal statute.

- Of those insiders who faced criminal charges, 83% were found guilty by trial or by plea and 5% entered a plea of no contest.

¹⁷ Data were only available for 43 of the 49 insiders studied.

- Seventy-one percent of these insiders were sentenced to probation or supervised release, 59% were ordered to pay restitution, 42% were sentenced to a period of incarceration, 7% were ordered to complete community service, and 5% were sentenced to home detention.
- The length of the probation periods ranged from 12 to 180 months, with an average probation period of 36 months.
- Insiders were ordered to pay restitution in far varying amounts, ranging from \$100 to \$2 million.
- The length of the incarceration periods ranged from 2 to 41 months, with an average incarceration period of 12 months.
- Most of the insiders (85%) did not consider the severity of the consequences that could result from their actions; however, in almost all cases (98%) the organizations responded to the incident by taking action external to the organization.

Almost all of the incidents resulted in law enforcement notification (96%) or the filing of a civil lawsuit (4%). In those cases in which law enforcement was notified, organizations contacted

- local police departments and/or local prosecutor's offices (39%)
- state law enforcement agencies and/or state prosecutor's offices (4%)
- federal law enforcement agencies and/or U.S. attorney's offices (37%)
- two or more of these groups (e.g., local and state police departments) (20%)

Section 3: Key Findings of the Insider Threat Study of Sabotage Across Critical Infrastructure Sectors

The key findings of the study of incidents of insider sabotage across critical infrastructure sectors are presented under five categories:

- The Insider's Motive
- Pre-attack Behavior and Planning
- Advancing the Attack
- Detecting the Attack
- Consequences for Targeted Organizations

The Insider's Motive

After noticing its graphic artist's aptitude with computers and computer programming, a company asked him to create its Internet website. A few months later, the company reprimanded the employee for absenteeism, and the company president notified the employee that the company planned to suspend him. Later that day, the employee remotely accessed the company's network, deleted information, and added other text and images to the company's website. The insider later admitted to law enforcement that he committed the offense because he was angry at the company for suspending him.

Key Findings

- A negative work-related event triggered most insiders' actions.
- Most insiders held a work-related grievance prior to the incident.
- The most frequently reported motive was revenge.

Supporting Data

In 92% of the cases, a specific event or a series of events triggered the insiders' actions. These events included, among others, various work-related events to include employment termination (47%), dispute with a current or former employer (20%), and employment related demotion or transfer (13%).

Eighty-five percent of the insiders held a grievance prior to the incident, and in 92% of these cases, the insider's grievance was work-related (including grievances against current and/or former employers, supervisors, and coworkers). Fifty-seven percent of the insiders were perceived by others as disgruntled employees.

Eighty-four percent of the incidents were motivated at least in part by a desire to seek revenge. Insiders also were motivated to

- address a grievance or issue held by the insider (41%)
- garner respect or acknowledgement (12%)
- address dissatisfaction with company policies (12%)
- address dissatisfaction with company culture (12%)

In 57% of the cases, insiders had more than one motive for carrying out their actions.

Pre-attack Behavior and Planning

In one case, an insider had become dissatisfied with his job installing software and hardware on the company's computers and with providing technical support to its employees. He emailed his employer the following message:

"I have no intention of taking ownership of modem troubleshooting. If you insist, so be it but I can assure you the job will be completed with very little effort and no attention to detail."

He also shared his negative feelings about his employer with a coworker in an email, stating "I hope [the company owner]'s not going to be coming to lunch tomorrow. I might wind up pummeling [him]." Eventually, the insider's dissatisfaction led to his resignation.

In spite of his negative communications, the company owner permitted the insider to retain email access as a paying customer following his resignation. Several weeks after he resigned, the insider used bogus accounts he had created to change all of the company's administrative passwords, alter the computer's registry, delete the entire billing system, and delete two internal databases. Prior to these activities, the insider had expressed his intent to harm the company in emails he sent to a relative and a former coworker at the company.

Key Findings

- Most of the insiders had acted out in a concerning manner in the workplace.
- The majority of insiders planned their activities in advance.
- Others had information about the insiders' intentions, plans, and/or ongoing activities.
- A majority of the insiders communicated negative sentiments to others, and in some cases they communicated direct threats of harm.

Supporting Data

Eighty percent of the insiders came to the attention of someone for behavior of concern or behavior that was inappropriate prior to the incidents. These behaviors included, among others, tardiness, truancy, arguments with coworkers, and poor job performance. Of those cases,

- In 97%, the insider's behavior came to the attention of others in the workplace, including supervisors, coworkers, and/or subordinates.
- In 74%, the insider's behavior had consequences.

Thirty-one percent of the insiders studied had a record of disciplinary actions within the organization prior to the incident.¹⁸

Sixty-two percent of the insiders developed plans to harm the organization. Forty-seven percent of the cases involved overt behaviors in preparation for the incident, such as stealing copies of back-ups. In 27% of the cases, the overt behaviors were technical actions taken to set up the attack, including constructing and testing a logic bomb on the network, centralizing critical assets and sabotaging backups, or installing backdoors.

In 37% of the cases, the insider's planning activity was noticeable. In those cases, the insiders' planning activities were noticeable online (67%), offline (11%), and, in some cases, both online and offline (22%).

In 31% of the cases, others had information about the insiders' plans, intentions, and/or activities. In these cases

- coworkers had information (64%)
- friends had information (21%)
- family members had information (14%)
- someone involved with the incident had information (14%)

Fifty-eight percent of the insiders communicated to others negative feelings, grievances, and/or an interest in causing harm. Of these, in 39% of the cases the insiders communicated their negative sentiments about the affected organization or individual directly to that organization or individual and in 69% they communicated these sentiments to someone else. Insiders communicated these negative sentiments in various ways, including verbally (92%) and via email (12%).

In 20% of the cases, the insider made a direct threat regarding harming the organization or an individual. In all of these cases, the insiders communicated these threats about their targets to others who were not affected by the incident. In only one case did the insider threaten his target directly. In 78% of these cases, the insider made verbal threats, although they may have used other means of communication as well.

Advancing the Attack

A system administrator was terminated and his account immediately disabled. However, his organization overlooked disabling his remote access to the organization's network through the firewall, and also failed to change the root password. These oversights enabled the insider, after business hours, to sabotage the system, making it inaccessible for three days. If his remote access had been disabled, and/or the root account password changed, his actions might have been prevented.

¹⁸ Data were only available for 42 out of the 49 cases.

Key Findings

- When hired, the majority of insiders were granted system administrator or privileged access, but less than half of all of the insiders had authorized access at the time of the incident.
- Insiders exploited systemic vulnerabilities in applications, processes, and/or procedures, but relatively sophisticated attack tools were also employed.
- The majority of insiders compromised computer accounts, created unauthorized backdoor accounts, or used shared accounts in their attacks.
- Remote access was used to carry out the majority of the attacks.
- The majority of attacks took place outside normal working hours.

