

Chronological Examination of Insider Threat Sabotage: Preliminary Observations

William R. Claycomb, Carly L. Huth, Lori Flynn, David M. McIntire, and Todd B. Lewellen
CERT Insider Threat Center
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA
(claycomb, clhuth, lflynn, dmmcintire, tblewellen)@cert.org

Abstract

The threat of malicious insiders to organizations is persistent and increasing. We examine 15 real cases of insider threat sabotage of IT systems to identify several key points in the attack time-line, such as when the insider clearly became disgruntled, began attack preparations, and carried out the attack. We also determine when the attack stopped, when it was detected, and when action was taken on the insider. We found that 7 of the insiders we studied clearly became disgruntled more than 28 days prior to attack, but 9 did not carry out malicious acts until less than a day prior to attack. Of the 15 attacks, 8 ended within a day, 12 were detected within a week, and in 10 cases action was taken on the insider within a month. This exercise is a proof-of-concept for future work on larger data sets, and in this paper we detail our study methods and results, discuss challenges we faced, and identify potential new research directions.

Keywords: insider threat, sabotage, security

1 Introduction

An employee of a telecommunications company, when asked to resign, responded by sabotaging company IT systems, shutting down their telecommunication system and blocking 911 services in four major cities. A disgruntled former employee, upset that he was not hired for a full-time position, remotely accessed SCADA systems for a sewage treatment plant and caused over 200,000 gallons of raw sewage to spill into nearby rivers and businesses. Both of these cases highlight the devastating impact insider sabotage can have on an organization and society in general. Unfortunately, the problem is not infrequent: in a 2011 survey by CyberSecurity Watch [9], 43% of participating organizations stated that they had experienced at least one insider incident in the past year.

To examine the problem and potential solutions, we must first define the “insider threat.” We consider a malicious insider to be a current or former employee, contractor, or other business partner who has or had authorized access to an organization’s network, system, or data and intentionally exceeded or misused that access in a manner that negatively affected the confidentiality, integrity, or availability of the organization’s information or information systems. Our current research work focuses specifically on insider information technology (IT) sabotage, which we define as an insider’s use of IT to direct specific harm at an organization or an individual.

The issue of insider threats is by nature complex; it examines humans, organizations, IT systems, and the interactions between them. Proposed solutions for detecting insiders generally fall among three main categories: those detecting *technical* indicators of insider threat, *behavioral* indicators of insider threat, and *socio-technical* indicators of insider threat, which combines both technical and non-technical input for a more holistic viewpoint. For all categories, empirical analysis of existing instances of known malicious activity is a useful method. From known cases of insider threats, characteristics and traits about attackers and attacks can be examined to identify potential indicators and patterns of behavior to detect future insider activity. The process of examining actual insider threat cases involves several steps; one method is as follows:

1. Collect source data (e.g., documents, reports, etc.) on instances of insider crime
2. Process case information using a repeatable and consistent process to store key information and events about the case.
3. Create chronological time-lines from case data.
4. Identify key events in the chronology of the attack.
5. Examine case chronologies to identify patterns or significant indicators of attack.
6. Compare results to baseline behaviors of assumed good populations.

A documented, consistent, and repeatable method for each component listed above is critical to sound empirical analysis of malicious insiders. In this paper, we discuss an approach to steps 4 and 5: the identification of key events within the time-line of an insider attack and preliminary analysis of those events. We will be detailing the first three steps in a later work, as a larger research effort is currently underway to identify and measure early indicators of insider threat IT sabotage. For this part of the study, however, we examine the notion that sabotage is a relatively fast-acting type of insider crime. That is, the time between when an insider decides to attack and actually carries out the attack is a matter of days at most. We hypothesized that of the cases we analyzed, more than 50% would have a time frame of 7 days or less between when the insider begins malicious actions and when damage to the victim organization's IT systems occurs.

2 Related Work

2.1 Foundational Insider Threat Research

The foundations that influence much of the current research into insider threats within the United States arose in the 1990s. In 1991, Seger et al. addressed the insider threat to nuclear material, although this work can be more generally applicable to any staff with access to restricted information [29]. This work discussed prevention, detection, and response to the insider threat through a Personnel Security Assurance Program, including annual reviews of personnel and a reporting process for suspicious behaviors. Around the same time, the Department of Defense Personnel Security Research Center (PERSEREC) published a study on American spies [33]. This research has continued, focusing on espionage and insider activity [18, 12, 11]. In the late 1990s, RAND Corporation began to hold workshops on the topic of insider threats, to better understand the research needs in the area [3, 7, 2]. Also during this time, models were being developed that used different perspectives in order to elucidate information technology related insider activity. Shaw et al. published a study identifying psychological characteristics that may indicate an increased risk of harmful behavior in IT insiders [30]. An inclusive model described by Schultz took into account factors including personal characteristics, organizational authority and group support [28]. Other models developed during this time period include the Capability, Motive, Opportunity model, a criminological and social model, and a normative behavioral model [14, 22, 23].

