

The “Big Picture” of Insider IT Sabotage Across U.S. Critical Infrastructures

Andrew P. Moore
Dawn M. Cappelli
Randall F. Trzeciak

May 2008

TECHNICAL REPORT
CMU/SEI-2008-TR-009
ESC-TR-2008-009

CERT Program
Unlimited distribution subject to the copyright.

<http://www.sei.cmu.edu>



This report was prepared for the

SEI Administrative Agent
ESC/XPK
5 Eglin Street
Hanscom AFB, MA 01731-2100

The ideas and findings in this report should not be construed as an official DoD position. It is published in the interest of scientific and technical information exchange.

This work is sponsored by the U.S. Department of Defense. The Software Engineering Institute is a federally funded research and development center sponsored by the U.S. Department of Defense.

Copyright 2008 Carnegie Mellon University.

NO WARRANTY

THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

Use of any trademarks in this report is not intended in any way to infringe on the rights of the trademark holder.

Internal use. Permission to reproduce this document and to prepare derivative works from this document for internal use is granted, provided the copyright and "No Warranty" statements are included with all reproductions and derivative works.

External use. This document may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other external and/or commercial use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

This work was created in the performance of Federal Government Contract Number FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. The Government of the United States has a royalty-free government-purpose license to use, duplicate, or disclose the work, in whole or in part and in any manner, and to have or permit others to do so, for government purposes pursuant to the copyright license under the clause at 252.227-7013.

For information about SEI reports, please visit the publications section of our website (<http://www.sei.cmu.edu/publications>).

Table of Contents

Acknowledgments	v
Abstract	vi
1 Introduction	1
2 General Observations about Insider IT Sabotage	3
3 Model of the Insider IT Sabotage Problem	7
4 Possible Leverage Points for Addressing the Problem	12
5 A Workshop on Insider IT Sabotage	18
6 Conclusion	22
Appendix A: System Dynamics Background	25
Appendix B: The Insider IT Sabotage Training Case	28
Appendix C: Model of the Insider IT Sabotage Problem	32
Appendix D: Insider Sabotage Mitigating Measures	33
References	35

List of Figures

Figure 1. Expectation Escalation	8
Figure 2. a) Typical Escalation of Disgruntlement b) Intended Effect of Sanctions	9
Figure 3. Technical Precursors due to Disgruntlement	9
Figure 4. Trust Trap	10
Figure 5. Early Mitigation through Expectation Setting	13
Figure 6. Handling Disgruntlement through Positive Intervention	14
Figure 7. Targeted Monitoring	14
Figure 8. Access Paths Available to Insider	15
Figure 9. Eliminating Unknown Access Paths	16
Figure 10. Measures Upon Demotion or Termination	17

Acknowledgments

CERT¹ would like to thank the Army Research Office and Carnegie Mellon University's CyLab for funding this project.

CERT would also like to thank the following individuals for their collaboration on the MERIT model:

- Dr. Eric D. Shaw—Consulting & Clinical Psychology, Ltd., and a visiting scientist at CERT
- Dr. Stephen R. Band—Counterintelligence Field Activity, Behavioral Science Directorate
- Dr. Lynn F. Fischer—U.S. Department of Defense Personnel Security Research Center
- Dr. Elise A. Weaver—jointly as faculty at Worcester Polytechnic Institute and visiting scientist at CERT

Their expertise and experience in the psychological and social sciences areas have enabled a much richer treatment of the insider threat problem than would have otherwise been possible.

CERT also appreciates the work and dedication of the *Insider Threat Study* team members from CERT and the U.S. Secret Service, National Threat Assessment Center; without the study none of our follow-on insider threat research would have been possible.

Finally, Christopher Nguyen—a student at the Information Networking Institute of Carnegie Mellon University, Eric Hayes—our CERT technical editor, and the anonymous reviewers for the 2007 International Conference of the System Dynamics Society provided many comments that improved both the content and presentation of this paper.

This paper was published in *Insider Attack and Cyber Security: Beyond the Hacker*, eds. Stolfo, S.J., et al., Springer Science + Business Media, LLC, 2008.

¹ CERT and CERT Coordination Center are registered in the U.S. Patent and Trademark Office by Carnegie Mellon University.

Abstract

A study conducted by the U.S. Secret Service and the Carnegie Mellon University Software Engineering Institute CERT Program analyzed 150 insider cyber crimes across U.S. critical infrastructure sectors. Follow-up work by CERT involved detailed group modeling and analysis of 30 cases of insider IT sabotage out of the 150 total cases. Insider IT sabotage includes incidents in which the insider's primary goal is to sabotage some aspect of the organization or direct specific harm toward an individual. This paper describes seven general observations about insider IT sabotage based on our empirical data and study findings. We describe a system dynamics model of the insider IT sabotage problem that elaborates complex interactions in the domain and unintended consequences of organizational policies, practices, technology, and culture on insider behavior. We describe the structure of an education and awareness workshop on insider IT sabotage that incorporates the previously mentioned artifacts as well as an interactive instructional case.

1 Introduction

Insiders, by virtue of legitimate access to their organizations' information, systems, and networks, pose a significant risk to employers. Employees experiencing financial problems have found it easy to use the systems they use at work everyday to commit fraud. Other employees, motivated by financial problems, greed, revenge, the desire to obtain a business advantage, or the wish to impress a new employer, have stolen confidential data, proprietary information, or intellectual property from their employers. Furthermore, technical employees have used their technical abilities to sabotage their employers' systems or networks in revenge for negative work-related events.

In January 2002, the Carnegie Mellon University Software Engineering Institute's CERT Program (CERT) and the United States Secret Service (USSS) National Threat Assessment Center (NTAC) started a joint project, the *Insider Threat Study*.² The study combined NTAC's expertise in behavioral psychology with CERT's technical security expertise to provide in-depth analysis of approximately 150 insider incidents that occurred in critical infrastructure sectors in the U.S. between 1996 and 2002. Analysis included the study of case documentation and interviews of personnel involved in the incident.

Two reports have been published to date as part of the *Insider Threat Study*. One analyzed malicious insider incidents in the banking and finance sector [Randazzo 2004]. The other analyzed insider attacks across all critical infrastructure sectors where the insider's intent was to harm the organization, an individual, or the organization's data, information system, or network [Keeney 2005]. Two additional reports will be published in the future: one pertaining to the information technology and telecommunications sector, and the other geared to the government sector.

The *Insider Threat Study* provided the first comprehensive analysis of the insider threat problem. CERT's technical security expertise was augmented with expertise from several experts in the areas of psychology, sociology, insider threat, espionage, cyber crime, and specific domains like the financial industry. The results of the study show that to detect insider threats as early as possible or to prevent them altogether, members of management, IT, human resources, security officers, and others in the organization must understand the psychological, organizational, and technical aspects of the problem, as well as how to coordinate their actions over time.

The CERT project team felt that it was important to further use the wealth of empirical data from the *Insider Threat Study* to next concentrate on conveying the "big picture" of the insider threat problem—the complex interactions, relative degree of risk, and unintended consequences of policies, practices, technology, insider psychological issues, and organizational culture over time. Thus, the MERIT project was initiated.³ MERIT stands for the Management and Education of the Risk of Insider Threat. As part of MERIT, we are developing a series of models and associated tools that can be used to better communicate the risks of the insider threat.

² The *Insider Threat Study* was funded by the USSS, as well as the Department of Homeland Security, Office of Science and Technology, which provided financial support for the study in fiscal years 2003 and 2004.

³ The MERIT project is supported by the Army Research Office through grant number DAAD19-02-1-0389 ("Perpetually Available and Secure Information Systems") to Carnegie Mellon University's CyLab.

This paper focuses on insider IT sabotage across the U.S. critical infrastructure sectors: insider incidents in which the insider's primary goal was to sabotage some aspect of the organization (e.g., business operations, information or data files, the system or network, organizational reputation) or to harm an individual. Section 2 describes key concepts for understanding the domain in the context of seven general observations about insider IT sabotage based on our empirical work. Section 3 presents the system dynamics model of the insider IT sabotage problem, bringing into sharper focus the concepts previously described, with an emphasis on their dynamic interrelationship. Section 4 identifies leverage points for the possible mitigation of the insider IT sabotage problem. Section 5 illustrates the structure of a workshop about insider IT sabotage that incorporates the previously mentioned artifacts. Section 6 concludes with an assessment of the value of our modeling efforts and a summary of our ongoing and future work in the area. Additionally, appendices describe details of the system dynamics approach we use, an instructional case used in our insider threat workshop, and an overview of the complete insider IT sabotage model.