Supporting Data

Fifty-seven percent of the insiders were granted system administrator access upon hire. Of these insiders, 85% no longer legitimately retained that level of access at the time of the incident. These insiders

- had been terminated or had resigned and their access had been disabled (38%)
- had been terminated or had resigned but their access was not disabled (27%)
- maintained their employment with the organization and remained authorized users, but had their level of access reduced (12%)
- had been terminated or had resigned but were permitted to retain limited system access (8%)

An additional 33% of the insiders were hired as privileged users, but 60% of them no longer retained authorized privileged access at the time of the incident. These insiders

- had been terminated or had resigned and their access had been disabled (20%)
- had been terminated or had resigned, but their access was not disabled (33%)
- had their access increased to administrator/root access (7%)

Only 43% of the insiders had authorized access to the system/network at the time of the incident. In 31% of the cases, the insider's access had been disabled by his or her employer prior to the attack. In 26% of cases, insiders were able to attack after termination because their employers did not disable their access.

In 57% of the cases, the insiders exploited or attempted to exploit systemic vulnerabilities in applications, processes, and/or procedures (e.g., business rule checks, authorized overrides).

In 61% of the cases, the insider's actions were limited to relatively unsophisticated methods of attack. These methods included user commands, information exchanges, and exploitation of physical security vulnerabilities. The remaining 39% of the insiders used one or more relatively sophisticated methods of attack, which included

- a script or program
- an autonomous agent
- a toolkit
- flooding

- probing
- scanning
- spoofing

In 60% of the cases, the insider compromised an account to carry out the attack. These compromises included the use of another's username and password (33%) or the use of an unauthorized account created by the insider (20%). In 92% of these cases, there were no indications of suspicious activity related to the account before the initial incident.

The insiders also used shared accounts to carry out their activities, including group accounts, for example, system administrator or database administrator (DBA) accounts (15%), and company accounts (13%). In 30% of the cases, the insiders used their own usernames and passwords. In 13% of the cases, the insiders used accounts in more than one of the above categories to carry out the attack.

In 87% of the cases, the victim organizations permitted employees remote access.¹⁹ In 56% of the cases, the attacks were conducted solely via remote access, 35% took place only from within the workplace, and 8% took place both from within the workplace and remotely.

Fifty eight percent of the attacks took place outside of normal working hours or on weekends or holidays, and 42% took place during normal working hours.

Detecting the Attack

A former consultant hired by an organization to set up its network was still able to log in to the admin account following termination of his contract. The new system administrator first detected the attack when he noticed probing of the organization's network, and suspected a potential security problem. Although he took steps to further secure the network, the former contractor was still able to log in, install a remote administration tool, and use the information gathered to compromise additional accounts. The insider was finally identified several weeks later when law enforcement used forensic examinations of the organization's server and system log files to trace his actions to remote VPN connections, then using ISP records, to the insider's home computer.

Key Findings

- The majority of the insider attacks were only detected once there was a noticeable irregularity in the information system or a system became unavailable.
- System logs were the most prevalent means by which the insider was identified.
- Insiders took steps to conceal their identities and their activities.
- Most of the incidents were detected by non-security personnel.

¹⁹ Data were only available for 39 of the 49 incidents studied.

- The majority of attacks were accomplished using company computer equipment.
- Forensic examinations were used to identify the insider and gather evidence in many of the cases.

Supporting Data

Sixty-three percent of the incidents were detected because of an irregularity in the information or system, 42% due to system failure, and 10% due to both an irregularity and a system failure.

In general, 75% of the insiders were identified through manual procedures only, and 19% were identified using a combination of automated and manual procedures. The various mechanisms used to identify the perpetrators included

- system logs (70%)
- insider's own source IP address (33%)
- phone records (28%)
- username (24%)
- auditing procedures (13%)

In those cases in which system logs were used to identify the insider as the perpetrator, the following logs were used

- remote access logs (73%)
- file access logs (37%)
- system file change logs (37%)
- database/application logs (30%)
- email logs (13%)

In 31% of the cases, multiple types of logs were used to identify the insider.

In 76% of the cases, the insiders took steps to conceal their identities (31%), actions (12%), or both (33%). These insiders engaged in activities to conceal their identities and/or actions, such as using technology to delete or modify records of the incidents, creating unauthorized accounts or backdoors, or spoofing the source IP address.

Seventy-one percent of the incidents were detected by individuals who were not part of the organization's security staff.²⁰ These individuals included customers (24%), supervisors (20%), and other non-security personnel (51%). The incidents that were detected by security personnel were discovered by a range of security professionals including information technology (IT) security staff or system administrators (22%), and staff responsible for information systems/data (7%).

Thirteen percent of the incidents were detected by non-technical means, such as observation by a coworker or notification by the insider.

²⁰ Note that in most cases the insiders were detected by multiple people.

In 55% of the cases, the insiders used company computer equipment to access information for the attack. In 56% of those cases, the insiders used computer equipment assigned to another employee, to a department, or to an employee group.

In 41% of the cases, the insider was identified through a forensic examination of the targeted network, data, or system, and in 24% of the cases that identity was confirmed through a forensic examination of the insider’s home equipment.

Consequences for Targeted Organizations

An insider had extensive control over the source code of a critical application used by the organization. As lead developer of the software, he made sure that he possessed the only copy of the source code. There were no backups, and very little documentation existed. Following a demotion in both position and pay, the insider “wiped” the hard drive of his company-provided laptop. In doing so, he deleted the only copy of the source code the organization possessed. It took several months to recover the source code from the insider, during which time the organization was unable to update the software.

Key Findings

- Insider activities caused organizations financial losses, negative impacts to their business operations and damage to their reputations.
- Incidents affected the organizations’ data, systems/networks, and components.
- Various aspects of organizations were targeted for sabotage by the insider.
- In addition to harming the organizations, the insiders caused harm to specific individuals.

Supporting Data

Eighty-one percent of the organizations experienced a negative financial impact as a result of the insiders’ activities. The losses ranged from a reported low of \$500 to a reported high of “tens of millions of dollars.” The chart below represents the percentage of organizations experiencing financial losses within broad categories.

Percentage of Organizations	Financial Loss
42	\$1 - \$20,000
9	\$20,001 - \$50,000
11	\$50,001 - \$100,000
2	\$100,001 - \$200,000
7	\$200,001 - \$300,000
9	\$1,000,001 - \$5,000,000
2	Greater than \$10,000,000

Seventy-five percent of the organizations experienced some impact on their business operations. These impacts included:

- severed communication with affected organizations due to shut-down networks, routers, servers, or dial-up access
- blocked sales due to blocked sales applications or deleted sales records
- blocked customer contact due to modified customer passwords
- damaged or destroyed critical information assets, such as proprietary software, data, computing systems, and storage media necessary to the organization's ability to contract work, produce product, or develop new products
- damaged supervisory integrity, including exposed personal or private communications embarrassing to a supervisor

In addition, 28% of the organizations experienced a negative impact to their reputations.