2.2 Case Study Methodology

The methodology employed by many researchers in this area is a multiple case study approach, which can focus on technical, socio-technical, or behavioral elements of the insider threat. As described in Yin [34], a case study inquiry is defined as, "an empirical inquiry that investigates a contemporary phenomenon within its real-life context, especially when the boundaries between phenomenon and context are not clearly evident." A comparative case design (also known as multiple case design), had the advantage of often offering more compelling evidence and therefore being considered more robust than a single case study [34].

Maybury et al. analyzed six cases in order to create a model of behaviors and observables for three different types of insiders: analysts, applications administrators and system administrators [20]. This model was then tested in a networked environment where sensors were deployed to monitor at the packet, system, and application level. Maybury et al., tested the assumption that malicious insiders can be detected through observable actions, then examined three different mechanisms for analysis of the monitoring data, studying the solutions for both timeliness and accuracy of detection. Some multiple case studies provide a framework that address both technical and behavioral aspects. A combined technical and behavioral approach was undertaken by Randazzo et al., who studied 23 cases of insider events in the banking and financial sectors [24]. The study found that insiders were most often motivated by financial gain, tended to plan their actions in advance, and that they often used legitimate system commands to perpetrate their crime, requiring little technical sophistication. A behavioral multiple-case study approach was taken by Shaw & Fischer, who studied 10 cases of malicious insiders in critical infrastructures [31]. The study found similarities among the cases including the presence of organizational and personal stressors in each case. In eight of the ten cases studied, organizational over-dependence on the insider and lack of organizational ability to address concerning behaviors were found. Other perspectives on a technical and behavioral framework have been proposed, including studies by Schultz [28] and Greitzer [13].

2.3 CounterProductive Workplace Behaviors

Behavioral studies such as the one conducted by Shaw & Fischer support the broader research concerning counterproductive work behavior (CWB). Defined as intentional behaviors that are contrary to legitimate organizational interests, CWB include sabotage [25]. The topic has been the subject of extensive research, as discussed by Sackett and Devore who organized the precursors of CWB into groups: personality variables, job characteristics, work group characteristics, organizational culture, control systems and injustice [26]. In several studies on insider threat sabotage, including CERT's Management and Education of Risks of Insider Threat (MERIT) project and a CERT/PERSEREC studying comparing IT sabotage and espionage, the concept of CWB was reinforced through findings of personal predispositions and stressors as precursors of malicious events [8]. In the CERT/PERSEREC study on espionage and IT sabotage, personal predispositions were further divided into several categories: serious mental health disorders, personality problems, social skills and decision-making biases, and a history of rule conflicts [4]. However, many CWB studies use the Five Factor model which describes openness to experience, extraversion, conscientiousness, agreeableness, neuroticism or emotional stability as the dimensions of personality to be measured. While not universally accepted, CWB studies have suggested that irresponsible behaviors can be correlated with several of the factors, namely agreeableness, openness, and achievement [10, 16, 6]. The CERT/PERSEREC study also found that organizational and individual stressors also play a role in espionage and sabotage [4]. Stressors have also been correlated with CWB, for example Baron and Neuman positively correlated organization changes (e.g., paycuts, changes in management) with aggression [5]. Perceived variations in justice is another potential stressor, and has been linked to cases of sabotage [1]. Ambrose et al. studied 122 cases of sabotage, finding perceived injustice as the most common cause of sabotage [1].

2.4 Previous Sabotage Research

Previous research into insider sabotage has yielded both behavioral and technical findings. In 2005, a joint study with the U.S. Secret Service examined 49 cases of malicious insiders in critical infrastructure sectors within the United States [17]. Researchers completed a questionnaire after examining primary source material, such as investigative reports and court records, and interviewing case investigators and

victim organizations. The research resulted in the discovery of some commonalities between cases including:

- Insider’s actions were often preceded by a negative-workplace event, with revenge being the most frequently reported motive.
- A majority of insiders planned their activities in advance, and more than a quarter of the time others had information about their plans.
- A majority of insiders held technical positions.
- Most insiders acted out in a concerning manner in the workplace.
- Insider attacks often were carried out through compromised computer accounts, unauthorized back doors, or shared user accounts.
- A majority of insiders used remote access to carry out attacks, often outside of normal working hours.

Also addressing sabotage, CERT’s MERIT project focused on mitigating the risk of sabotage to an organization’s information, systems or networks [8]. The research applied system dynamics modeling to both technical and behavioral aspects of sabotage. Specifically, the research focused on developing a model of disgruntlement, illustrating both the insider’s expectation of freedom and the disgruntlement escalation due to organizational sanctions. The research also focused on modeling the insider’s attack, including acquiring unknown access paths as part of attack setup and attack escalation as insider disgruntlement increases. CERT then collaborated with PERSEREC to develop a model of espionage and compare it to the insider IT sabotage model [4]. Analysis yielded the following findings:

- Most saboteurs and spies had common personal predispositions that contributed to their risk of committing malicious acts.
- In most cases, stressful events, including organizational sanctions, contributed to the likelihood of insider IT sabotage and espionage.
- Concerning behaviors were often observable before and during insider IT sabotage and espionage.
- Technical actions by many insiders could have alerted the organization to planned or ongoing malicious acts.
- In many cases, organizations ignored or failed to detect rule violations.
- Lack of physical and electronic access controls facilitated both IT sabotage and espionage.