2 General Observations about Insider IT Sabotage

The cases of insider IT sabotage were among the more technically sophisticated attacks examined in the *Insider Threat Study* and resulted in substantial harm to people and organizations. Forty-nine cases were studied, as described in [Keeney 2005]. Eighty-six percent of the insiders held technical positions. Ninety percent of them were granted system administrator or privileged system access when hired by the organization. In those cases, 81 percent of the organizations that were attacked experienced a negative financial impact as a result of insider activities. The losses ranged from a low of five hundred dollars to a high of “tens of millions of dollars.” Seventy-five percent of the organizations experienced some impact on their business operations. Twenty-eight percent of the organizations experienced a negative impact to their reputations.

The *Insider Threat Study* focused on analysis of individual components of insider incidents, such as characteristics of the insider, technical details, planning and communication before the incident, detection and identification of the insider, and consequences of the attack. The purpose of the MERIT models is to analyze the cases in a different way. Rather than focusing on individual details of the cases, MERIT attempts to identify common patterns in the evolution of the cases over time. Although 49 insider IT sabotage cases were examined for the *Insider Threat Study*, not all of the case files contained enough information for this modeling effort. In the end, 30 IT sabotage cases were selected for use in this project based on availability of pertinent information.

In performing the “big picture” analysis of insider IT sabotage, we identified seven general observations about the cases. We then validated those observations against the empirical data from the *Insider Threat Study*. We have used the comparative case study methodology [Yin 2003], in our research. The findings from case study comparisons cannot be generalized with any degree of confidence to a larger universe of cases of the same class or category. What this method can provide, however, is an understanding of the contextual factors that surround and influence the event. We briefly describe each of those observations below, along with the percentage of cases that supports the observation. Band describes these observations in more detail, including their relevance to the problem of espionage [Band et al. 2006].

OBSERVATION 1: MOST INSIDERS HAD PERSONAL PREDISPOSITIONS THAT CONTRIBUTED TO THEIR RISK OF COMMITTING IT SABOTAGE

Personal predisposition: a characteristic historically linked to a propensity to exhibit malicious insider behavior.

Personal predispositions explain why some insiders carry out malicious acts, while coworkers who are exposed to the same conditions do not act maliciously. Personal predispositions can be recognized by certain types of observable characteristics [Band et al. 2006]:

- Serious mental health disorders—Sample observables from cases include alcohol and drug addiction, panic attacks, physical spouse abuse, and seizure disorders.

- Social skills and decision-making bias—Sample observables from cases include bullying and intimidation of coworkers, serious personality conflicts, unprofessional behavior, personal hygiene problems, and inability to conform to rules.
- A history of rule violations—Sample observables from cases include arrests, hacking, security violations, harassment complaints, and misuse of travel, time, and expenses.

All of the insiders in the MERIT cases who committed IT sabotage exhibited the influence of personal predispositions.

OBSERVATION 2: MOST INSIDERS WHO COMMITTED IT SABOTAGE WERE DISGRUNTLED DUE TO UNMET EXPECTATIONS

Unmet expectation: an unsatisfied assumption by an individual that an organization action or event will (or will not) happen, or a condition will (or will not) exist.

All of the insiders in the MERIT cases who committed IT sabotage had unmet expectations. In the *Insider Threat Study* IT sabotage cases, 57 percent of the insiders were perceived as being disgruntled. Eighty-four percent were motivated by revenge, and 92 percent of all of the insiders attacked following a negative work-related event such as termination, dispute with a current or former employer, demotion, or transfer.

Unmet expectations observed in cases include insufficient salary/bonus, lack of promotion, restriction of online actions, limitations on use of company resources, violations of privacy in the workplace, diminished authority/responsibilities, perceived unfair work requirements, and poor coworker relations.

OBSERVATION 3: IN MOST CASES STRESSFUL EVENTS, INCLUDING ORGANIZATIONAL SANCTIONS, CONTRIBUTED TO THE LIKELIHOOD OF INSIDER IT SABOTAGE

Stressful events: those events that cause concerning behaviors in individuals predisposed to malicious acts.

Ninety-seven percent of the insiders in the MERIT cases who committed IT sabotage experienced one or more stressful events, including sanctions and other negative work-related events, prior to their attack. The majority of insiders who committed IT sabotage in the *Insider Threat Study* cases attacked after termination or suspension from duties.

Stressful events observed in cases include poor performance evaluations, reprimands for unacceptable behavior, suspensions for excessive absenteeism, demotions due to poor performance, restricted responsibilities and Internet access, disagreements about salary or bonuses, lack of severance packages, new supervisors hired, divorce, and death in the family.

OBSERVATION 4: BEHAVIORAL PRECURSORS WERE OFTEN OBSERVABLE IN INSIDER IT SABOTAGE CASES BUT IGNORED BY THE ORGANIZATION

Behavioral precursor: an individual action, event, or condition that involves personal or interpersonal behaviors and that precedes and is associated with malicious insider activity.

Ninety-seven percent of the insiders in the MERIT cases who committed IT sabotage came to the attention of supervisors or coworkers for concerning behavior prior to the attack.

Behavioral precursors observed in cases include drug use, conflicts with coworkers, aggressive or violent behavior, inappropriate purchases on company accounts, mood swings, poor job performance, absence or tardiness, sexual harassment, deception about qualifications, violations of dress code, and poor hygiene. Many behavioral precursors were direct violations of explicit organizational policies and rules.

OBSERVATION 5: IN MANY CASES ORGANIZATIONS FAILED TO DETECT TECHNICAL PRECURSORS

Technical precursor: an individual action, event, or condition that involves computer or electronic media and that precedes and is associated with malicious insider activity.

Eighty-seven percent of the insiders in the MERIT cases of insider IT sabotage performed technical precursors prior to the attack that were undetected by the organization.

Technical precursors observed in cases include the download and use of hacker tools, failure to create backups, failure to document systems or software, unauthorized access of customers' or coworkers' systems, system access after termination, inappropriate Internet access at work, and the setup and use of backdoor accounts.

OBSERVATION 6: INSIDERS CREATED OR USED ACCESS PATHS UNKNOWN TO MANAGEMENT TO SET UP THEIR ATTACK AND CONCEAL THEIR IDENTITY OR ACTIONS. THE MAJORITY OF INSIDERS ATTACKED AFTER TERMINATION

Access path: a sequence of one or more access points that lead to a critical system.

Seventy-five percent of the insiders in the MERIT cases who committed IT sabotage created access paths unknown to the organization. In the *Insider Threat Study* IT sabotage cases, 59 percent of the insiders were former employees, 57 percent did not have authorized system access at the time of the attack, and 64 percent used remote access.

Many insiders in the cases analyzed used privileged system access to take technical steps to set up the attack before termination. For example, insiders created backdoor accounts,⁴ installed and ran password crackers,⁵ installed remote network administration tools, installed modems to access organization systems, and took advantage of ineffective security controls in termination processes. Many of these steps created or allowed the use of unknown access paths.

⁴ A backdoor account is an unauthorized account created for gaining access to a system or network known only to the person who created it.

⁵ A password cracker is a program used to identify passwords to a computer or network resource; used to obtain passwords for other employee accounts.

OBSERVATION 7: LACK OF PHYSICAL AND ELECTRONIC ACCESS CONTROLS FACILITATED IT SABOTAGE

Electronic access controls: the rules and mechanisms that control electronic access to information systems.

Physical access controls: the rules and mechanisms that control physical access to premises.

Ninety-three percent of the insiders in the MERIT IT sabotage cases exploited insufficient access controls. Access control vulnerabilities observed in cases include coworkers' computers unattended while logged in, ability to create accounts unknown to organization, ability to release code into production systems without verification or knowledge by the organization, and insufficient disabling of electronic and physical access at termination.

3 Model of the Insider IT Sabotage Problem

The *Insider Threat Study* investigated cases of actual insider attack. It brought to light how the problem of malicious insider retribution arises and escalates within the organization. This section describes the key elements of the insider IT sabotage problem that we saw in a majority of cases. The patterns embodied by the model were not seen in all cases, but in a sufficient number to raise concern. The next section describes the measures that an organization can take to prevent, detect, and respond to malicious insider actions based on our group's experience with the psychology of insiders as well as the managerial and technical aspects of organizational and information security. In the course of our study, we learned much more about what organizations should *not* do than what they should. Further research is needed to understand the effectiveness of various countermeasures for the insider threat problem.

After researching potential methods and tools that could be used for this purpose, system dynamics was chosen for its strengths in modeling and simulation of complex problems [Sterman 2000]. This paper is written for readers who are not familiar with system dynamics modeling. An explanation of system dynamics is provided in Appendix A and will be helpful for understanding the following description. For those readers who are familiar with system dynamics, we emphasize that we do not use the traditional causal loop diagramming notation in this paper. In our experience, the traditional notation using positive and negative signs can be confusing to audiences not familiar with system dynamics; non-technical people generally have been intimidated by the notation and technical people often read too much into the signs. In the following section, we use a more subtle notation of dashed arrows for negative influence and solid arrows for positive influence.

