A Method For Characterizing Sociotechnical Events Related to Insider Threat Sabotage

William R. Claycomb* and Carly L. Huth
CERT® Insider Threat Center
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA
{claycomb, clhuth}@cert.org

Abstract

Analyzing historical cases of insider crimes to identify patterns or specific indicators of attack is a challenging task, particularly when using large volumes of free-text input sources, such as court documents and media reports. In this workshop paper, we offer a new process for processing, or *coding*, free-text descriptions of insider crimes for future analysis; specifically, we study cases of insider threat sabotage. Our method is based on a triad of discrete descriptors which allow for a quick, accurate, and repeatable characterizations of any event in the timeline of an insider attack. While the majority of this paper is concerned with reporting our development efforts and describing the current state of the project, we will briefly address some initial findings based on analysis conducted on the results of our coding efforts. In general, we found our new method increased the ease with which analysts could distinguish between technical events (those involving IT systems) and behavioral events (individual or interpersonal events not involving IT systems). Also, this coding technique also allowed for consistent comparison of events across cases. For instance, from 49 cases of insider threat sabotage, we determined that the majority had behavioral events prior to technical events, indicating a potential area for further study.

Keywords: insider, security, sabotage

1 Introduction

One of the most difficult cyber crimes to address is an act committed by a trusted insider. Many definitions of the "insider threat" exist [1]; for the purpose of our previous and current studies, we define a malicious insider threat to an organization as "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." Addressing the insider threat issue is challenging not only because attacks usually happen within the constraints of authorized system and data access, but also because attacks often involve combinations of events observable on IT systems (*technical* events), as well as those not observable on IT systems (non-technical or *behavioral* events.) Technical events include logon/logoff, data access, read/write/execute, etc. Behavioral events include social interactions between people and organizations.

This combination of social and technical, or *sociotechnical*, data related to the crime presents both challenges and opportunities for insider threat research. Challenges include identifying and observing behavioral events that indicate increased risk of insider attack - doing the same for technical events is challenging enough! Opportunities arise, however, because behavioral events often provide critical

Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications, volume: 4, number: 4, pp. 1-19
*Corresponding author: 4500 Fifth Avenue, Pittsburgh, PA 15213-2612 USA, Tel: +1-412-268-5800

context necessary to correctly identify technical events as malicious or not. For example, a system administrator modifying a backup script may not seem malicious under normal circumstances, but when viewed in light of a concurrent conversation about wanting to "get back at the company," and including the phrase "wait until you see the fireworks this weekend," script modification suddenly becomes much more concerning. This particular example is from one of over 150 cases of insider IT sabotage we have collected for analysis, and fortunately, the insider's attempt to sabotage the company was unsuccessful.

Therefore, we believe a sociotechnical approach is critical to identifying indicators of malicious insider activity, because it provides a more holistic view point of the insider, the organization, and the crime. For each case we analyzed in this study, we used publicly available information about the crime to create a chronological timeline which allowed us to "trace events over time" [2]. In a recent study of information technology (IT) sabotage cases, it was noted that one difficulty with this methodology is that chronological timelines are free-text, making a detailed analysis of the chronology difficult [3]. To address this issue, we first attempted to adapt an existing mechanism that applies numerical codes to each event, depending on how the event fits into a hierarchical coding structure. Our effort even including developing an automated process for coding each event. However, we found that the coding process we developed was too complex and unrefined to be used in the current study. Our next step was to attempt to simplify the coding process by developing a discrete and finite set of descriptors to use in characterizing each chronological event. Because this set of events included three components, the *actor*, the *action*, and the *target*, we refer to it as the *triad*. We tested our method by using the triad to characterize 49 actual cases of IT sabotage. We found that it allowed us to more easily characterize the types of events, particularly when distinguishing between behavioral or technical events.