2.5 Coding Methodologies

As described by Yin, chronologies can be considered a “special form of time-series analysis” [34]. Chronologies allow the researcher to “trace events over time”, investigating earlier events that may have precipitated later events. The chronology can be compared to a theory which lays out conditions related to the sequence of events, time periods in which events take place, etc. After the chronology of events is created, the next step, as addressed in this paper, is to prepare the events for analysis through coding. Numerous coding schemas exist for qualitative research, many of which multiple phases of coding. In a review of coding processes, Saldana grouped them into first cycle, which deals with the initial coding, and second cycle, which addresses conceptualizing, prioritizing, and building a theory [27]. An often used example of the multi-phased coding process can be found in Strauss and Corbin’s open, axial, and selective coding as part of a grounded theory [32]. In the open coding process, events are scrutinized and compared to find common concepts. Next, axial coding develops groups of concepts and subcategories for some of the groups. Finally, selective coding is used to refine the groupings. This coding process is part of the development of a grounded theory in which, rather than viewing the data from the lens of an existing theory, a theory is “derived from data” which makes the theory “more likely to resemble

	Num. of cases	Average # of events per case
15 analyzed cases	15	30.0
Cases not analyzed	127	12.9
All sabotage cases	142	14.7

Table 1: Comparison of cases selected for analysis to source data set

the reality.” [32]. Concerns over reliability may arise when more than one individual is coding the data. Researchers have employed several techniques to address this concern. One technique is inter-rater reliability, in which more than one coder encodes each case and the level of agreement between coders is measured [19]. The technique of group consensus has also been employed to increase reliability, where differences are debated until agreement is reached [15].

3 Methods

Our approach explores the process of identifying key events in a chronological time-line of insider threat activity. Though important, the topics of case identification, gathering source material, and extracting chronological events are outside the scope of this paper. We describe our work based on events extracted from 15 cases of known insider sabotage of IT systems.

3.1 Our Coding Methodology

The data used for the study described in this paper was taken from a large database of actual cases of insider activity, covering the crimes of fraud, intellectual property theft, and sabotage. Case information was collected from both public sources such as court documents and non public sources such as law enforcement investigations and interviews with convicted insiders. Information was collected about the organizations involved, the perpetrator, and other details of the incident. With respect to the organizational data, information was collected such as the industry sector, work environment (e.g., layoffs, or mergers), and opportunity provided to the insider by the organizational action or inaction. The information collected on a convicted insider included demographic information, potential motives, concerning behavior, and violation history. Information was also collected about the perpetrator and organizational actions taken prior to the attack as well as any vulnerabilities exploited, detection of the attack, and the impact of the attack. The section of the database that is most relevant for this current work is the incident chronology. Each chronology contains a sequence of events including the date and time, place, and a detailed description of each of the known organizational and perpetrator actions starting with any information known prior to the attack up through any known legal adjudication.

3.2 Case Selection Method

15 cases of insider IT sabotage were chosen from over 130 previously collected cases. Additional descriptions of some of those cases, the case identification process, and results of previous studies can be found in [21, 17, 4, 8]. For this study, we selected cases based on a score calculated as the normalized sum of data availability (based on the overall number of observed elements, other than chronological events) and the number of events in the case chronology. The top-scoring 15 cases were used as our data set. While this selection method would skew any generalized results, we only intend to perform a proof-of-concept study at this point. Therefore we would like the richest set of data available to drive future hypothesis development. A comparison of the number of events for the chosen cases compared to

Seq.	Date	Event Description
1	7/22/2008	Insider lies on employment questionnaire
2	10/1/2008	Insider starts work at victim organization as assistant system admin.
3	6/16/2010	Insider reprimanded by manager for harassing coworker.
4	11/20/2010	Insider receives a below-average performance review rating overall.
5	1/18/2011	Insider is demoted and moved to smaller office.
6	1/18/2011	Insider threatens coworker that insider could 'mess with your user account and make you look really bad at work, if I felt like it'.
7	2/3/2011	11:54PM Insider installs a service and sets it to open a port during particular hours, known only to insider.
8	2/4/2011	9:05AM Insider fired by manager.
9	2/5/2011	11:03PM Insider connects to open port on file server at victim organization.
10	2/5/2011	11:20PM-11:40PM Insider deletes 5 work files from coworker #2.
11	2/18/2011	11:21PM Insider connects to open port on file server at victim organization.
12	2/18/2011	11:25PM Insider installs a logic bomb on the victim organization IT system.
13	2/19/2011	Victim organization recovers missing files from 2/5/2011 attack, but fails to detect logic bomb.
14	3/05/2011	Logic bomb executes, deleting critical configuration files from victim organization servers and deleting all user accounts.
15	3/07/2011	Victim organization IT system discovers the missing accounts when employees return to work and are unable to log in.
16	3/07/2011	Victim organization checks logs to find the unusual connections.
17	3/07/2011	Victim organization discovers missing configuration files.
18	3/10/2011	Victim organization calls in law enforcement.
19	4/10/2011	Law enforcement finds evidence of attack on insider's laptop.
20	6/15/2012	Insider found guilty of IT sabotage.