INSIDER EXPECTATION ESCALATION

Employee disgruntlement was a recurring factor in the insider IT sabotage cases, predominately due to some unmet expectation by the insider. For example

1. The insider expected certain technical freedoms in his⁶ use of the organization's computer and network systems, such as storing personal files, but was reprimanded by management for exercising those freedoms.
2. The insider expected to have control over the organization's computer and network system, but that control was revoked or never initially granted.
3. The insider expected a certain financial reward for his work, but bonuses were lower than expected due to the company's financial status.

Figure 1 represents the escalation of expectations that often leads to insider disgruntlement. As shown in the lower left side of the figure, the insider's personal predisposition may lead to heightened expectation. This predisposition differs from one person to the next, and influences the rate that expectations rise and fall.

⁶ Ninety-six percent of the insiders in the *Insider Threat Study* who committed IT sabotage were male. Therefore, male gender is used to describe the generic insider throughout this paper.

The rise of an insider's expectations is influenced heavily by the expectation fulfillment. Policies and management controls are needed to keep employee expectations in check. As illustrated in reinforcing loop (R1), with lax management controls the insider's expectation grows commensurate with the expectation fulfillment. Fulfillment of expectation only serves to stimulate greater expectation.




Figure 1. Expectation Escalation

Lax management that permits continually increasing employee expectation can result in major problems later, especially if the insider is so predisposed. The trigger for those major problems, which we call the precipitating event, tends to be anything that removes or restricts the freedom or recognition to which the insider has become accustomed. For instance, the hiring of a new supervisor who suddenly enforces the organization's acceptable use policy may cause extreme disgruntlement in certain employees. Other precipitating events include the insider being passed up for a promotion, sanctions by management, or termination of the insider.

ESCALATION OF DISGRUNTLEMENT

Often, the first sign of disgruntlement is the onset of behavioral precursors. Behavioral precursors include observable aspects of the insider's social (non-technical) behavior inside or outside the workplace that might be deemed inappropriate or disruptive in some way. Some examples of behavioral precursors in the MERIT cases were conflicts with coworkers; a sudden pattern of missing work, arriving late, or leaving early; or a sudden decline in job performance.

As shown in Figure 2a, the degree of disgruntlement influences the insider's exhibition of behavioral precursors, which can be discovered provided the organization has sufficient behavioral monitoring in place. An organization's punitive response to inappropriate behaviors in the form of sanctions can be technical, such as restricting system privileges or right to access the organization's systems from home, or non-technical, such as demotion or formal reprimand. The intended effect of sanctions, as shown in the balancing loop B1 of Figure 2b, is to prevent additional behavioral precursors. Feedback loop R2, however, shows that sanctions can have unintended con-

sequences, such as escalation of disgruntlement. Whether sanctions curb behavioral precursor activity or spur the insider to greater disgruntlement and disruption depends largely on the personal predispositions of the insider.




Figure 2. a) Typical Escalation of Disgruntlement b) Intended Effect of Sanctions

ATTACK SETUP AND CONCEALMENT

Given an insider with personal predispositions, unmet expectations can lead to increased disgruntlement which, if left unchecked, can spur not just behavioral precursors but technical disruptions and attacks on the organization’s computer and network systems. Prior to the actual attack, there are typically technical precursors—actions by the insider to either set up the attack (for example, installing malicious software programs) or to put in place mechanisms to facilitate a future attack (for example, creation of backdoor accounts—secret, unauthorized accounts to be used later for the attack). These technical precursors could serve as an indicator of a pending attack if detected by the organization.

Figure 3 depicts the influence that insider disgruntlement can have on the occurrence of technical precursors that may indicate a pending attack. Some of these actions also contribute to the damage potential of the attack. Examples include sabotage of backups and decreases in the redundancy of critical services or software. As shown in loop R3, insiders may also acquire access paths unknown to the organization. This increases the insiders’ ability to conceal their activity, making it more difficult for the organization to discover the precursors. The feedback loop is reinforcing, since the ability to hide their actions may embolden the risk-averse insiders to continue, or even increase, their efforts to attack.

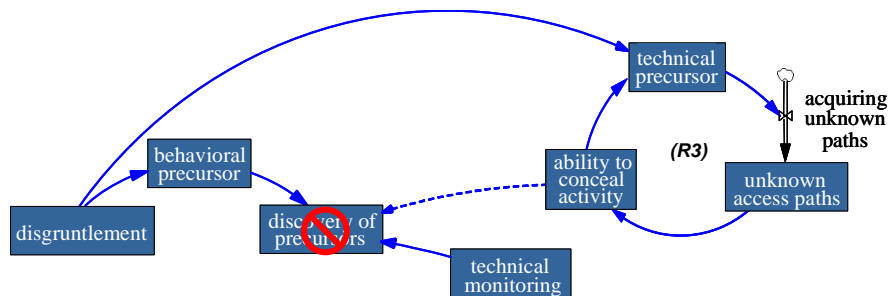


Figure 3. Technical Precursors due to Disgruntlement

The extent to which insiders rely on unknown access paths to set up and execute their attack depends on their risk tolerance. Insiders who do not care whether they are caught, or insiders acting

impulsively (often out of the passion of the moment), may use both known and unknown paths in their attack. Insiders who are particularly risk averse may only attack using access paths that are unknown to the organization. Of course, an insider may not know whether the organization is aware of a particular access path. Nevertheless, in either case, insiders who commit IT sabotage generate technical precursors that suggest suspicious activity. Just as for behavioral precursors, the detection of technical precursors depends on having a sufficient level of technical monitoring in place.

THE TRUST TRAP

In addition to insider predispositions and behaviors, organizational predispositions and behaviors—such as excessive trust of employees, a reluctance to “blow the whistle” on coworkers, or inconsistent enforcement of organization policies—can also influence an organization’s exposure to malicious insider acts. Figure 4 depicts a trap in which organizations sometimes find themselves. We call this the Trust Trap and have described its role in previous models [Anderson et al. 2004, Cappelli et al. 2006a, Band et al. 2006].

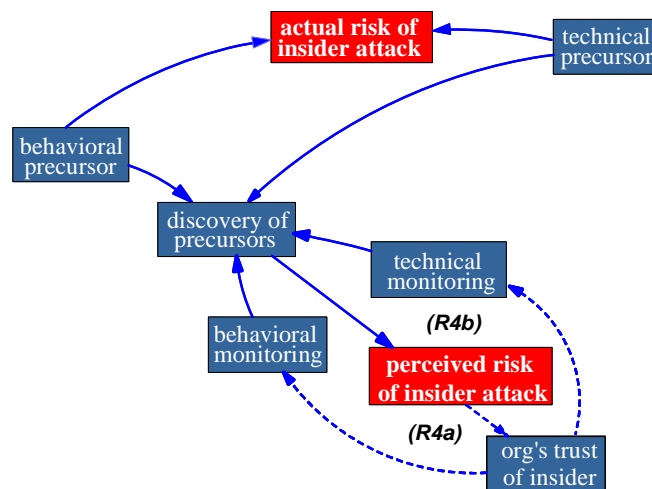


Figure 4. Trust Trap

To understand the Trust Trap, we need to distinguish between the actual and perceived risk of an insider attack. As shown in the top portion of Figure 4, actual risk depends on the behavioral and technical precursors exhibited by the insider. However, the risk of insider attack is only perceived by the organization to the extent that it discovers those precursors.

A key factor in the Trust Trap is the organization’s trust of the insider, as shown in loops R4a and R4b. Clearly, there are good reasons why managers want to create a workplace in which individuals can trust each other and there is a good trust relationship between the organization and its employees (e.g., to increase morale and productivity). However, managers who strive to promote trusting workplace relationships sometimes shortcut essential behavioral and technical monitoring procedures, or allow them to erode over time due to competing pressures and priorities. Lower levels of monitoring lead to undiscovered precursors, resulting in an overall lower perceived risk

of attack. This false sense of security reinforces managers' trust in the individuals working for them. The cycle continues, with the organization's monitoring capability steadily deteriorating until a major compromise becomes obvious to all involved.

4 Possible Leverage Points for Addressing the Problem

The intent of the MERIT project is to communicate the severity of the insider threat problem and describe it using system dynamics models based upon empirical data. Although our research has focused on the insider threat problem, we would be remiss to leave participants with the impression that the organization is helpless to defend itself against someone from within. We can propose effective countermeasures based on our team's expert opinions in behavioral psychology and information security.⁷ All levels of management should recognize and acknowledge the threat posed by insiders and take appropriate steps to mitigate malicious attacks. While it may not be realistic to expect that every attempt at insider IT sabotage will be stopped before damage is inflicted, it is realistic to expect that organizations can build resiliency into their infrastructure and business processes to allow them to detect the attacks earlier, thereby minimizing the financial and operational impact.