Ninety four percent of the incidents affected the integrity, confidentiality, and/or availability of the organizations' data, including modification and/or deletion of data (88%); reading, copying, and/or stealing data (27%); corrupting data (65%); and engaging in unauthorized disclosure of data (20%).

Fifty seven percent of the incidents affected the integrity, confidentiality, availability, or authentication of the organizations' information systems/networks. These actions included denial of service in 73% of the cases and unauthorized increase in system/network access in 27% of the cases. Forty seven percent of the incidents affected the organizations' network, components, or external connectivity (47%).

For inclusion in this report, the insiders' primary goals for their activities had to be a desire to sabotage some aspect of the organization and/or harm a specific individual. Insiders sought to

- sabotage information systems/networks (51%)
- sabotage business operations (51%)
- sabotage information and/or data (49%)
- harm specific individuals (35%)
- sabotage the organization's reputation (25%)

Thirty-five percent of the cases involved harm to a specific individual(s). Examples of such cases include ones in which the insider maligned the reputation of a company owner in email communications, and threatened company officers while also posting social security numbers associated with the company on the Internet.

Section 4: Implications of the KEY Findings for the Prevention of Insider Sabotage

In this section, the authors discuss the implications that key findings of the study of insider sabotage across critical infrastructure sectors may have for developing strategies for the prevention of such incidents.

As in the previous section, the discussion of implications is organized under five categories: The Insider's Motive; Pre-attack Behavior and Planning; Advancing the Attack; Detecting the Attack; and Consequences for Targeted Organizations. For the reader's convenience, the findings for each category are repeated in this section. The findings then are followed by brief descriptions of the relevant implications and a commentary that describes the implications in greater detail.

The Insider's Motive

Key Findings

- A negative work-related event triggered most insiders' actions.
- Most insiders held a work-related grievance prior to the incident.
- The most frequently reported motive was revenge.

Implications

- Management attention is needed for employees who experience negative work-related events.
- Establish formal grievance procedures and additional forums for employees to voice concerns.

Commentary

Management attention is needed for employees who experience negative work-related events. The findings suggest organizations may benefit from paying attention to employees who experience negative employment-related events, including employment termination, demotion, or conflicts with coworkers and management. Developing organizational policies that outline steps management may take when an employee has been negatively impacted at work is an important step. Knowledge that such policies exist may diminish an employee's concern that he or she has been selected or targeted for attention following a negative event at work. Rather the employee would understand that the organization's response was part of an established protocol.

Such policies may provide an opportunity for constructive discourse between the employee, management and/or coworkers in an effort to decrease the employee's dissatisfaction or other negative affect. These discussions can explore how the event(s) has affected the employee and what he or she perceives as options in reacting to and responding to the event, including information that may provide clues as to whether the employee intends to harm the organization. Reaching out to these employees and

exploring how these events have affected them and their perceived options may deter some employees from doing harm.

Establish formal grievance procedures and additional forums for employees to voice concerns. The findings suggest companies may want to provide formal grievance procedures and additional forums in which employees can voice their concerns. Formally addressing an employee's grievance may diminish the employee's desire to address the grievance in another manner that harms the organization. Additionally, providing an opportunity for the employee to be heard and demonstrating that his or her concerns will be considered may reduce an employee's motive to harm the organization.

Pre-attack Behavior and Planning

Key Findings

- Most of the insiders had acted out in a concerning manner in the workplace.
- The majority of insiders planned their activities in advance.
- Others had information about the insiders' intentions, plans, and/or ongoing activities.
- A majority of the insiders communicated negative sentiments to others, and in some cases they communicated direct threats of harm.

Implications

- Establish a formal process for reporting and sharing information.
- Document reports of problematic behavior and develop procedures to respond to such reports.

Commentary

Establish a formal process for reporting and sharing information. The findings suggest organizations have opportunities to detect problem and planning behavior before harm occurs. Developing a formal process for the reporting of such behavior in the workplace is important, including the consideration of whether a mechanism for anonymous reporting should be provided. Employees should be informed of the process and encouraged to avail themselves of the opportunity to report suspicious or inappropriate behavior.

In addition, appropriate procedures should be developed that allow various departments within the organization to share information regarding problematic or suspicious employee behavior. Representatives from management, security, human resources, and legal counsel would benefit from sharing information they have learned regarding particular behavior of concern.

Document reports of problematic behavior and develop procedures to respond to such reports. Employees may be more likely to report suspicious or problematic behavior if they believe such information will be documented and appropriate action will be taken. An official documentation procedure should be developed to ensure that

relevant information is gathered in a standard and thorough manner. Such documentation will serve to create a recorded history of various concerning behaviors to both track the course of behavior over time and also provide a recorded basis for disciplinary action, if such action is warranted.

Additionally, organizations may benefit from developing policies that govern who will respond to such reports and how they will be handled. An established procedure that is followed in a responsible and professional manner will lay the groundwork for fair and responsive action by the organization. This procedure will also inform employees that the organization has developed specific ways to address concerning behavior in the workplace, which may encourage their reporting of such behavior.

Advancing the Attack

Key Findings

- When hired, the majority of insiders were granted system administrator or privileged access, but less than half of all of the insiders had authorized access at the time of the incident.
- Insiders used unsophisticated methods for exploiting systemic vulnerabilities in applications, processes, and/or procedures, but relatively sophisticated attack tools were also employed.
- The majority of insiders compromised computer accounts, created unauthorized backdoor accounts, or used shared accounts in their attacks.
- Remote access was used to carry out the majority of the attacks.
- The majority of attacks took place outside normal working hours.
- The majority of attacks were accomplished using company computer equipment.

Implications

- Formal policies and procedures for disabling access upon an employee's termination or resignation should be established and followed.
- Procedural and technical controls should be established for system administrator and privileged system functions.
- Attacks that are not technically sophisticated require policies, practices, and technology for prevention or early detection.
- Comprehensive password policies should be established and communicated to employees.
- Comprehensive computer account management policies and practices should be implemented.
- Characterization of hardware and software, and configuration management practices should be followed for detection of logic bombs and other malicious code.
- Layered security for remote access should be considered.