Table 3: Summary of Hypothetical Case of IT Sabotage Found in Appendix A

the overall data set is shown in Table 1. A subset of events from a hypothetical case chronology is shown in Table 3. The full series of case events is described in Appendix A.

3.3 Points of Interest

We identified several key points of interest among the events for each case we studied. These points denote significant events in the case, and are defined as follows:

Tippling Point (TP) The *TP* event is the first observed event at which the insider clearly became disgruntled. This point was particularly difficult to define operationally, as we lacked operational definitions of “disgruntled” and metrics for measuring “clearly.” However, as this is a preliminary study and not meant as a definitive work, we proceeded with this fairly subjective definition for the *TP* data point.

Malicious Act (MA) The *MA* event is the first observed event that clearly enabled the attack. This is meant to denote the point at which the insider put the attack in motion, but not necessarily the moment of impact. For example, this would include an insider testing a logic bomb or creating back-door accounts, in anticipation of a future attack.

Key event designator	Case event number
Tipping Point (<i>TP</i>)	5
Malicious Act (<i>MA</i>)	7
Attack Occurs (<i>OH</i>)	10
Attack Detected (<i>AD</i>)	15
Attack Ends (<i>AE</i>)	14
Action on Insider (<i>AI</i>)	18

Table 4: Key points associated with specific events in the example case shown in Table 3

Attack Occurs (OH) The *OH* event is the attack “zero hour.” This is the point where cyber damage actually begins. For example, the point where the insider deletes files, a logic bomb executes, or backup tapes are stolen (damaging data availability).

Attack Detected (AD) The *AD* event denotes the point where the organization realized something is wrong. This does not indicate the point at which the insider is identified as the attacker or even when the event is considered an attack. Rather, this is the point at which the organization begins to feel the effect of the attack. For instance, employees unable to log in, customers unable to access resources, or processes failing to run and automatically notifying system administrators.

Attack Ends (AE) The *AE* event is the point where cyber damage stops. It does not denote when recovery begins or ends. It could be very soon after the *OH* event (e.g., if the attack is simply to delete a few files.) Or, it could be hours, days, or even weeks later, if the insider remains undetected and continues to cause harm to the organization’s IT systems.

Action on Insider (AI) The *AI* event is the first instance of organizational response to the insider, such as the insider being fired, arrested, etc. This is the first *observed* event denoting a response. As our case data is based on publicly available sources, events such as the insider being fired are not always clearly denoted. While some might assume the insider was terminated, we did not record that as an event unless explicitly stated in the source material.

3.4 Method for Assigning Points of Interest

For each case, key points were associated with events based on group consensus among five internal insider threat researchers. When a consensus could not be reached among all five researchers, the key point was not assigned to any point in the case. The group decisions were not tested or validated with external researchers. Table 4 shows the results of applying our technique to the example case in Table 3.

4 Preliminary Results

We analyzed the points identified in each case to determine how close to the attack (*OH*) the insider reached the tipping point (*TP*) and began setting the attack in motion (*MA*). Additionally, we looked at attack duration, when it was detected (*AD*), and when action was taken on the insider (*AI*).