This section of the paper describes potential countermeasures that we believe could be effective in mitigating insider IT sabotage. The identification of these countermeasures is based on expert opinion and our analysis of the problem.

EARLY MITIGATION THROUGH EXPECTATION SETTING

First, managers should recognize the personal predispositions of their employees and understand the impact they can have on insider threat risk. Second, organizations should attempt to manage the expectations of employees to minimize unmet expectations. This can be achieved through communication between managers and employees (especially in the form of regular employee reviews), taking action to address employee dissatisfaction when possible, and consistent enforcement of policies for all employees so that individual employees do not come to feel that they are above the rules or that the rules are inconsistently enforced.

Figure 5 describes the influence expectation setting can have on the insider's unmet expectations. When the expectations of the insider are in line with the organization's practices and policies, unmet expectations are not an issue. However, if a precipitating event impacts expectation fulfillment, actions by management to reset expectations may decrease the level of unmet expectations. If the organization fails to reset expectations, the level of unmet expectations may continue to rise, increasing disgruntlement on the part of the insider.

⁷ The effectiveness of the countermeasures proposed in this section is not supported in the case data since we were rarely able to obtain that kind of data during the coding process.




Figure 5. Early Mitigation through Expectation Setting

For example, the organization can attempt to lower the level of unmet expectations regarding system use and job responsibilities by a number of proactive countermeasures:

- The organization institutes an acceptable use policy, describing the employee's roles and responsibilities when using the organization's information systems. The policy should be given to each employee as part of their orientation to the organization. As changes to the policy occur, employees need to be made aware of the changes and the impact to them. In addition, the policy should be consistently enforced for all employees so that no employees may feel that they are "above the rules."
- Management, in conjunction with Human Resources, should clearly define job responsibilities for each employee in the organization. Processes such as performance reviews should be used to check and set expectations periodically.

HANDLING DISGRUNTLEMENT THROUGH POSITIVE INTERVENTION

As the organization discovers the behavioral precursors exhibited by the insider, it can employ positive intervention strategies to lower the disgruntlement of the insider. While the intent of employee sanctioning may be to reduce undesirable behaviors, it may backfire in some cases, causing disgruntlement to increase and leading to more disruptive behaviors. Figure 6 describes the influence positive intervention strategies may have on the disgruntlement of the insider. When positive intervention is used, the disgruntlement may be reduced, eliminating additional behavioral precursors, as well as the escalation to technical precursor behaviors (see Figure 3).

One positive intervention strategy is an employee assistance program (EAP). EAPs are sometimes offered by organizations as an employee benefit, to assist employees in dealing with personal or work-related issues that may affect job performance, health, and general well-being. EAPs can include counseling services for employees and/or their family members.




Figure 6. Handling Disgruntlement through Positive Intervention

TARGETED MONITORING

It is usually not practical for an organization to monitor every behavioral and technical action taken by each employee. However, a reasonable level of proactive logging of online activity across the organization's network provides data that can be monitored or audited for suspicious activity proactively, or targeted to monitor people who have raised the suspicions of their managers. Based on findings from the *Insider Threat Study*, for example, periodic account audits could be effective in detecting backdoor accounts that could be used for malicious insider activity.

Figure 7 describes the relationship between the perceived risk of an insider attack and the amount of technical and behavioral monitoring organizations institute. As the perceived risk of an insider attack increases, due to detection of behavioral or technical precursors, the amount of technical and behavioral monitoring should also increase. Increased monitoring could lead to discovery of precursor activity, enabling the organization to identify individuals at a higher risk for malicious behavior and implement more targeted individual monitoring.

If a manager notices an employee progressing through the pattern of behavior described above, he might consider an audit of that employee's online activity. If the employee's behaviors, either technical or non-technical, are extreme enough, managers may need to escalate the level of logging and monitoring of that employee's online activity. Note that policies should be in place in advance of such targeted logging and monitoring; an organization should not perform these actions without consulting with their legal department in advance. Band identifies specific observable behaviors that should impact an organization's trust level [Band et al. 2006].

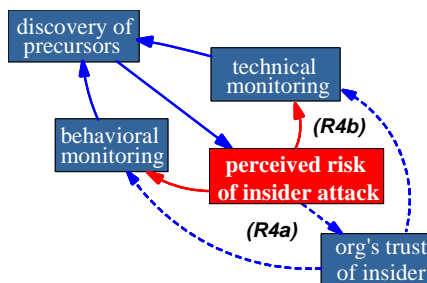


Figure 7. Targeted Monitoring

ELIMINATING UNKNOWN ACCESS PATHS

An organization's full awareness of access paths available to an insider is critical to being able to disable those access paths when needed. Figure 8 reflects the relationship between two variables: insider access paths unknown to the organization and insider access paths known to the organization. Important relationships between the two variables include

- *forgetting paths*: Management or the IT staff may forget about known paths, making them unknown. The forgetting paths variable represents the rate at which access paths move from the known to the unknown category. For example, a manager might authorize a software developer's request for the system administrator password during a time of heavy development. Therefore, the system administrator password is an access path known to the organization at that point in time. If a formal list of employees with access to that password is not maintained, the manager could forget that decision over time. The manager may also simply resign from the organization, leaving no organizational memory of the decision to share the system administrator password. In either case, the software developer's knowledge of the system administrator password has become an access path unknown to the organization.
- *discovering paths*: The discovering paths variable represents the rate at which management or the IT staff discover unknown paths, making them known. Access paths can be discovered by monitoring network traffic or by computer system account auditing, for example. Monitoring network traffic facilitates discovery of suspicious network traffic for further investigation. Account auditing facilitates discovery of unauthorized accounts directly.

Insiders can acquire new access paths unknown to the organization (through the acquiring unknown paths variable shown in Figure 8) by, for example, installing a backdoor account or stealing passwords. Finally, organizations can disable access paths that they know about (through the disabling known paths variable shown in Figure 8) by, for example, removing backdoor accounts or changing shared passwords. The critical concept is that an organization may not know about all of the access paths each of their employees has to its critical systems.




Figure 8. Access Paths Available to Insider

Access paths unknown to the organization provide a mechanism that can be used by the insider to facilitate a future attack, even following termination. For example, organizations often did not know about (or did not think about) insiders' access to shared accounts like system administrator or database administrator accounts; overlooking such accounts during an insider's termination process often allowed an insider's attack following termination. In addition, unknown access

paths can make it more difficult for the organization to attribute the attack to the insider. If the organization is unaware of the paths that can be used by an insider for attacks, the task of protecting itself is significantly more complex.

Figure 9 emphasizes the importance of diligent tracking and management of access paths into the organization's system and networks. As tracking increases, the likelihood an organization will forget about the existence of specific access paths and who has access to them decreases. If precursor technical activity is detected, unknown access paths can be discovered and disabled, further reducing the number of unknown access paths available to the insider. As the number of unknown access paths decreases, the ability to conceal malicious activity by the insider decreases. As the ability to conceal decreases, the discovery of technical precursors increases. This makes it more and more difficult for the insider to conceal unauthorized or malicious online activity. Conversely, if technical precursors are not discovered, the insider can accumulate unknown access paths, making it easier for him to conceal his actions.

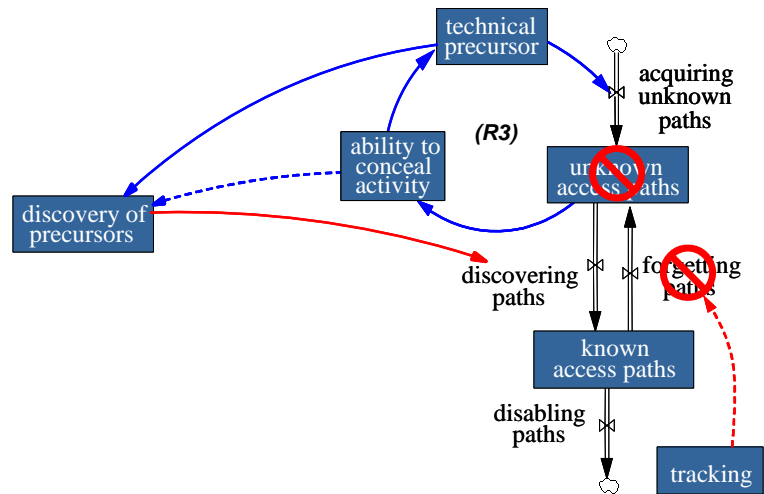


Figure 9. Eliminating Unknown Access Paths

In the cases we examined, the access paths often used by the insider but unknown to management were accounts secretly created by the insider or shared with other coworkers. Compounding the problem, lack of tracking led to unknown access paths that were overlooked in the termination process and later used by the insider to attack. Therefore, we consider ongoing and thorough account management an important practice for tracking access paths and reducing the occurrence of unknown access paths. Account management is a complex task that encompasses verifying new accounts, changing account authorization levels, tracking access to shared accounts, and decommissioning old accounts. Unfortunately, it takes a significant amount of time and resources for an organization to recover from obsolete account management practices.