Commentary

Formal policies and procedures for disabling access upon an employee's termination or resignation should be established and followed. Some former employees were able to attack because their access to the organizations' systems or networks was not disabled. Organizations may prevent such attacks by establishing formal policies and procedures for disabling access to the network upon an employee's departure. The practice of deactivating computer accounts, system authorizations, and remote access should occur immediately when employment, consulting, or contracting agreements are terminated, regardless of the reason for the termination.

Ad hoc or outdated termination policies and procedures also increase the risk of access control gaps inadvertently left open for terminated personnel. Some of the insiders were able to attack even though their access had been disabled upon firing or resignation. Insider knowledge of the organization's policies, procedures, technology, and security measures provided them an enhanced ability to attack the organization even though their access was disabled. For example, one insider was aware of a system vulnerability, and even reported it to management. He was able to exploit the vulnerability after his termination as management never addressed it. Formal policies and procedures that are periodically reviewed and modified if necessary provide consistent and comprehensive measures to reduce that risk.

Organizations should also consider the risk posed by former employees who retain authorized access to the company's system or network. In two cases, the organizations deliberately permitted the insiders to retain limited system access following their departure from the companies. In both cases, the insiders used this access to carry out their attacks.

Disabling Access to Shared Accounts

Shared, privileged accounts provide increased access to an organization's network, systems, applications, or data. Employee and contractor termination procedures must include disabling access to *all* accounts to which that user had access, including shared accounts. Although information was only available for 26 cases, of those, 65% involved victim organizations that did not change group passwords for shared accounts, including system administrator accounts, upon an employee's departure. This oversight exposes the organization to a vulnerability easily exploited by a former employee or contractor possessing the shared password. It is critical that all group passwords be changed immediately upon termination of any user previously authorized to have the password. These include group passwords for remote access, and company, customer, database administrator, application, and system administrator accounts.

Disabling Remote Access

Disabling remote access also is a critical part of the employee termination process that is sometimes overlooked. It is essential that employee termination procedures include

- disabling remote access accounts (Virtual Private Network, dial-up accounts, etc.)
- disabling firewall access

- disconnecting any remote connections previously opened by the employee

Additional Termination Tasks

Some specific items that should be added to the termination procedures, based on information gleaned from individual cases in the study, include:

- Reminding all coworkers of a departed employee to change their passwords if there is the slightest chance they may have shared their passwords with that employee. Organizations should recognize that sometimes, even in violation of policy, employees share their passwords with their coworkers.
- Terminating physical access as it can be used to facilitate access to the organization's network. Organizations should keep in mind that coworkers might provide physical access to a former employee if they are unaware of the employee's termination or resignation.
- Disabling access for temporary employees and contractors should be handled as thoroughly as that of permanent employees.

Procedural and technical controls should be established for system administrator and privileged system functions. The power of system administrators should not be underestimated: almost all of the insiders in this study were granted system administrator or privileged access when hired. Because of their elevated access level, they have the ability to cause catastrophic system failure or gradually compromise system or data confidentiality, integrity, or availability over time.

Forty percent of the attacks used technically sophisticated tools or methods commonly associated with the hacker community requiring elevated access levels, such as scripts or programs (such as logic bombs), autonomous agents, and toolkits. In addition, some insiders planted malicious code on the system while still employed with the organization, setting the execution to occur at a later time. If technical controls are used to enforce separation of duties and a two person rule for system administration functions, then release of such malicious code requires collaboration among multiple employees with the necessary access levels; thus, reducing the risk of release of such code by a single insider.

However, separation of duties for system administrators and privileged users also requires that responsibilities be shared, but with individual accountability. Procedural and technical controls must be enforced to ensure that all actions, including the use of shared privileged or system administrator accounts, can be traced back to an individual user.

Attacks that are not technically sophisticated require policies, practices, and technology for prevention or early detection. Technical sophistication was not used in 61% of the cases; those cases involved only simple attack methods (legitimate user commands, information exchanges, or physical attacks). Attacks were successful primarily due to systemic vulnerabilities in technology and/or policies, processes or procedures, such as

- coarse access control restrictions²¹
- sloppy accounting procedures
- infrequent or non-existent monitoring procedures (including transaction verification)
- insufficient physical access controls

Cases in which only simple attack methods were used emphasize the need for fine-grained access controls, two-person control of critical system and data modifications, and integrity checking to track system modifications after the fact. Fine-grained access controls prohibit personnel in one role from accessing functions authorized for personnel in other roles, limiting the damage one insider could inflict if he or she became disgruntled or otherwise decided to exploit the organization for personal gain. Similarly, integrity checking flags critical or suspicious actions for investigation, which may serve to prevent harm or mitigate additional damage that may occur if the activity is not stopped.

Other non-technical techniques, social engineering and physical sabotage to harm electronic assets, were employed in individual cases and should be considered by organizations in strategies to prevent insider activity. The risk of social engineering can be mitigated by employee security awareness training, and physical sabotage resulting in harm to the organization's network, components, systems or data might be prevented via physical security methods.

Comprehensive password policies should be established and communicated to employees. In one third of the cases, the insider utilized another employee's account to commit the attack. Compromised accounts included not only computer log-in accounts, but also remote access/VPN accounts, and involved the accounts of various personnel to include those of

- coworkers
- supervisors
- chief operating officers (COO)
- chief financial officers (CFO)
- senior management
- system administrators
- database administrators
- customers

One insider utilized an expired customer account, some obtained passwords for the accounts of a coworker, supervisor, or manager, and one intimidated a subordinate into revealing her password. In some cases, insiders used a root, system administrator, or other group account. Some insiders took advantage of physical access to a computer to

²¹ The coarseness of access control restrictions is a relative measure. For example, user group level access control is coarser than individual user level access control. Another example, unconstrained remote access is coarser than remote access limited to non-critical functions. In general, coarse access controls are coarser than fine-grained access controls. See the definition of "access control" in the glossary for more information.

compromise an account, suggesting the use of password-protected screen savers to prevent account compromises by insiders.

The purpose of passwords is to protect computer accounts from unauthorized use; these findings suggest that careful attention to password policies is warranted. Formal policies should prohibit sharing of passwords and require the selection of strong passwords. A procedure should be instituted whereby attempts by anyone to obtain passwords should be reported and investigated immediately. In addition, organizations should educate all employees about the importance of protecting all of their passwords.

Comprehensive computer account management policies and practices should be implemented. Account management policies and practices serve many purposes, including:

- detection of backdoor accounts
- tracking and management of access to shared accounts
- association of individual users with shared account actions
- tracking of every account to which every user has access
- disabling of account access

Several attacks were carried out using backdoor accounts and shared accounts, which made identification of the insider difficult. Some insiders created unauthorized, backdoor accounts that they later used to commit their attacks. These attacks, planned and executed by users with system administrator privileges, were harmful and difficult to trace. Malicious acts can usually be associated with the computer accounts used to commit them, however, backdoor accounts, because unauthorized and not tied to a legitimate user, do not provide clear, traceable paths back to those persons using them.