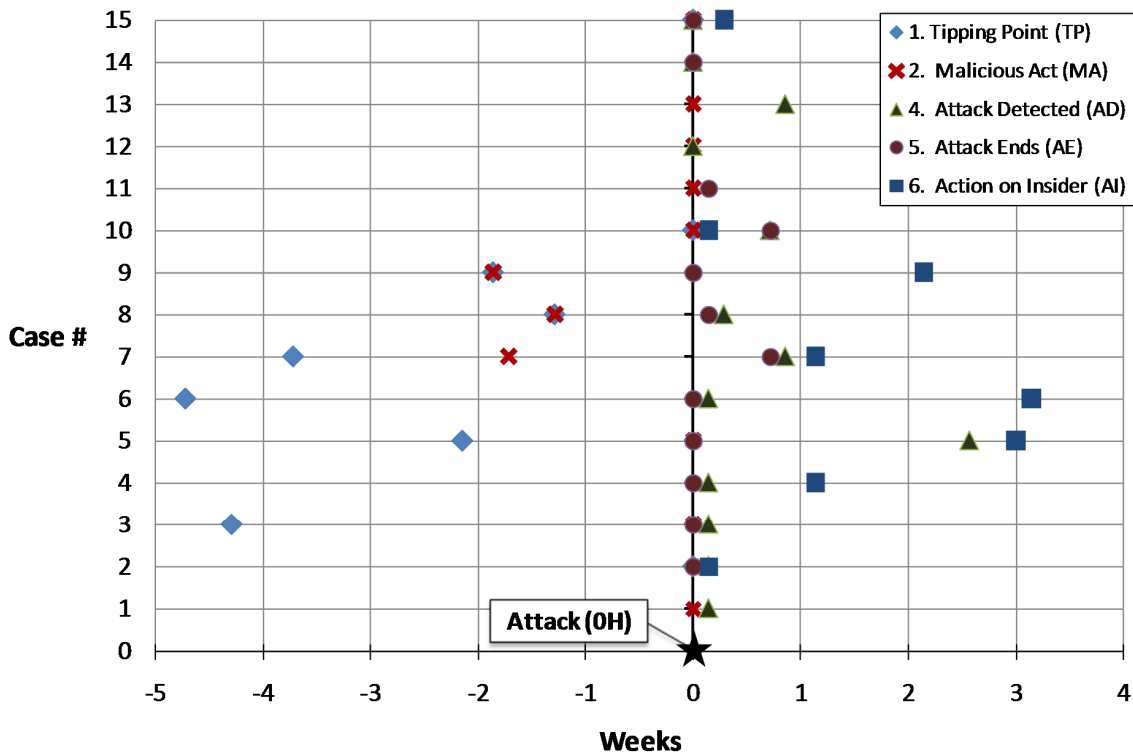


Figure 1: Timeline of key events relative to attack (OH).

4.1 Case Time-lines

We can make several observations by considering when key events occurred relative to OH , as shown in Figure 1, where $time = 0 = OH$. First, we note that 7 of 14 known TP events and 12 of 15 known MA events occurred fewer than 28 days prior to OH . 9 of 15 known MA events actually occurred less than 24 hour prior to attack, supporting our hypothesis that IT sabotage is often a crime where little time is spent planning and carrying out the attack. Table 5 shows a summary of these results.

> 28 days prior	7	3
8 – 28 days prior	4	3
1 – 7 days prior	0	0
< 1 day prior	3	9
Unknown	1	0
Time relative to attack (OH)	Tipping Point (TP)	Malicious Act (MA)

Table 5: Distribution of events prior to attack (OH)

Next, we note that post-attack events AD and AE typically occur within a week after attack. AD occurred within one day in 3 of 13 cases, and within a week in 12 of the 13 cases where these data points were known. AE was usually within a day (8 of 15 cases) to a week (12 of 15 cases), though some attacks lasted longer than 28 days (3 of 15 cases). AI occurred between 1 and 7 days after attack in 2 of 13 cases, between 8 and 28 days after attack in 5 of 13 cases, and more than 28 days after attack in 3

of the 13 cases where these data points were known. These figures are further detailed in Table 6. The maximum, minimum, mean, and median times (in days) relative to OH for each key event category is shown in Table 7.

> 28 days after	0	3	3
8 – 28 days after	1	0	5
1 – 7 days after	9	4	2
< 1 day after	3	8	3
Unknown	2	0	2
Time relative to attack (OH)	Attack Detected (AD)	Attack Ends (AE)	Action on Insider (AI)

Table 6: Distribution of events after attack (OH)

Maximum	0	0	18	116	152
Minimum	-305	-180	0	0	1
Mean	-55	-26	3	19	37
Median	-28	0	1	0	18
	Tipping Point (TP)	Malicious Act (MA)	Attack Detected (AD)	Attack Ends (AE)	Action on Insider (AI)

Table 7: Event occurrence relative to OH (days)

4.2 Insiders Past TP and MA

It could be of particular interest to security practitioners to take a closer look at when the insiders reached a critical point prior to attack (TP or MA). Figure 2 helps to illustrate this point, showing the percentage of all insiders past these key points prior to attack for the cases we studied. Note that one month prior to attack, nearly half of the insiders had passed TP , but only three had engaged in malicious acts related to the attack. One week prior to attack 11 of 14 insiders were past TP , but only 6 were past the MA event. This would suggest that detecting a potential insider threat’s tipping point would be much more effective in preventing attacks than simply monitoring for technical activity that may indicate a malicious event.

5 Discussion of Results

We noticed several interesting items during our study. Among them was the difficulty in developing operational definitions for each key point in an insider’s timeline. For instance, consider case 6 in Appendix B, where MA (-63 days) occurs before TP (-33 days). It seems counterintuitive that an insider would begin committing malicious acts before reaching the tipping point of deciding to attack. In this particular case, the insider covertly moved several key software applications from distributed servers to one centralized server 63 days prior to attack. Though this action was unauthorized, we do not believe the insider intended to attack at that point (i.e., the insider had not reached TP). Subsequent events, including demotion and downsizing, caused the insider to reach TP 30 days later, and the original act of centralizing key applications enabled the insider to easily delete all of them at once during the attack.




Figure 2: Percent of Insiders Past Key Points Prior to Attack

Our operational definition for *MA* is “the first observed event that clearly enabled the attack,” so in this case we concluded that *MA* did in fact occur prior to *TP*.