The primary goal of this paper is to describe the triad coding process development and application using insider threat as a use case; the process could be applied to other types of empirical event-based analyses. We will also briefly share some initial results that followed from our analysis of cases of insider IT sabotage. Specifically, we found that behavioral events preceded technical events in 85% of cases where both types of events were present. Also, we noted that the most common insider-initiated actions included absenteeism, abuse or intimidation of coworkers, and installation of software on victim organization IT systems. The most common actions by the victim organization included denying requests from the insider (i.e. for access, permanent employment) and expressing concern about the insider.

2 Background and Related Work

Previous research has been conducted to study insider IT sabotage, defined by Band, et. al as "an insider's use of IT to direct specific harm at an organization or an individual." [4]. These studies provide some initial insight into potential early indicators of insider IT sabotage. A U.S. Secret Service 2005 study found that a majority of insiders planned their activities in advance, and many also acted out in a concerning manner in the workplace [5]. In 2006, the CERT Program's MERIT project studied sociotechnical aspects of sabotage, focusing on modeling insider disgruntlement and steps in the attack. The CERT Program also collaborated with PERSEREC to compare of espionage and IT sabotage [6]. The study found that concerning behaviors were often visible before and during the crime and technical actions could have alerted the organization to the crime. Finally, a 2012 preliminary study looked at 15 cases of IT sabotage, coding key points in the chronology including insider disgruntlement, attack preparation and time of attack. This study found indications that sabotage may be a very fast-acting crime, and that many insiders became disgruntled well before they carried out the attack [3].

We chose the cases for this study from an internal database of over 800 real-world insider threat cases. The database contains two other types of crimes in addition to insider IT sabotage, specifically, fraud and theft of intellectual property. Each case includes information about the perpetrator, organiza-

Table 1: Sample Insider Sabotage Case Chronology

Sequence	Date	Event
1	10/1/2008	Insider starts work at victim organization as an assistant system administrator.
2	11/30/2010	Insider receives a below-average performance review rating overall.
3	1/18/2011	Insider threatens co-worker, saying that insider could "mess with your user account and make you look really bad."
4	2/3/2011	10:16 PM Insider logs into a shared workstation using co-workers userID and password.
5	2/3/2011	10:20 PM Insider exploits security flaw to elevate privileges.
6	2/3/2011	10:22 PM Insider deletes critical project files.
7	2/4/2011	Insider fired.

tions involved, and incident details. This information is derived mainly from public sources (e.g. media reports and court documents) but also includes some data from non-public sources (e.g. law enforcement investigator notes). This approach, known as a comparative case design, has been employed by insider threat researchers studying technical, behavioral, and sociotechnical aspects of insider threat [7]. However, limitations exist, primarily due to the difficulty in obtaining real-world data. For instance, Greitzer and Paulson note that researchers lack real-world data to validate their approaches. A recent insider threat literature review by Hunker and Probst also called out the need for real-world data in analysis [8].

Patterns among IT sabotage cases have been the topic of previous research efforts, including a 2005 joint study with the U.S. Secret Service, the CERT Program's MERIT project, and a collaboration with PERSEREC to compare espionage and sabotage [9, 5, 6]. These studies provide some initial insight into potential early indicators of IT sabotage. In the 2005 joint study, findings included the discovery that a majority of insiders planned out their activities in advance and acted out in a concerning way in the workplace [5]. The CERT Program and PERSEREC's comparison of espionage and IT sabotage yielded several observations, including that concerning behaviors were often observable before insider IT sabotage and espionage [6]. This study also noted that insider technical actions could have alerted the organization to planned or ongoing malicious events.

3 Methodology

3.1 Case Chronologies

This study selected 49 cases of malicious insider IT sabotage from the larger database. These cases represented those with the greatest number of chronological events and diversity of information sources from the more than 140 available cases of IT sabotage. Each chronology contained a sequence of discrete events including (when known) the date, time, place of the event, as well as a detailed description of the known organizational and perpetrator actions, starting with any information known prior to the attack up through any known legal adjudication.