Insiders also used shared accounts to carry out their attacks. These actions can be difficult to trace if shared accounts are not implemented with technical controls to associate account usage with individual users.

The findings suggest that *all* computer accounts should be carefully tracked to ensure that all access can be disabled for terminated employees. Accounts should be audited immediately prior to and following termination of any employee or contractor. These audits should check for unneeded or unauthorized accounts. In addition, periodic account audits should be conducted for all computer accounts, including

- remote access accounts
- login accounts
- DBA accounts
- other application, customer, and company accounts

The audit should include verifications by all account owners so that unauthorized accounts are discovered before they are used for malicious activity.

Characterization of hardware and software, and configuration management practices can facilitate detection of logic bombs and other malicious code.

Several insiders used scripts (including logic bombs) and/or autonomous agents to delete and corrupt files on which their companies' critical operations depended. Such technical attacks are stealthy and can be quite difficult to detect head of time.²² However, organizations can institute practices for detection of malicious code that might be inserted onto their network by either insiders or outsiders. Characterization of software, hardware, and information assets is one of the most effective methods for detecting such code. Organizations should establish a trusted baseline for each machine on the network, store it in a secure location, and periodically compare the current 'footprint' or configuration of each machine with that baseline. Unexpected files or changes to files revealed by this comparison can be analyzed to determine whether they are legitimate changes to the baseline or rather, malicious code.²³ The release process for any new files must include an update to the trusted baseline. Technical controls forcing multiple levels of authorization also should be required for such updates.²⁴

Configuration and change management procedures and software do not provide a foolproof solution.

They require strict monitoring, follow-through, and separation of duties. If these procedures are implemented effectively, however, they can detect otherwise obscure system changes or releases of new code. A comparison of system modifications to previous versions can highlight the insertion of malicious code. In addition, a configuration management system that enforces separation of duties, requiring separate authorization for release of changes to critical files, can prevent insiders from releasing malicious code in the first place.

Layered security for remote access should be considered. Another predominant means of insider attack was remote access. Remote access was widely available to the insiders in these cases. While there are benefits to remote access, organizations should carefully consider the security implications of providing remote access to *critical* data, processes, or information systems.

A layered security approach is suggested that allows remote access to email and non-critical data, but restricts access to the most critical data and information systems to only those employees physically located inside the workplace. In addition, remote

²² Of the 28 cases in which information was available on insider coding/scripting ability, 86% possessed coding/scripting skills; 71% had experience in Windows/DOS or UNIX scripting; and 25% could program in C++.

²³ Corporate Information Security Working Group Report of the Best Practices and Metrics Teams - Subcommittee on Technology, Information Policy, Intergovernmental Relations and the Census, Government Reform Committee, United States House of Representatives November 17, 2004 (Revised January 10, 2005)
http://www.educause.edu/content.asp?page_id=666&ID=CSD3661&bhcp=1&bhav=6.00&bhsh=1024&bhsw=1280&bhiw=1278&bhih=848&bhq=1

²⁴ For a complete description of security improvement practices for characterizing software, hardware, and information assets, see <http://www.cert.org/security-improvement/practices/p091.html>.

system administrator access should be limited to the smallest group possible, and all such access closely logged and monitored on a regular basis. Information such as login account, date/time connected and disconnected, and IP address should be logged for all remote logins. It is also 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.

Detecting the Attack

Key Findings

- The majority of the insider attacks were only detected once there was a noticeable irregularity in the information system or a system became unavailable.
- System logs were the most prevalent means by which the insider was identified.
- Insiders took steps to conceal their identities and their activities.
- Most of the incidents were detected by non-security personnel.
- Forensic examinations were used to identify the insider and gather evidence in many of the cases.

Implications

- Logging and monitoring can be used for proactive detection rather than solely for identification of the perpetrator following an incident.
- Technical controls should protect system logs.
- Law enforcement should be notified for investigation assistance.

Commentary

Logging and monitoring can be used for proactive detection rather than solely for identification of the perpetrator following an incident. Many of the insider attacks were only detected once a system became unavailable or there was a noticeable irregularity in the information or system. In some cases, the system administrators were discovering the damage to the system at the same time that customers were discovering their inability to access it. In others, customers provided the initial indication that there was a problem. Once the incident was detected, a variety of system logs were key in identifying the perpetrator. Proactive monitoring of system logs therefore might facilitate detection of an incident before it becomes apparent externally. The fact that so many insiders used company computer equipment for the attack suggests that host-based countermeasures installed on the equipment could be effective in preventing or detecting attacks.

Automated Versus Manual Processes

Most attacks were detected through manual means or procedures, rather than automated mechanisms. Automated detection of precursor or attack activity can be greatly improved where that activity involves spurious system or data changes.

Anomaly detection tools that monitor and flag individual actions for user activity deviating significantly from a pre-defined profile may also be useful. However, these tools are known to be expensive to operate, only minimally effective, and not widely

available. Therefore, it is likely that the early detection of insider incidents will be accomplished most effectively for the foreseeable future using a combination of automated tools for logging, monitoring, and flagging suspicious activity, along with manual diagnosis and analysis.

For example, characterizing software, hardware, and information assets and tracking their modification, as described earlier, would provide an automated method for flagging system changes. However, those changes require a manual review process for identification of malicious system changes such as logic bombs. In addition, automated logging of account creations would facilitate detection of rogue accounts, but also require a combination of automated and manual analysis mechanisms.

Monitoring Remote Access

Remote access was used by many of the insiders to commit their attacks; they were often identified using a combination of remote access logs, source IP address, and phone records. In some cases identification was straightforward because the user name of the intruder pointed directly to the insider. Of course, corroboration of this information is required since the intruders may have been trying to frame other users or cast attention away from their own misdeeds by using other users' accounts or otherwise manipulating the monitoring process.

When to Monitor

Attacks occurred both during and outside normal working hours, suggesting that organizations cannot afford to ignore either possibility. Organizations may need to institute increased physical and remote access controls, with monitoring and (possibly automated) alerting of organizational security personnel during off hours. Improved employee awareness of behaviors indicative of malicious insider actions also helps to increase the chances of catching insiders that are engaging in harmful activity during work hours.

Technical controls should protect system logs. Because system logs were so crucial to the identification of the perpetrators, it is critical these logs be secured from manipulation and backed up. In most of the attacks, the insiders took steps to conceal their identities and/or actions. Investigators and organizational representatives should be cognizant that insiders attempt to conceal their identities and/or actions as they piece together what may have happened and who the perpetrator may be. This information is also helpful for configuring the network to effectively log future illicit activities and for tracing those activities to their source. Many insiders used technology to delete or modify records of the incident. Implementing measures to prevent such alterations will assist with identification and potential prosecution of the insider should he or she succeed in carrying out the incident.