MEASURES UPON DEMOTION OR TERMINATION

Termination or demotion was the final precipitating event in many cases we examined. It is important for organizations to recognize that such precipitating events may cause the insider to take technical actions to set up and carry out the attack, possibly using previously acquired unknown

5 A Workshop on Insider IT Sabotage

Our insider IT sabotage workshop has the following structure:

- *Insider Threat Study* overview
- interactive discussion of an instructional case of insider IT sabotage
- general observations from the insider IT sabotage cases
- system dynamics model: problem, prevention, and mitigation
- recommendations for countering the threat

We introduce domain concepts during the *Insider Threat Study* overview and in the interactive instructional case discussion, which is described below. By the time we introduce the model, participants are very familiar with the primitive domain concepts. The model then serves to bring into sharper focus the concepts previously described, with an emphasis on the dynamic interrelationship of those concepts. This approach helps ensure that workshop participants are not overwhelmed with too many new concepts, both modeling and domain, at the same time. The rest of this section provides an overview of our instructional case and general observations about insider IT sabotage. We present this material to illustrate how domain concepts are presented prior to presenting the model.⁸

THE INSTRUCTIONAL CASE

A concrete case example helps to clearly illustrate the relationship between the aspects of the insider threat and the effectiveness of various measures to counter the threat. However, the sensitivity of actual *Insider Threat Study* case data precludes the use of actual cases for training. We, therefore, developed a fictional case scenario that is representative of a preponderance of actual cases of insider IT sabotage from the *Insider Threat Study*. The fictional organization is called iAssemble, Inc.⁹ The full text of the iAssemble case example provided in Appendix B is a substantial revision of a previously published version [Cappelli et al. 2006a] based on specific guidelines in the area [Naumes and Naumes 1999]. Moore provides guidelines for using the instructional case in a classroom setting [Moore et al. 2007].

We believe that the iAssemble case provides a coherent and well-grounded basis for training on the important issues relevant to insider IT sabotage and is representative in character (but not necessarily detail) of many of the actual cases that we have seen. This fictional case deals with the events surrounding an insider IT sabotage case at iAssemble. Ian was hired during the company startup as a computer specialist and technical assistant to the original founders. With hard work and dedication, he became the sole system administrator at iAssemble, a position he held for four years. He was also responsible for building the software that ran the company's computer sys-

⁸ Anyone interested in attending CERT's insider threat workshop can contact them at insider-threat-feedback@cert.org.

⁹ The iAssemble organization and case example are completely fictional; any resemblance to a real organization or insider threat case is unintentional.

tems. With the increase in sales at iAssemble and its focus on making profits and meeting deadlines, iAssemble began hiring new people.

After being passed over for the new lead system administrator position, Ian began acting out in the workplace. In the next few months, Ian's disruptions to iAssemble operations grew to the point that iAssemble managers decided they had no choice but to let him go. On the day he was fired, Ian installed a malicious software program, which is generically referred to as a logic bomb, on iAssemble's central server. The logic bomb, which detonated one week after Ian's dismissal, deleted all of the programs supporting iAssemble's mission critical processes that Ian himself had developed. The instructional case describes in more detail the motive behind the attack, the technical and/or non-technical actions taken by the insider and the organization, and the impact of those actions on iAssemble.

The following four questions are used to focus the workshop participant's attention on the issues and concepts central to understanding insider IT sabotage. They focus on

- identifying behavioral and technical precursors exhibited by the insider
- understanding how Ian's personal predispositions and unmet expectations caused an escalation of disgruntlement that was triggered by a precipitating event
- discussing the technical aspects of the access paths into the organization's systems that are available to the insider

QUESTION 1: WHY DID IAN ATTACK IASSEMBLE?

The concepts of unmet expectation and personal predisposition are critical to understanding why Ian attacked iAssemble. The root cause of Ian's disgruntlement was his unmet expectation – his expectation for recognition and for control of the system. Ian enjoyed four years at iAssemble in which he had total control over the design and evolution of the company's systems and networks. During that time, his expectation of continued control and prominence within the organization grew and became firmly entrenched. Ian's personal predispositions exacerbated his sense of entitlement. Personal risk factors included Ian's arrogant behavior in the workplace and his alcohol addiction problem. Ian was also under great personal stress due to family issues, which further amplified his disgruntlement at work.

QUESTION 2: WHY WAS IAN ABLE TO HARM IASSEMBLE AFTER FIRING?

Discussion about this question typically focuses on how Ian accessed the system. Also relevant is the organization's focus, prior to the attack, almost exclusively on the growth of the company with little or no recognition of the risks associated with that growth or with Ian's actions in particular. Another key question is why iAssemble fired Ian before cutting off all access. This naturally leads to the definition of an access path, and the fact that Ian had access paths into the computer networks about which the organization did not know. For example, Ian was able to use his coworker's account to plant the logic bomb because they had shared passwords months earlier. In addition, the logic bomb itself can be viewed as an (unknown) access path in that it allows the insider to take action within the organization's network even when all of his direct connections have been severed.

QUESTION 3: WHAT COULD IASSEMBLE HAVE DONE TO PREVENT THE ATTACK?

One focus of this question is to understand the importance of actions or events that occurred, or conditions that existed, prior to the insider's attack. Precursors may be both technical and behavioral in nature. For example, the sharing of passwords between Ian and a coworker facilitated the attack. The password sharing also opened an access path to the insider that the organization did not know about. In this case, the organization may have closed this avenue of attack by

- prohibiting password sharing by policy, reinforced through periodic security awareness training
- instituting regular password changes, including administrator or other group accounts
- requiring all employees to change their passwords when Ian was fired

Another focus of this question is to discuss how the organization could have used the knowledge of precursors to prevent the insider attack or otherwise mitigate the risk of attack. Behavioral precursors are often one of the first indicators of employee disgruntlement. If an organization is successful in identifying these precursors and taking measures to address them in a timely fashion, they might be successful in preventing attacks. This depends on perceptive management and targeted behavioral monitoring.

Technical precursors to an attack are even more serious and usually follow, but may come in parallel with, behavioral precursors. They may, by themselves, cause disruption in the organization's systems. They often indicate steps taken to set up a future attack on the organization's systems, possibly unbeknownst to the organization, such as creation of malicious code. Other technical precursors simply enable the insider to conceal his malicious acts. For instance, insiders often create fictitious (backdoor) accounts for their surreptitious entry to the system at a later date. This is an example of an access path that is not known by management. The organization needs to have technical monitoring in place to be able to detect such precursors at an early stage and they must take appropriate actions. While behavioral precursors, by themselves, are indicative of insider threat risk, the combination of technical and behavioral precursors indicates an even greater risk of insider attack.

QUESTION 4: WHAT SHOULD IASSEMBLE DO IN THE FUTURE?

This question requires participants to take a step back from the details of the particular scenario to describe what iAssemble should do in the future to ensure that the risk of insider IT sabotage is acceptably mitigated. Effective risk mitigation strategies should focus as much on understanding and reducing the impact of possible attacks as they do on preventing them in the first place. Organizational priority should be given to those malicious acts with the largest potential impact to the organization.

Of course, organizations cannot prevent all malicious acts. They cannot even always monitor for all precursors equally. Difficult decisions must be made, especially regarding monitoring for technical precursors, due to the associated costs. Resources used to audit and monitor technical accounts and activities may divert effort from project deliverables. Comprehensive monitoring of all employees is often not cost effective. Organizations can, however, implement proactive moni-

toring and logging of all staff for a key set of precursors. When circumstances dictate an increased risk, they should engage in more detailed, targeted monitoring.

There are also difficult questions regarding which measures should be used to mitigate risks. Should the organization use technical measures, like restricting access to curtail the risk of insider attack? Should it use non-technical measures, like a warning or reprimand for concerning behaviors? An organization needs to take into account the effect of these technical measures on morale and productivity as well as risk. The organization also needs to be aware of access paths available to the insider, including indirect access paths like malicious code.

If choosing non-technical measures to reduce risk, the organization needs to consider positive intervention, such as constructive dialogue with employees or employee assistance program (EAP) referrals, in addition to punitive techniques, such as reprimands or sanctions. For certain insiders, punitive techniques may increase the insider's disgruntlement to such a degree that they have little regard for future repercussions. Excessive or unreasonable monitoring within the workplace may make employees feel like they are being watched by "Big Brother." Low levels of trust within the workplace can discourage employees, creating an environment of low morale and low productivity. Of course, different workplace cultures accept different levels of monitoring. Organizational management has to find the right balance between providing a trusting workplace environment and managing risks associated with insider IT sabotage.