Defining the point of attack (*OH*) was also challenging. For example, consider insiders who test and deploy malicious code designed to execute the attack at some later point in time (“logic bombs.”) Given the intention to attack, should the act of testing a logic bomb be considered the beginning of the attack? Or perhaps the act of placing the logic bomb on corporate systems prior to execution should be considered the attack? One of our goals was to study observable events that could have alerted an organization to danger before an insider strikes, so we decided the attack zero-hour (*OH*) should be defined as the point at which cyber damage (i.e., affecting the confidentiality, availability, and/or integrity of corporate data) begins.

Tipping point (*TP*) was the most difficult point to define and identify in the cases we studied. It is a very subjective data point, and we must note that our results are based solely on the data available, which often included very few events describing interactions and behavioral events of the insider prior to attack. We felt it was important to attempt to define this key event, however, as previous work has noted that behavioral precursors (i.e., *TP* to sabotage generally occur prior to technical precursors, and we wanted to examine that time difference for these cases, if possible. [21].)

5.1 Caveats on Results

The 15 cases we studied were chosen from over 130 cases of insider threat sabotage (see Section 3.2 for case selection criteria.) Those cases are part of a larger data set of over 800 cases of insider crimes collected over the last eleven years. Most of the information in the larger data set was collected from publicly available sources, such as court records or media articles, though some of the case information was collected from non-publicly available sources, such as investigator reports or subject interviews. An overwhelming majority of the insiders documented in these cases were convicted of crimes related to their malicious insider activity.

It is important to note several caveats regarding use of this type of data. First, this is a qualitative

study, and due to the sampling method used, is not generalizable. That is, the set of insiders represents only those whose crimes were reported by the victim organization. As the 2011 CyberSecurity Watch survey notes, 76% of insider cases are handled without legal action or law enforcement, so the available case data is already limited to approximately 24% of known insider activity [9]. When cases are reported by victim organizations, it is often due to the magnitude of the crime, either in terms of impact to the organization or the number of people outside the organization who are affected. So there is a high likelihood that cases reported externally are not entirely representative of all known insider crimes. Furthermore, it is unclear how many insider attacks are undetected, or are detected but not attributed to insider activity.

Another caveat to note is the limitation on data collection. The data-set we used is limited by the scarcity of detailed source materials. Not only do these source materials differ in terms of credibility (e.g., media reports vs. expert testimony), but they also tend to lack the technical details necessary for in-depth analysis of the nuances of insider attacks on IT systems. Additionally, the interpersonal relationships that are critical to helping identify potential insider threats are also often overlooked by the source material available.

6 Conclusion and Observations for Future Work

Several observations were made during this study that will guide future work. The need for unambiguous operational definitions of various elements of case analysis is essential for high-confidence results. But even with clearly defined events, examining chronologies of multiple cases can be difficult without consistent procedures for extracting chronological information from source materials. Even the subtle difference between “Insider hired by the victim organization” and “Victim organization hired insider” can affect how analysts perceive certain data points, and makes automated data extraction more difficult. A repeatable method for extracting case events is needed. Furthermore, once events have been entered into a chronology, detailed analysis is difficult while events remain in free-text. A process for describing events using a discrete and finite set of descriptors would also be very helpful. Finally, our case selection procedure was somewhat limited (amount of data available and the number of events in the case). A more rigorous case selection methodology, resulting in higher confidence results might instead be based on attributes such as quality of case data, number of source documents, confidence in source documents, corroboration of source documents, and level of detail of events in the case.

Despite the limitations, we believe our study revealed interesting characteristics of insider threat sabotage that will guide future work. Notably, we saw that among the cases examined almost half of the insiders clearly became disgruntled more than four weeks prior, but more than half did not commit malicious acts until one day or less prior to actual attack. Attacks were generally over quickly, detected within a week, and some form of action on the insider occurred within a month. Again, these results should not be generalized, but the methods we used may guide other researchers performing similar studies. As we understand more about how insiders behave, we can come closer to developing effective techniques in identifying and stopping potential threats before they attack.

Acknowledgments This material is based upon work funded and supported by the Department of Defense under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