Law enforcement should be notified for investigation assistance. Forensic examination of computers in the workplace helped to identify some of the insiders, and forensic examination of the insiders' home computers helped to corroborate the identification. Therefore, it is critical for identification and successful prosecution that

organizations contact a forensic specialist to advise the organization on maintaining the integrity of the evidence for law enforcement or other investigation.

Consequences for Targeted Organizations

Key Findings:

- Insider activities caused organizations financial losses, negative impacts to their business operations and damage to their reputations.
- Incidents affected the organizations' data, systems/networks, and components.
- Various aspects of organizations were targeted for sabotage by the insider.
- In addition to harming the organizations, the insiders caused harm to specific individuals.

Implications

- Policies and procedures should be designed for survivability of critical assets.
- Backup and recovery procedures must be followed and periodically tested.

Commentary

Policies and procedures should be designed for survivability of critical assets.

Since sabotage through insider manipulation of computer systems and networks is the focus of this report, it is unsurprising that loss of critical data, critical applications, and critical computer systems was frequent. Insiders had ample opportunity to learn which assets were most critical to company operations, and to exploit weaknesses in those assets or their management to cause harm. Consequently, because critical applications and computer systems may be specific targets of insider attacks, risk management for survivability of those assets should be considered.

For instance, in some cases insiders corrupted applications to deny service or denigrate performance, causing problems for employees and customers. This damage included inserting code to delete files, engineering transient failures in communication, and encrypting critical applications. By attacking applications, the insiders damaged organizations by directly interfering with the critical processes associated with those applications. Some organizations were forced to bargain with insiders for restoration of this information, or bring in outside help to try to recover it.

Actions directed against network performance and access were also very common. The insider attacks included cases where denial of service was attempted or successful, where customer access to organizations was specifically halted, and where employee access to network service was halted. In other cases, while network access was still possible, the attack slowed performance for customers and employees.

Attacks against networks not only cause direct harm, but 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 an attack, may specifically seek to interfere with such communication. Organizations might mitigate this effect by multi homing: maintaining trusted communication paths outside of the network with sufficient capacity to assure critical operations in the event of a network outage. This kind of protection would have a mitigating effect as insiders would be less likely to strike against connectivity because the impact and cost of strikes against the network could be mitigated.

Several insiders manipulated organization or customer web pages in the course of their attack – particularly insiders trusted with web page maintenance. While web page defacement is not usually devastating, the annoyance, embarrassment, and cost of dealing with it is not negligible.

This study highlights the need for organizations to deal proactively with potential insider threats. It suggests that action should be taken to ensure a workplace in which processes are regularly examined to assure their survivability and resiliency in the face of attacks by well-placed individuals. Critical business data must be carefully protected and appropriately validated at regular intervals. The degree of impact observed in these cases serves to justify the cost of such measures.

Backup and recovery procedures must be followed and periodically tested.

Normal defensive measures, in particular backups, may not always be effective against insider threats. This study identified cases in which attackers deleted backups, stole backup media, or performed actions that could not be undone due to faulty backup systems. In one case, the individual deleted critical data that he knew was not backed up, since he was the person responsible for backups. To guard against insider threat, organizations need to assure that backups are not only performed and periodically tested, but that the media and content is protected against modification, theft, or destruction.

Centralization of critical assets and sabotage of backups enabled some insiders to amplify the impact of their attack by eliminating redundant copies and avenues for recovery. Centralization reduces the distribution and redundancy of assets, both of which are an important part of survivable operations. While centralization can contribute to the efficiency of an organization, care needs to be taken that backups are performed regularly and are protected to ensure business continuity in the event of damage or loss to centralized data.

Section 5: Conclusion: Reflections on the Findings for the Prevention of Insider Sabotage

Although insiders in this report tended to be former technical employees, there is no demographic “profile” of a malicious insider. Ages of perpetrators ranged from late teens to retirement. Both men and women were malicious insiders. Their positions included programmers, graphic artists, system and network administrators, managers, and executives. They were currently employed and recently terminated employees, contractors, and temporary employees. As such, security awareness training needs to encourage employees to identify malicious insiders by behavior, not by stereotypical characteristics. For example, behaviors that should be a source of concern include making threats against the organization, bragging about the damage one could do to the organization, or discussing plans to work against the organization. Also of concern are attempts to gain other employees’ passwords and to fraudulently obtain access through trickery or exploitation of a trusted relationship.

Organizations need to provide training programs that create a culture of security that is appropriate for them and that includes *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. For example, if an organization places a high value on customer service quality, it may view security as protection of individual customer information, as well as the ability to serve customers. That organization could train its members about malicious employee actions on customer data and service, focusing on the key findings discussed above. Training content would be based on documented policy including a confidential means of reporting security issues with appropriate follow-up to security reports.

Employees need to understand that the organization has policies and procedures in place and will respond to detected security issues in a fair and prompt manner. Separation of duties and remote access monitoring should be explained. While employee alertness is key to detecting many insider attacks, several cases have been detected because of abnormal system activity (including changes in system configuration and illicitly escalated user privileges). Employees should be notified that system activity is monitored, especially system administration, privileged, and remote activity. All employees should be trained in their personal responsibility, such as protection of their own passwords and work products.

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. Organizations must look beyond information technology to the organization’s overall business processes and the interplay between those processes and the technologies used.

APPENDICES

APPENDIX A: Incident Date and Location

Insider Incidents by Year of Initial Damage

<i>YEAR</i>	<i>Number of Incidents</i>
1996	9
1997	5
1998	8
1999	7
2000	8
2001	6
2002	6

Locations of Insider Incidents

<i>State</i>	<i>Number of Incidents</i>
California	8
Connecticut	1
District of Columbia	1
Florida	7
Illinois	4
Indiana	1
Kansas	1
Maryland	2
Minnesota	1
Nevada	2
New Hampshire	1
New Jersey	4
New York	3
North Carolina	1
South Carolina	1
Tennessee	1
Texas	1
Utah	1
Virginia	2
Washington	5
West Virginia	1

APPENDIX B: Organizational Size

Size of Organizations²⁵

<i>Number of Employees</i>	<i>Number of Incidents</i>
1 – 100	19
101 – 500	8
501 – 3,000	7
3,001 – 10,000	2
10,001 – 50,000	3
Over 50,000	7

²⁵ Data were only available for 46 of the 49 cases studied.