6 Conclusion

This paper describes how we have used the empirical data, in combination with our library of insider threat cases, to create materials for raising awareness on insider threat and risk mitigation. We describe our use of system dynamics modeling to better communicate the nature of the insider IT sabotage problem and potential leverage points for its mitigation. The MERIT model effectively combines psychological and technical findings from the joint CERT/U.S. Secret Service *Insider Threat Study* with other insider threat expertise from participating researchers. MERIT has greatly enhanced our ability to lead facilitated training on the risk of insider IT sabotage. This section describes the value of system dynamics modeling to achieving a better understanding of the insider IT sabotage problem. It also highlights the directions for our ongoing and future work.

VALUE OF MODELING FOR INSIGHT

We found that the system dynamics approach helped to structure and focus the team's discussion. This was important since members of the team, by necessity, came from the different disciplines of psychology and information security.

By identifying the primary variables of interest, the influences between these variables, and the feedback loops that are so important for understanding complex behavior, the team found itself able to communicate much more effectively. The group modeling process enabled the team to step back and consider the "big picture" at times and focus on individual concepts at other times. The rigorous notation helped identify commonalities to simplify the models and prevent misunderstandings that could have otherwise hindered progress. In addition, it was immensely valuable for each team member to be able to come away with the models that we developed after our group sessions and devote individual thought to each. It not only documented our progress but also helped us pick up from where we left off after a period of downtime and reflection on what we had accomplished. The models also provided a concrete target for validation through mapping to observables exhibited by the real-world cases.

Significant methodological and data challenges must be overcome before research on insider activity can be soundly prescriptive for mitigation policies, practices, and technology. However, we cannot overestimate the importance of looking at the total context of adverse insider behavior for understanding why these events happened and how they might be prevented in the future. By using the system dynamics approach, we attempt to assess the weight and interrelatedness of personal, organizational, social, and technical factors as well as the effectiveness of deterrent measures in the workplace. Prospective studies of these phenomena will always be challenging because of low base rates. In the meantime, system dynamics modeling using available empirical data can bridge this methodological gap and translate the best available data into implications for policies, practices, and technologies to mitigate insider threat.

RELATED CERT/CYLAB RESEARCH

Ongoing and future research into insider threat at CERT and CyLab includes three areas:

- a broader study of insider threat

- the development of an insider threat diagnostic
- the development of a training simulation for improved insider threat risk education, awareness, and mitigation

This section summarizes each of these areas of work.

Broader Study of Insider Threats

The library of assets produced by the MERIT project provides a collection of tools that have been very effective in transitioning our knowledge of insider IT sabotage to an international audience of security experts, IT practitioners, all levels of government and business managers, and law enforcement. Insider threat workshop participants appreciate the interactive nature of the initial discussions and the use of the model to interrelate important, but complex, insider IT sabotage domain concepts. They have one primary complaint: They need to understand insider fraud and theft of confidential or sensitive information in the same depth that the workshop provides for insider IT sabotage. Results of the *Insider Threat Study* show that insider fraud using IT is a significant problem in industry, especially in the banking and finance sector [Randazzo 04]. Likewise, theft of information using IT, including crimes like identity theft and corporate espionage, is a significant problem in today's privacy-conscious and competitive corporate world. Case data collected on these two crimes bolster this need by showing their significant difference as compared to IT sabotage crimes, especially in insider motivation, insider characteristics, and the technical nature of the malicious activity [Cappelli 06b].

The primary objective of our broader study is to extend MERIT to include a comprehensive pattern analysis and transition mechanism for all types of insider threat, including fraud, theft of confidential or sensitive information, and IT sabotage. Outputs of this project will include a complete package of empirically based insider threat system dynamics models, as well as a full-day insider threat workshop that includes in-depth analysis and interactive discussion of the behavioral and technical aspects of insider fraud, theft of confidential or sensitive information, and IT sabotage. We expect that participation in the workshop will empower corporate and government personnel to develop comprehensive, efficient, and justifiable defenses to insider threats along with the organizational understanding and support needed to maintain a strong security posture over time. In addition to looking in detail at insider fraud and theft of information cases, we are now collecting and analyzing insider compromises that have occurred since 2002. This extends the terms of analysis of the original *Insider Threat Study*, which analyzed insider compromises against U.S. critical infrastructure sectors occurring from 1996 to 2002. A focus of this broader analysis will be to determine how the insider threat is evolving as well as to generate a larger dataset on which to base findings. We plan to update our common sense guide to insider threat prevention and detection based on the results of this work [Cappelli 2006b].

Insider Threat Diagnostic Instrument

The objective of this project is to build a comprehensive diagnostic instrument empirically based on all of our prior insider threat research. This instrument can be used by organizations to self-assess their insider threat risk, with the ultimate goal of improving the resiliency and survivability of the organization. The insider threat diagnostic will enable organizations to gain a better understanding of current insider threat activity and an enhanced ability to assess and manage associated risks. It will merge technical, organizational, personnel, and business security and process issues

into a single, actionable framework. As in our past projects, our project team includes psychological and technical expertise. The instrument will be structured to encompass all stakeholders in the fight against insider threat: management, information technology, human resources, and physical security.

We will build a pilot instrument based on over 250 insider threat cases in the CERT case library and continue to expand our library with recent cases for inclusion in this research. We welcome collaboration with external organizations on this project. Collaboration opportunities range from review of the instrument to confidential sharing of insider case and/or best practice information for inclusion in the instrument. In return for participation, we will offer those organizations opportunities to pilot the insider threat risk assessment diagnostic. Following each pilot, we will provide a confidential report on the findings of the pilot and suggestions for improvement. As with all of our insider threat research, all collaborations will remain confidential and no references will ever be made to any organizations and/or individuals.

Training Simulation for Insider Threat

A project that started in September 2006, called *MERIT-Interactive*, builds upon the MERIT foundation to develop a stand-alone tool that can be used for more effective widespread training on insider threat risk education, awareness, and mitigation. In collaboration with Carnegie Mellon's Entertainment Technology Center, we used state-of-the-art multi-media technologies to develop a compelling training simulation (*MERIT-Interactive*) that immerses players in a realistic business setting from which they first make decisions regarding how to prevent, detect, and respond to insider actions and then see the impacts of their decisions in terms of key performance metrics.

The *MERIT-Interactive* proof of concept provides

- a stand-alone, multi-media training simulation for interactive and independent hands-on analysis of the effects of decisions regarding policies, practices, and technology on malicious insider activity based on the MERIT model for IT sabotage
- an effective means to communicate insider threat risks and tradeoffs, useful for both technical and non-technical personnel, from system administrators to corporate CEOs
- the state of the practice information regarding insider threats and effective countermeasures

We finished the proof of concept and now seek funding to develop a production version of the tool. We believe that *MERIT-Interactive* will ultimately help decision makers better understand risk from insider threat and the role their decisions play in promoting or mitigating that risk.

Appendix A: System Dynamics Background

System dynamics is a method for modeling and analyzing the holistic behavior of complex problems as they evolve over time. System dynamics has been used to gain insight into some of the most challenging strategy questions facing businesses and government for several decades. System dynamics provides particularly useful insight into difficult management situations in which the best efforts to solve a problem actually make it worse. Examples of these apparently paradoxical effects include the following [Sterman 2000]:

- low-nicotine cigarettes, supposedly introduced to the benefit of smokers' health, that only result in people smoking more cigarettes and taking longer, deeper drags to meet their nicotine needs
- levees and dams constructed to control floods that only produce more severe flooding by preventing the natural dissipation of excess water in flood plains

The *Insider Threat Study* found that intuitive solutions to problems with employees often reduce the problem in the short term but make it much worse in the long term. For example, employee termination might solve an immediate problem, but it may also lead to long-term problems for the organization if the insider has the technical means to attack the system following termination. System dynamics is a valuable analysis tool for gaining insight into long-term solutions and for demonstrating their benefits.

A powerful tenet of system dynamics is that the dynamic complexity of problematic behavior is captured by the underlying feedback structure of that behavior. We decompose the causal structure of the problematic behavior into its feedback loops to understand which loop is strongest (i.e., which loop's influence on behavior dominates all others) at particular points through time. We can then thoroughly understand and communicate the nature of the problematic behavior and the benefits of alternative mitigations.

System dynamics model boundaries are drawn so that all the enterprise elements necessary to generate and understand problematic behavior are contained within them. This approach encourages the inclusion of soft (as well as hard) factors in the model, such as policy-related, procedural, administrative, or cultural factors. The exclusion of soft factors in other modeling techniques essentially treats their influence as negligible, which is often not the case. This endogenous viewpoint helps show the benefits of mitigations to the problematic behavior that are often overlooked, partly due to a narrow focus in resolving problems.