While chronologies were already populated for some cases, initial examination of event contents revealed issues that hindered development of a repeatable analysis methodology. For instance, a single event entry sometimes described more than one distinct action by the subject of the entry, for example, "The insider logged on to the system and developed a logic bomb." In this example, "logged on" and

"developed a logic bomb" are two distinct actions. Therefore, prior to analysis of chronological events, two trained analysts examined the existing chronologies, modifying them if necessary, to ensure that each event contained only one action.

In addition, some effort was made to streamline and regularize the free-text description of each event. For example, the 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 [3]. To mitigate this issue, two trained analysts developed a standardized method for expressing common events such as hiring, firing, damage, and use of drugs. Table 1 provides an example of a hypothetical partial chronology sequence.

3.2 Adapting Previous Codebook

All of the cases we studied had been entered into our database according to procedures in the CERT Insider Threat Database Codebook [10], which contains codes representing data about the victim organization, the insider, and the incident, as well as coding and quality control guidance. However, we found the existing codebook lacked the expressiveness needed for in-depth analysis of the events prior to attack, so we determined that a new coding structure should be developed for this study. We initially created a new coding structure following a critical pathway model of malicious insider behavior, first proposed by Shaw and Fischer [11], which describes how insiders often harbor personal predispositions that contribute to their risk of committing attack when faced with stressful events. Under such conditions, these individuals may be more likely to exhibit concerning behaviors that either precede or are part of an insider attack. Our effort to capture this structure in our coding techniques included both the reorganization of existing event codes into a new hierarchical structure, as well as the addition of new event codes to describe additional technical and behavioral observables that researchers hypothesized might be present in the chronologies. The highest level categories were descriptive of events related to the following: personal predispositions, stressful events, concerning actions, the incident, the organization, and the response. Three additional sub-levels were added to each category to provide granularity, for example, "2.1.7.3: Poor Performance Review" is part of category 2 - "Stressful Events", and subcategories 2.1 ("Work-related Stressful Events") and 2.1.7 ("Negative Employment Events").

We performed an initial pilot study to test the effectiveness of this new coding procedure; specifically aiming to measure the usefulness of our operational definitions for specific event codes. In this test, analysts were given an updated codebook, including definitions of the codes and examples for each code. Analysts were asked to identify appropriate codes for several chronology sequences. Agreement among analysts was low due to three significant observations:

- The operational definitions for event codes were ambiguous.
- The actual chronology items lacked a standard structure, which led to different interpretations of the same events across multiple cases.
- At 16 pages, the new coding guide was too lengthy and lacked the structure necessary to lead analysts to the same codes consistently.

To address the last point, in an effort to aid the analysts, we began development of an automated, streamlined analysis process, with the goals of reducing cognitive load [12] and increasing agreement between analysts. We first developed a set of decision trees based on the hierarchical structure of the codebook. For example, the analyst would be asked, "Does this event indicate something that was known about the insider previous to employment?" If the analyst answered "yes," the analyst would be guided to codebook items addressing pre-employment information. A pilot test was conducted using

this decision tree approach. While initial results were promising, development and testing of a full-scale implementation, covering all codes required for this project, was outside the scope of time and resources available. We concluded that our new codebook was simply too complex and unrefined to use in an accurate and timely manner, though future work includes creating the actual automated query process.

Based on the lessons learned during this first attempt, we determined that we needed a more efficient way of parsing event data as an initial gateway into analysis. To address this, we developed a three-component structure, dubbed the "triad," for describing incident events. This method is loosely based on a method for describing events related to foreign policy described in East and Herman [13]. Rather than assign a unique code to each event, we simply focused on accurately describing the contents of the event. To put it plainly, we wanted to know very clearly and simply *what happened* during each event described in a case. We believed that focusing more on understanding semantics rather than determining specific codes would yield more consistent results than our first attempt, and we anticipated that with the semantics of events determined, specific coding would be much easier - perhaps even automated.