APPENDIX C: Additional Case Examples

The Insider's Motive

Negative Work-Related Events

After more than four years of successful service marked by stellar performance reviews, management commendations, and nomination for the organization's executive training program, a female employee filed multiple complaints with human resources against her male supervisor and male coworkers. She claimed her coworkers had made sexual remarks, overridden her technical decisions regarding databases (an area in which she was considered an expert), and contacted her team's contractors regarding her projects without her knowledge. No action was taken by human resources, and the actions by her coworkers continued. The employee's performance reviews declined sharply in the next two years, and she was demoted. Subsequent complaints to her supervisor resulted in a suspension for insubordination. Almost a year following her written complaint to human resources, she resigned and began employment with another organization. Two months later, she learned that only her more recent, negative performance reviews had been forwarded to her new employer. She used one of several shared DBA accounts to delete critical table spaces in the organization's Oracle database, deleting crucial data. Due to a coincidental problem with database backups during the same time period, 115 employees had to spend 1800 hours to recover and re-enter the lost data.

Pre-Attack Behavior and Planning

Responding to Problematic Behavior

An insider began working with his organization as a network engineer but within a few months became network manager. Coworkers began to notice concerning behavior, which included leaving work without notice and leaving the office if management had left early or was out of town. He also set up a web cam on one of the computers in the computer room, which he positioned to observe everyone. He would watch his coworkers from home and call them to report what he witnessed. There were also rumors that the insider was taking office equipment. On at least one occasion, he showed a coworker one of the items he was going to take. During an interview for this study, the company representatives stated that the insider's coworkers attributed most of his behavior to what one would expect from an average "weird tech guy." They also were hesitant to report the behavior because they were worried about being fired due to recent management and staffing changes.

Importance of Information Sharing Mechanisms

The insider discussed in the example above, *Responding to Problematic Behavior*, gained access to the company's network following termination and caused a server to fail, shutting down all of the company's interstate and foreign commerce and communications for two and a half days. Prior to the incident, the insider telephoned a

coworker and expressed his anger over the company's decision to place a stop payment on his severance check. The company had discovered he still had some of its equipment in his possession. The insider notified his coworker that he had backdoors into the network that he would use to take it down. The coworker did not report the threat at the time, but did share it during a conversation with company representatives as part of this study. When the coworker reported his knowledge of the insider's threat during a study-related teleconference, it was the first time the vice president of information services and many others on the call had heard of it.

Advancing the Attack

Closing Open Connections for a Terminated Employee

A company disabled access for a software engineer just prior to his firing. However, he had logged into the system from home earlier in the week and maintained his connection. After being terminated, he went home to find his remote connection still open. His remote access had been disabled so he could not make any new connections, but he used the existing open connection to delete several critical files from the company's manufacturing application. The company lost over four hours of manufacturing time and had to load backup data to restart the manufacturing process.

Disabling Access for Temporary Employees

A temporary employee with system administrator access applied for a permanent position and was rejected. He reacted with an angry email and was dismissed as a result. However, the organization did not change the system administrator passwords. This enabled him to log in and delete accounts, change passwords, and clear the security logs.

Changing Passwords of Shared Accounts

A shared account was used to manage a company's voice mail system. The account required a password for administrative access. Upon the departure of one of its employees, the company overlooked changing the password of the account. The terminated employee, who possessed that password, used the password to enter the account and make changes that directed certain customers to a pornographic telephone service.

Shared Passwords

A fired employee had privileged access to the company's collaborative workspace application, which was used to maintain clients' websites. Although his access was disabled upon termination, employees in his group typically shared their passwords among the team for testing purposes. Because of this, he was able to log into the application following his termination using his supervisor's username and password. Having done so, he made malicious, embarrassing changes to the content of their clients' websites – particularly the "high-profile" clients.

Password-Protected Screensavers

A contractor was able to gain physical access to the organization's Network Operations Center where consoles were left logged in as root with no password-protected screensavers. He then deleted system files, a database, and all software from three of the company's servers, resulting in over two hundred thousand dollars in damage.

Denying Physical Access to Terminated Employees

An insider with system administrator privileges was terminated from a cancer research project that used a single, stand-alone computer. His physical access to the building was immediately disabled. However, he returned to the building after normal working hours, and when his access card denied him entry, another employee let him into the building believing that the insider's access card had malfunctioned. The insider then deleted 18 months of data from the cancer research on which his office had been working.

Lack of Fine-Grained Access Control and Separation of Duties

A programmer was given system administrator access even though system administration was not his responsibility. He used that access to plant a logic bomb on the organization's network that interrupted customer access to the organization's systems.

Lack of Separation of Duties or Two-Person Rule

The sole system administrator at an organization terminated his employment without warning and refused to divulge the administrator passwords until they met his financial demands. Furthermore, he proceeded to change the passwords for all user accounts, preventing anyone in the organization from logging into any of the company's systems. Next, he changed the IP address of its web server so no one could access its website. Finally, he created a backdoor account for later use. After revealing the administrator passwords to the organization two days later, he utilized that backdoor account to run a password sniffer on the organization's network.

Absence of Procedural and Technical Controls for System Administrators

A UNIX network administrator was reprimanded for behavioral issues; his computer accounts and remote access were then disabled. He returned the following business day to turn in his letter of resignation. Before doing so, however, he gained physical access to restricted workstations, logged in as root, and planted a time bomb that deleted all of the files on three company servers several days later. Recovery required the assistance of an outside consultant for five days. Two days following their recovery, the servers were again sabotaged in the same manner. On their second visit, the consultants discovered a destructive script on three of the company's UNIX file servers that was scheduled to run at 3 a.m. every Wednesday. The company estimated the total loss sustained by the business due to the incident at \$237,550. During the investigation, the company learned from a coworker that he and the insider had discovered a backdoor on twenty restricted workstations. On any of those workstations they could gain root access. Because of the trusted host system configuration, they could then access any of the organization's file servers as root.

Undetected Use of Sophisticated Attack Tools and Methods

An insider used a toolkit to install unauthorized backdoors on his employer's systems. These backdoors allowed the insider to gain remote access to the system after his termination, delete the computer accounts of several company executives, change passwords, and clear security logs to conceal his actions.

Use of Backdoor Accounts

One insider, while employed at the victim organization, created VPN accounts for his supervisor, the Chief Financial Officer, and the vice president of sales, but never told them. After his termination, he used those accounts to gain remote access to the system, logging in undetected for two weeks before using them to commit his attack.

Physical Sabotage to Harm Electronic Assets

A well respected employee who prided himself on his outstanding reputation in the company sabotaged the project on which he was working to disguise the fact he was failing to meet his own deadlines. He terminated processes, reformatted disks and cut computer cables, all of which caused failure of the servers on which the company ran its tests of the project.

Social Engineering

An insider used a combination of coercion and intimidation to convince the human resources (HR) manager to give him the only backup tape for the organization's mission critical software, even though the HR manager was involved in the pending firing of the employee. This action amplified his later attack, when he successfully deleted the software from the production systems.