References

- [1] M. L. Ambrose, M.A. Seabright, and M. Schiminke. Sabotage in the workplace: The role of organizational injustice. *Organizational Behavior and Human Decision Processes*, 80, 2002.
- [2] R.H. Anderson. Research and development initiatives focused on preventing, detecting, and responding to insider misuse of critical defense information systems (RAND CF-151-OSD). Technical report, RAND Corporation, 1999.
- [3] R.H. Anderson, T. Bozek, T. Longstaff, W. Meitzler, M. Skroch, and K. Van Wyk. Research on mitigating the insider threat to information systems #2 (RAND CF-163-DARPA). Technical report, RAND Corporation, 2000.
- [4] S. R. Band, L. F. Fischer D. M. Cappelli, A. P. Moore, E.D. Shaw, and R.F. Trzeciak. Comparing insider IT sabotage and espionage: A model-based analysis. Technical report, Software Engineering Institute, 2006.
- [5] R.A. Baron and J.H. Neuman. Workplace violence and workplace aggression: Evidence on their relative frequency and potential causes. *Aggressive Behavior*, 22, 1996.
- [6] J. Block. A contrarian view of the five-factor approach to personality description. *Psychological Bulletin*, 117, 1995.
- [7] R.C. Brackney and R. H. Anderson. Understanding the insider threat (RAND CF-196-ARDA). Technical report, RAND Corporation, 2004.
- [8] D. Cappelli, A. Desai, A. Moore, T. Shimeall, E. Weaver, and B.J. Willke. Management and education of the risk of insider threat (MERIT): System dynamics modeling of computer system sabotage. Technical report, Carnegie Mellon University/Software Engineering Institute - CERT, 2006.
- [9] CSO Magazine, US Secret Service, Software Engineering Insistute CERT Program at Carnegie Mellon University, and Deloitte. 2011 cybersecurity watch survey. *CSO Magazine*, January 2011.
- [10] H. J. Eysenck. Four ways five factors are not basic. *Personality and Individual Differences*, 13, 1992.
- [11] L. F. Fischer. Characterizing information systems insider offenders. In *Proceedings of the 1993 Annual Conference of the Intl. Military Testing Association*, 2003.
- [12] L.F. Fischer, J.A. Riedel, and M. Wiskoff. A new personnel security issue: Trustworthiness of defense information systems insiders. In *Proceedings of the 42nd Intl. Military Research Association Conference*, pages 333–339, 2000.
- [13] Frank L Greitzer and Ryan E. Hohimer. Modeling human behavior to anticipate insider attacks. *Journal of Strategic Security*, IV(2):25–48, 2011.
- [14] T.M. Gudaitis. The missing link in information security: Three dimensional profiling. *CyberPsychology and Behavior*, 1, 1998.
- [15] B. Harry, K.M. Sturges, and J. K. Klinger. Mapping the process: An exemplar of process and challenge in grounded theory analysis. *Educational Researcher*, 34, 2005.
- [16] L.M. Hough. The ‘big five’ personality variables-. construct confusion: Description versus prediction. *Human Performance*, 5, 1992.
- [17] M. Keeney, E. Kowalski, D. Cappelli, A. Moore, T. Shimeall, and S. Rogers. Insider threat study: Computer system sabotage in critical infrastructure sectors. Technical report, Carnegie Mellon University/Software Engineering Institute, 2005.
- [18] K.Herbig and M.F. Wiskoff. Espionage against the united states by american citizens 1947-2001. Technical report, Defense Personnel Security Research and Education Center (PERSEREC), 2002.
- [19] M.G. Maxfield and C.S. Widom. The cycle of violence: Revisited six years later. *Archives of Pediatrics and Adolescent Medicine*, 150, 1996.
- [20] M. Maybury, P. Chase, B. Cheikes, D. Brackney, S. Matzner, T. Hetherington, B. Wood, C. Sibley, J. Marin, T. Longstaff, L. Spitzner, J. Haile, J. Copeland, and S. Lewandowski. Analysis and detection of malicious insiders. In *Proceedings of the 2005 Intl. Conference on Intelligence Analysis*, 2005.
- [21] A. Moore, D. Cappelli, and R. Trzeciak. The “big picture” of insider IT sabotage across U.S. critical infrastructures. Technical report, Carnegie Mellon University/Software Engineering Institute, May 2008.
- [22] J. Morahan-Martin. Women and girls last: Females and the internet. In *In Proceedings of the IRISS Confer-*

- ence., 1998.
- [23] D.B. Parker. *Fighting Computer Crime: A new framework for protecting information*. John Wiley and Sons, 1998.
- [24] M.R. Randazzo, M. Keeney, E. Kowalski, D. Cappelli, and A. Moore. Insider threat study: Illicit cyber activity in the banking and finance sector. Technical report, Carnegie Mellon University/Software Engineering Institute - CERT, 2004.
- [25] P.R. Sackett. The structure of counterproductive work behaviors: Dimensionality and relationships with facets of job performance. *Intl. Journal of Selection and Assessment*, 10, 2002.
- [26] P.R. Sackett and C.J. DeVore. *Handbook of Industrial, Work and Organizational Psychology*, volume 1, chapter 8: Counterproductive Behaviors at Work. SAGE Publications Ltd, 2002.
- [27] J. Saldana. *The Coding Manual for Qualitative Research*. SAGE, 2009.
- [28] E.Eugene Schultz. A framework for understanding and predicting insider attacks. *Computers & Security*, 21(6):526 – 531, 2002.
- [29] K.A. Seger, T.E. Childress, and G.R. Eisele. Personnel security assurance program: Countering the insider threat (CRISE 92-G-18). Technical report, Oak Ridge, Tennessee: Center for Human Reliability Studies, 1991.
- [30] E. Shaw, K.G. Ruby, and J.M. Post. The insider threat to information systems: The psychology of the dangerous insider. *Security Awareness Bulletin*, 2, 1998.
- [31] E.D. Shaw and L.F. Fischer. Ten tales of betrayal: The threat to corporate infrastructure by information technology insiders. Technical report, Defense Personnel Security Research Center (PERSEREC), 2005.
- [32] A. Strauss and J. Corbin. *Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory*. SAGE, 2 edition, 1998.
- [33] S. Wood and M.F. Wiskoff. Americans who spied against their country since World War II. Technical Report PERS-TR-92-005, Defense Personnel Security Research and Education Center (PERSEREC), 2002.
- [34] R.K. Yin. *Case Study Research: Design and Methods*. Sage, 4th edition, 2009.