In this project we rely on system dynamics as a tool to help understand and communicate contributing factors to insider IT sabotage and espionage threats and implications for various mitigation strategies and tactics. It is tempting to use the simulation of the model to help predict the effect of mitigation strategies, but what is the nature of the types of predictions that system dynamics facilitates? Dennis Meadows offers a concise answer by categorizing outputs from models as follows [Meadows et al. 1974].

1. Absolute and precise predictions (Exactly when and where will the next cyber attack take place?)
2. Conditional precise predictions (How much will it cost my organization if a cyber attack occurs?)
3. Conditional imprecise projections of dynamic behavior modes (If a bank mandates background checks for all new employees, will its damages from insider fraud be less than they would have been otherwise?)
4. Current trends that may influence future behavior (What effect will current trends in espionage have on national security in five years?)
5. Philosophical explorations of the consequences of a set of assumptions, without regard for the real-world accuracy or usefulness of those assumptions (If a foreign country succeeds in human cloning, how would this affect the United State's risk of espionage?)

Our models, and system dynamics models in general, provide information of the third sort. Meadows explains further that “this level of knowledge is less satisfactory than a perfect, precise prediction would be, but it is still a significant advance over the level of understanding permitted by current mental models.”

As described in the main body of this paper, we have modified the system dynamics causal loop diagram notation to better suit the expected participants of our workshop. Arrows still represent the pair-wise influence of the variable at the source of the arrow on the variable at the target of the arrow, but their look indicates how they should be interpreted:

- Roughly, a *solid* arrow indicates a *positive* influence—that the value of the source and target variables moves in the same direction.¹⁰
- Roughly, a *dashed* arrow indicates a *negative* influence—that the value of the source and target variables moves in the *opposite* direction.¹¹

As mentioned, dynamically complex problems can often be best understood in terms of the feedback loops underlying those problems. There are two types of feedback loops: *balancing* and *reinforcing*.

- Balancing loops (labeled B# in the figures presented in this report) describe the system aspects that oppose change, tending to drive organizational variables to some goal state. In other words, balancing loops tend to move the system to an equilibrium state even in the face of change. The behavior of a thermostat is an example of a balancing loop. It continually changes the air flow into a room based on the temperature of the room, with the goal of maintaining an equilibrium temperature.

¹⁰ More formally, a *solid* arrow indicates that if the value of the source variable increases, then the value of the target variable increases above what it would otherwise have been, all other things being equal. And, if the value of the source variable decreases, then the value of the target variable decreases below what it would otherwise have been, all other things being equal.

¹¹ More formally, a *dashed* arrow indicates that if the value of the source variable increases, then the value of the target variable decreases below what it would otherwise have been, all other things being equal. And, if the value of the source variable decreases, then the value of the target variable increases above what it would otherwise have been, all other things being equal.

- Reinforcing loops (labeled R# in the figures presented in this report) describe the system aspects that tend to drive variable values consistently upward or consistently downward. In other words, reinforcing loops can “spiral out of control.” A flu epidemic is an example of a reinforcing loop. It spirals out of control as more and more people contract the flu.

The type of a feedback loop is determined by counting the number of negative influences along the path of the loop. An odd number of negative influences indicates a balancing loop, and an even (or zero) number of negative influences indicates a reinforcing loop.

System dynamics models are described as a sequence of feedback loops that characterize how the problem unfolds over time. Each feedback loop describes a single aspect of the problem. Multiple feedback loops interact to capture the complexities of the problem domain.

Appendix B: The Insider IT Sabotage Training Case¹²

INTRODUCTION

Chris, president of the computer systems sales company iAssemble, felt like he had just been hit by a Mack truck.

A partner in the company, Caroline, explained to him that something had just wiped out their system configuration and assembly programs. “And to top it off,” she continued, “the only backups were given to Ian before he was fired and we haven’t seen them since. Given the circumstances of Ian’s departure, we suspect that he might be responsible.”

“I just can’t believe that Ian would do something like that,” said Chris. “He’s been with the company since the beginning; he wrote most of those programs himself, for crying out loud!”

Chris paused and looked back at Caroline, “What the heck do we do now?”

BACKGROUND

iAssemble sold computer systems directly to customers, building each system to order and offering competitive prices. iAssemble had been doing extremely well and conducted an initial public offering (IPO) in 2001, after which its stock doubled.

Chris started the company in 1997 with his friend Caroline, who is now the Chief Technology Officer (CTO). The company has had success hiring experienced managers and employees since the beginning.

Ian was among the few employees who had been with iAssemble since its establishment. Ian started out as computer specialist and technical assistant to the two original founders, Chris and Caroline. When hired, Ian held certifications in personal computer (PC) hardware maintenance and operating system administration. Although he did not possess a college degree, with hard work and dedication, he became the sole system administrator at iAssemble, a position he held for four years. He also built the software that ran the computer system assembly machinery pretty much from the ground up. This software was the foundation for automating the PC assembly processes that allowed iAssemble’s rapid growth.

iAssemble grew at an increasing rate. Recognizing the need for qualified personnel, Chris and Caroline began to hire experienced system administrators who could also function as project managers. Lance was hired as lead system administrator because of his education and experience, and, James was hired as a junior system administrator to share Ian’s growing systems administration workload and responsibilities.

Ian was assigned to mentor James and ensure his smooth assimilation within the company. Ian and James worked on several projects together. Ian respected James’s abilities, finding him to be

¹² The iAssemble organization and case example are completely fictional, any resemblance to a real organization or insider threat case is unintentional.

nearly as technically competent as himself. The two of them got along fine, but James began to notice that Ian was becoming increasingly short-tempered over a period of several months.

One day after Ian's moodiness started annoying him, James decided to find out what the problem was. "Hey, what's bugging you Ian?" inquired James.

As if he had been waiting for someone to ask that question, he dumped on James. "James, I can't take it any longer. That idiot Lance thinks he knows how to run these networks better than me. I know these systems inside and out. The changes that he is suggesting will bring the network to its knees. They've got me on these piddling projects while they are destroying the foundations that I laid for iAssemble."

"You should be managing these networks, Ian," suggested James. "Why didn't they make you lead admin?"

"I have no idea, but who wants to be pushing papers all day long, anyway" Ian interrupted. "Lance is perfectly suited to that, but he doesn't know the first thing about running these networks. They simply don't appreciate what I do around here and some day they may just regret it."

Ian's disgruntlement grew and it became obvious in his hostile dealings with coworkers. He even bottlenecked projects on purpose on several occasions, stalling his work to ensure Lance and the team missed project milestones. Ian received a written warning from Caroline after several coworkers formally complained. Enraged by this, he had a heated argument with a team member who quit the very next day, citing Ian as the reason for his resignation. Caroline sent Ian a letter of final warning that put him on probation for his conduct—any more such problems would result in his immediate termination from iAssemble.

This seemed to resolve the situation, at least for a while. During the subsequent months, iAssemble continued to thrive. With a whopping 68 percent growth in sales over the previous quarter, iAssemble was forced to hire additional people. The staff adopted a "do whatever it takes" attitude to their job in order to keep up with the demands placed on them due to the growth. One staffer described it as follows:

We were one lean coding machine in those days. We had to be to extend the systems to support the company's amazing growth. Ian thought of and implemented the idea to centralize the core software on a central server to coordinate updates more efficiently. We made vast improvements to the flexibility and sophistication of the assembly programs over a very short period of time. And the extensions worked well with very few glitches. Of course, we had to cut some corners, giving people access when and where they needed it to make things happen. If something did not contribute to extending the systems, it just did not seem worth doing. This was how we were able to accomplish as much as we did.

Unbeknownst to management, during these months Ian was busy developing and testing a logic bomb that would delete all of the files on the central server. He did the testing and some of the development on his office desktop machine to make sure that it would really work. He also planted a backdoor account, with administrator privileges, on the main machinery server that provided him with unconstrained access from home, just in case he needed it.

Ian tried to get along with his coworkers, knowing that he would eventually get even. But he viewed most of his coworkers as incompetents; he just could not help "letting loose" on them

every once in a while. He had already started looking for another job, one where his abilities were recognized and valued, so he felt that he would not be around iAssemble much longer.

THE FINAL WEEKS

Management decided that it needed to deal with Ian's lingering performance problems. In a meeting with Chris and Caroline, Lance complained, "Ian is an arrogant jerk. He harasses and bullies his coworkers, treating them like they are dirt under his feet. Larry, the new programmer that we hired a few months ago, suspects that Ian is messing with the code that he is developing to make him look bad. Most of the staff walks on eggshells around him. We've got to do something."

"How big a hole is it going to leave in development and operations if we fire him?" asked Chris.

"Virtually none," replied Caroline, "with all of our hiring and aggressive training programs over the past few months, the rest of the staff is well up to speed on how things run and the directions that we are going. Both James and I think that we'd be a whole lot better off without him."