Danger of No Software Characterization or Configuration Management

An insider modified his employer's premier product, an inter-network communication interface, to insert the character "i" at random places in the supported transmission stream and during protocol initialization. The malicious code was inserted as a logic bomb set to detonate more than six months after the insider left the company.

Use of Remote Access with Undetected Precursor Activity

A computer technician sabotaged a number of customer systems installed by his company. First, he accessed five of the customers' systems using remote access from his home. He replaced the company's programs on the customer systems with new executables that would not run. Second, he planted a boot virus on the systems so that a reboot would render them useless. As planned, the customers could not run the program the next morning, rebooted their computers, and lost the programs completely. As a result, retail operations were shut down at two of the customer sites for two days and his organization's IT staff had to fly to each customer site and spend several days recovering the systems. The investigation showed that he had unsuccessfully attempted to access one of the customer systems remotely on three different occasions prior to the attack.

Detecting the Attack

Altering System Logs to Cast Suspicion on Someone Else

An insider responded to a network disruption, diagnosed the problem, and brought the network back up fairly quickly. At the same time, however, he framed his supervisor by manipulating the system log. First, he downloaded a logic bomb script onto his organization's system, next he created a fictitious entry in the network log falsely showing that his supervisor had downloaded the logic bomb. Although the logic bomb never detonated, he used the presence of the logic bomb and the fictitious log as evidence to implicate his supervisor in the network crash.

Danger of no Technical Controls for Protecting System Logs

An insider was able to insert a malicious code into programs created by his organization and concealed his identity by turning off security settings and erasing log files.

APPENDIX D: Glossary of Technical Terms²⁶

access controls: the rules and mechanisms that control access to information systems and physical access to premises; *fine-grain access controls* describe rules and mechanisms that apply at the individual file level or below.

admin account: a computer account with system administrator privileges used to install software and manage the system.

audit: verify, independently, the quality and integrity of the work that has been undertaken within a particular area, with reference to accepted procedures.

authorized access: explicit permission to use.

autonomous agent: a self-contained program that is capable of making independent decisions and taking actions to satisfy internal goals based upon its perceived environment²⁷.

backdoors: in computer and network systems, unauthorized means for gaining access to the system known only to the person who installed them.

boot virus: an MS-DOS virus that infects the boot record program on hard disks and floppy disks or the master boot record on hard disks. The virus gets loaded into memory before MS-DOS and takes control of the computer, infecting any floppy disks subsequently accessed.

business rule check: method for comparing data to a definition reflected in the design of a database for checking data integrity and flagging inconsistencies.

critical assets: an organizational entity essential to the organization's mission or efficient functioning.

coding: composing sequences of instructions for execution on a computing system (also known as programming).

configuration and change management procedures and software: procedures or software for tracking releases and changes to software components so that previous versions can be recreated. It can also prevent unauthorized access to files or alert appropriate users when a file has been modified or released. Hardware configuration management can be facilitated through maintenance of a database containing

²⁶ Many of these definitions are taken from the Information Security Glossary (<http://www.yourwindow.to/information-security/>)

²⁷ http://en.wikipedia.org/wiki/Autonomous_agent

information about the workstations, servers, bridges, routers, and other equipment on the network.

cron: a Unix command for scheduling a program, script or command to be executed sometime in the future.

Database Administrator (DBA): the person responsible for the architecture, configuration, operation, security, maintenance, performance, and/or backup/recovery of a database.

DBA account: a privileged account that is used to perform Database Administrator functions.

denial of service attack: an attack directed towards a service, computer system or network with the objective of making it inaccessible to legitimate users.

encrypt: the process by which data is temporarily re-arranged into an unreadable or unintelligible form for confidentiality, transmission, or other security purposes.

executable: a binary file containing a program in machine language which is ready to be executed (run).

Information exchanges: a means of obtaining information either from other attackers (such as through an electronic bulletin board), or from the people being attacked (commonly called social engineering).

IP address (Internet Protocol address): number used to uniquely identify computers on the Internet.

logic bomb: malicious code implanted on a target system and configured to execute after a designated period of time or on the occurrence of a specified system action.

packet sniffer: a network monitoring tool that captures data packets and decodes them using built-in knowledge of common protocols. Sniffers are used to debug and monitor networking problems.

password sniffer: program that captures passwords of network users by analyzing traffic on the network.

privileged users: users who have an elevated level of access to a network, computer system, or application. For example, system administrators, network administrators, and Database Administrators (DBAs) have the ability to create new user accounts and control the access rights of users within their domain.

probing: accessing a target in order to determine its characteristics and/or vulnerabilities.

remote administration tool: software that enables a user to remotely monitor or work on other networked computers from the user's own workstation.

rogue accounts: accounts that exist without authorization.

root access: system administrator access on Unix-based operating systems.

scripting: preparing collections of individual commands to a computer that accomplish a more complex purpose.

spam: unsolicited or inappropriate e-mail. Also, to send such email.

time bomb: a subspecies of logic bomb that is triggered by reaching some preset time, either once or periodically.

toolkit: a software package often distributed on hacker websites containing a variety of malicious or obfuscating programs.

trusted baseline: a state of a computer system which is known or assumed not to contain malicious content or programs.

trusted host system configuration: the trusted host configuration simplifies access by bypassing password security checks otherwise required. Network administrators can define which remote hosts and which users on those hosts are trusted.

two-person control: the close surveillance and control of a system, process, or materials at all times by a minimum of two authorized persons, each capable of detecting incorrect and unauthorized procedures with respect to the tasks to be performed and each familiar with established security requirements.²⁸

Virtual Private Network (VPN): encrypted sequence of communication often used to provide secure remote access to an organization's network via the Internet.

²⁸ This definition was taken from R.W. Shirley, "Internet Security Glossary", The Internet Society/Internet Engineering Task Force RFC 2828, May 2000.

APPENDIX E: Acknowledgements

The Secret Service and CERT appreciate the work and dedication of the following personnel, without whose efforts this study would not have been possible. With many thanks to the Insider Threat Study research staff:

U.S. Secret Service, National Threat Assessment Center

Brandi Justice
Diana McCauley
Eileen Kowalski
Georgeann Rooney
Jim McKinney
Karen Damato
Lea Bauer
Lisa Eckl
Marisa Reddy Randazzo
Megan Williams
Michelle Keeney
Susan Keverline
Tara Conway

Carnegie Mellon University, Software Engineering Institute, CERT Program

Andy Moore
Bill Wilson
Bradford Willke
Casey Dunlevy
Chris Bateman
Dave Iacovetti (USSS/CERT Liaison)
David Mundie
Dawn Cappelli
Mark Zajicek
Stephanie Rogers
Tim Shimeall
Tom Longstaff
Wayne Peterson (USSS/CERT Liaison)