A Sample Case Chronology

The following is an example of a case chronology similar to those analyzed in this study. Due to confidentiality agreements with data providers, this does not represent an actual case, but is modeled on actual insider activity.

Seq	Date	Event
1	1/25/2006	Insider found guilty on assault charge.
2	2/15/2006	Insider sentenced to probation on assault charge.
3	2/16/2007	Insider's probation period ended.
4	7/22/2008	Insider lies on employment questionnaire: writes "no criminal convictions".
5	9/30/2008	Organization does not do a pre-employment check to verify prior record.
6	10/1/2008	Insider starts work at victim organization as assistant system admin.
7	8/2/2009	Coworker #2 verbally complains to manager about harassment from insider.
8	11/15/2009	Insider receives a below-average performance review rating specific to teamwork, but otherwise above-average performance review rating.
9	3/20/2010	Coworker #1 observes insider breaking company rule by copying files from the insider's smartphone to the work system, but doesn't report it even though coworker knows about the rule.
10	6/15/2010	Insider alters coworker #2's system to show offensive images on the desktop.
11	6/16/2010	Insider reprimanded by manager.
12	11/20/2010	Insider receives a below-average performance review rating overall.

13	1/18/2011	Insider is demoted and moved to smaller office.
14	1/18/2011	Insider threatens coworker #3 that insider could 'mess with your user account and make you look really bad at work, if I felt like it'.
15	1/18/2011	Coworker #3 tells coworker #4, but asks coworker #4 not to tell anyone else.
16	2/3/2011	10:16PM Insider logs into a shared workstation using coworker #3's userID and password (method of acquisition unknown).
17	2/3/2011	10:20PM Insider logs out without doing anything else.
18	2/3/2011	10:45PM Coworker #5 notices supposed coworker #3's login, from a login message on the shared machine. Knows it was done by the insider, because knew coworker #3 was gone but saw insider alone in room at that time.
19	2/3/2011	Coworker #5 tells insider that the insider's falsified login will be reported.
20	2/3/2011	11:54PM Insider installs a service and sets it to open a port during particular hours, known only to insider.
21	2/3/2011	Coworker #5 reports the falsified login to insider's manager.
22	2/4/2011	9:05AM Insider fired by manager.
23	2/5/2011	11:03PM Insider connects to open port on file server at victim organization.
24	2/5/2011	11:20PM-11:40PM Insider deletes 5 work files from coworker #2.
25	2/9/2011	8:52PM Insider connects to open port on file server at victim organization.
26	2/9/2011	9:01PM Insider modifies a presentation of insider's manager meant for customers, to make it look like the organization was performing poorly on the job for the customer.
27	2/18/2011	11:21PM Insider connects to open port on file server at victim organization.
28	2/18/2011	11:25PM Insider installs a logic bomb on the victim organization IT system.
29	2/19/2011	Victim organization recovers missing files, but fails to detect logic bomb.
30	3/05/2011	Logic bomb executes, deleting critical configuration files from victim organization servers and deleting all user accounts.
31	3/07/2011	Victim organization IT system discovers the missing accounts when employees return to work and are unable to log in.
32	3/07/2011	Victim organization checks logs to find the unusual connections.
33	3/07/2011	Victim organization discovers missing configuration files.
34	3/07/2011	Victim organization calls in law enforcement.
35	3/08/2011	Using backup tapes, victim organization restores user accounts and files.
36	4/10/2011	Law enforcement investigation finds evidence insider's laptop was used for the sabotage acts.
37	6/15/2012	Insider found guilty of IT sabotage.

B Complete Results

Table 9 shows the number of days before or after attack (*OH*) for each key event identified in all cases analyzed during this study.

Case	Tipping Point (<i>TP</i>)	Malicious Acts (<i>MA</i>)	Attack Begins (<i>OH</i>)	Attack Ends (<i>AE</i>)	Attack Detected (<i>AD</i>)	Action on Insider (<i>AI</i>)
1	-79	0	0	116	1	152
2	0	0	0	0	1	1
3	-30	0	0	0	1	59
4	-305	-180	0	0	1	8
5	-15	0	0	0	18	21
6	-33	-63	0	0	1	22
7	-26	-12	0	5	6	8
8	-9	-9	0	1	2	N/A
9	-13	-13	0	0	N/A	15
10	0	0	0	5	5	N/A
11	-69	0	0	1	N/A	N/A
12	-62	0	0	74	0	N/A
13	N/A	0	0	76	6	80
14	-124	-113	0	0	0	N/A
15	0	0	0	0	0	2

Table 9: Number of days before or after attack for each key event

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.

This material has been approved for public release and unlimited distribution.

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 and 252.227-7013 Alternate I.