"OK," replied Chris, "Caroline, you take care of this yourself. Make sure to coordinate with James to be darned sure you cut off his access before you let Ian know. I feel bad about this—Ian has been with us since the beginning, but he has brought this on himself. So let's make it happen. How soon should we do it?"

"The sooner the better in my book." said Caroline, "I'll schedule to meet with him this morning. Lance, you disable his access while I'm meeting with him, and I'll have security escort him from the building after the meeting."

Chris's comment about disabling Ian's access had left Lance with concerns. Security practices at iAssemble had been less than rigorous lately, with the push to get the new software out. But Lance decided not to voice them at the meeting and later that day on July 10, 2001, Ian was fired, his access was disabled, and he was escorted from the building just as they had planned. Passwords for all shared accounts, including the system administrator accounts, were changed while Caroline was meeting with Ian.

Unfortunately, iAssemble managers were not aware that James had shared his password with Ian months earlier in order to make the development process easier. Ian went home the night he was fired and successfully logged into James's account. He then used his backdoor account on the machinery server to plant the logic bomb. He set it to go off one week later.

After the logic bomb detonated, Caroline was in Chris's office explaining that their critical software had been wiped out. Chris was puzzled as to how that was possible in light of the monitoring, policies, and security practices in place at iAssemble. After numerous hours of investigation, the system logs were used to trace the access of the machinery server to James's account. The evidence pointed to James as the saboteur. James claimed that he was not responsible for the deleted software and explained that he had given the password to his computer to Ian when they worked together. According to James, it had been such a long time ago that it had slipped his mind.


Management decided to call in law enforcement. Forensics analysis revealed that James's account was accessed remotely from Ian's home. Analysis of Ian's computer showed that he tested the

logic bomb four times over a period of three months. When questioned, Ian continued to maintain his innocence, even though the evidence against him was substantial. Investigation into Ian's background revealed that his father had been suffering from lung cancer over the last year and that he had recently lost his driver's license due to a conviction for driving under the influence of alcohol.

In a meeting with Caroline and Lance, Chris exploded. "We know who did it, but how do we recover from something like this? It will take months to recover operations even close to what we had. When this gets out, stockholders are going to demand a detailed explanation."

Slamming his fist on his desk, Chris demanded, "We must not only understand how this happened, but why, and make sure it does not happen again!"

Appendix C: Model of the Insider IT Sabotage Problem



References

- [Anderson et al. 2004]** Anderson, D.F.; Cappelli, D.M.; Gonzalez, J.J.; Mojta-hedzadeh, M.; Moore, A.P.; Rich, E.; Sarriegui, J.M.; Shimeall, T.J.; Stanton, J.M.; Weaver, E.; & Zagonel, A. "Preliminary System dynamics Maps of the Insider Cyber-Threat Problem." *Proceedings of the 22nd International Conference of the System dynamics Society*. Oxford, England, 2004.
<http://www.cert.org/archive/pdf/InsiderThreatSystemDynamics.pdf>.
- [Band et al. 2006]** Band, S.R.; Cappelli, D. M.; Fischer, L.F.; Moore, A. P.; Shaw, E.D.; & Trzeciak, R.F. *Comparing Insider IT Sabotage and Espionage: A Model-Based Analysis* (CMU/SEI-2006-TR-026). Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, 2006.
<http://www.sei.cmu.edu/publications/documents/06.reports/06tr026.html>.
- [Cappelli et al. 2006a]** Cappelli, D. M.; Desai, A. G.; Moore, A. P.; Shimeall, T. J.; Weaver, E. A.; & Willke, B. J. "Management and Education of the Risk of Insider Threat (MERIT): Mitigating the Risk of Sabotage to Employers' Information, Systems, or Networks." *Proceedings of the 24th International System dynamics Conference*. Nijmegen, Netherlands, 2006.
<http://www.albany.edu/cpr/sds/conf2006/proceed/proceed.pdf>.
- [Cappelli et al. 2006b]** Cappelli, D.M.; Moore, A.P.; Shimeall, T.J.; & Trzeciak, R.J. 2006b. "Common Sense Guide to Prevention and Detection of Insider Threats: Version 2.1." Pittsburgh, PA: CyLab and the Internet Security Alliance, Carnegie Mellon University, 2006.
<http://www.cert.org/archive/pdf/CommonSenseInsiderThreatsV2.1-1-070118.pdf>.
- [Keeney 2005]** Keeney, M.M.; Kowalski, E.F.; Cappelli, D.M.; Moore, A.P.; Shimeall, T.J.; & Rogers, S.N. *Insider Threat Study: Computer System Sabotage in Critical Infrastructure Sectors*. Pittsburgh, PA: Software Engineering Institute and U.S. Secret Service, Carnegie Mellon University, 2005.
<http://www.cert.org/archive/pdf/insidercross051105.pdf>.

- [Meadows et al. 1974]** Meadows, D. L.; Behrens, W. W.; Meadows D. H.; Naill, R. F.; Randers, J.; & Zahn, E. K. O. *Dynamics of Growth in a Finite World*. Cambridge, MA: Wright-Allen Press, Inc., 1974.
- [Melara et al. 2003]** Melara, C.; Sarriegui, J.M.; Gonzalez, J.J.; Sawicka, A.; & Cooke, D.L. "A System dynamics Model of an Insider Attack on an Information System." *Proceedings of the 21st International Conference of the System dynamics Society*. New York, NY, 2003.
- [Naumes & Naumes 1999]** Naumes, W. & Naumes, M.J. *The Art & Craft of Case Writing*. Thousand Oaks, California: SAGE Publications, 1999.
- [Rich et al. 2005]** Rich, E.; Martinez-Moyano, I.J.; Conrad, S.; Cappelli, D.M.; Moore, A.P.; Shimeall, T.J.; Andersen, D.F.; Gonzalez, J.J.; Ellison, R.J.; Lipson, H.F.; Mundie, D.A.; Sarriegui, J.M.; Sawicka, A.; Stewart, T.R.; Torres, J.M.; Weaver, E.A.; & Wiik, J. "Simulating Insider Cyber-Threat Risks: A Model-Based Case and a Case-Based Model." *Proceedings of the 23rd International Conference of the System dynamics Society*. Boston, MA, 2005.
- [Sterman 2000]** Sterman, J.D. *Business Dynamics: Systems Thinking and Modeling for a Complex World*. New York, NY: McGraw-Hill, 2000.
- [Yin 2003]** Yin, R.K. *Case Study Research, 3 ed.* Thousand Oaks, CA: Sage Publications, 2003.

REPORT DOCUMENTATION PAGE			<i>Form Approved</i> <i>OMB No. 0704-0188</i>	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE May 2008	3. REPORT TYPE AND DATES COVERED Final		
4. TITLE AND SUBTITLE The "Big Picture" of Insider IT Sabotage Infrastructures		5. FUNDING NUMBERS FA8721-05-C-0003		
6. AUTHOR(S) Andrew P. Moore, Dawn M. Cappelli, Randall F. Trzeciak				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Software Engineering Institute Carnegie Mellon University Pittsburgh, PA 15213		8. PERFORMING ORGANIZATION REPORT NUMBER CMU/SEI-2008-TR-009		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) HQ ESC/XPK 5 Eglin Street Hanscom AFB, MA 01731-2116		10. SPONSORING/MONITORING AGENCY REPORT NUMBER ESC-TR-2008-009		
11. SUPPLEMENTARY NOTES				
12A DISTRIBUTION/AVAILABILITY STATEMENT Unclassified/Unlimited, DTIC, NTIS		12B DISTRIBUTION CODE		
13. ABSTRACT (MAXIMUM 200 WORDS) A study conducted by the U.S. Secret Service and the Carnegie Mellon University Software Engineering Institute CERT Program analyzed 150 insider cyber crimes across U.S. critical infrastructure sectors. Follow-up work by CERT involved detailed group modeling and analysis of 30 cases of insider IT sabotage out of the 150 total cases. Insider IT sabotage includes incidents in which the insider's primary goal is to sabotage some aspect of the organization or direct specific harm toward an individual. This paper describes seven general observations about insider IT sabotage based on our empirical data and study findings. We describe a system dynamics model of the insider IT sabotage problem that elaborates complex interactions in the domain and unintended consequences of organizational policies, practices, technology, and culture on insider behavior. We describe the structure of an education and awareness workshop on insider IT sabotage that incorporates the previously mentioned artifacts as well as an interactive instructional case.				
14. SUBJECT TERMS Insider Threat Study, insider attack, MERIT, system dynamics, simulation-based learning, access control, information security, corporate security, risk management, critical infrastructure sectors, electronic crime, law enforcement		15. NUMBER OF PAGES 46		
16. PRICE CODE				
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	