3.3 The Triad

The first step was for a group of senior analysts to develop an initial set of possible choices for each component, organizing them into a two- or three-level hierarchical structure. The goal of this activity was to develop a finite and discrete list of descriptors that adequately and accurately describe all possible actors, targets, and actions represented in the cases we studied. However, we did not want this set of descriptors to be so detailed that analysts would face the same challenges they faced during our first attempt to codify events. To this point, we reasoned that the sets of possible actors and targets are likely identical, and that determining the actor and target for each event automatically determines the type of event that occurred. Analysts would then be asked to apply an event descriptor consisting of an actor, action, and target to each chronological event, as illustrated in Figure 1. In our study, we developed the triad related to insider IT sabotage, but the same process could be applied to other domains as well.

3.3.1 Actors and Targets

Because our cases involved insider sabotage of IT systems, we chose the following top-level categories to represent both the actor and target of an event: *IT, NOT IT, Unknown,* and *Null*. The *IT* category represents IT systems, which the *NOT IT* category represents everything else, such as people and organizations. The *Null* category is used for events without a target, such as life events (i.e., individual feelings, change in location, changes in financial situation, changes in health, change in family status). At the second level, we broke the *IT* category down into components such as the IT systems of the victim or other organizations, the insider's personal IT assets, and the internet. *NOT IT* sub-categories included the insider, the victim organization, other outside entities, and the criminal justice system.

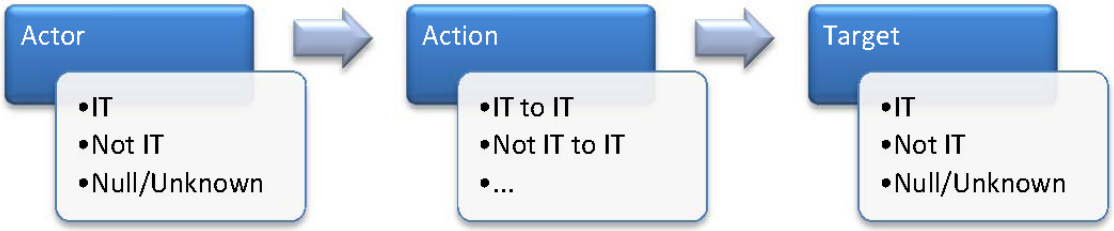

Figure 1: The Triad

Event Seq.	**Actor Level 1**	**Actor Level 2**	**Actor Level 3**	**Action Level 2**	**Action Level 3**	**Target Level 1**	**Target Level 2**	**Target Level 3**
1	NOT IT	Victim Org	Organi-zation/ Group	Inter-personal	Hired	NOT IT	Insider	Insider Individual
2	NOT IT	Victim Org	Super-visor	Inter-personal	Commun-icated	NOT IT	Insider	Insider Individual
3	NOT IT	Insider	Insider Individual	Inter-personal	Bullied/ Intimidated/ Threatened	NOT IT	Victim Org	Co-worker
4	NOT IT	Insider	Insider Individual	Logged On	N/A	IT	Victim Org IT	N/A
5	NOT IT	Insider	Insider Individual	Exploited	N/A	IT	Victim Org IT	N/A
6	NOT IT	Insider	Insider Individual	Deleted	N/A	IT	Victim Org IT	N/A
7	NOT IT	Victim Org	Organi-zation/ Group	Inter-personal	Fired	NOT IT	Insider	Insider Individual

Table 2: Triad Values for the Example Event Shown in Figure 1

3.3.2 Actions

Knowing the actor and target types identified the type of event represented, for instance *NOT IT* to *IT* describes an action a person or organization took upon IT systems (e.g. "the insider (*NOT IT*) installed keystroke logging software on his supervisor's computer (*IT*).") With the top-level action pre-determined, analysts began by choosing the second and third level actions from a subset of choices only applicable to the top-level action type. For instance *NOT IT* to *IT* actions include "logged in," "executed," "downloaded," and "printed." *NOT IT* to *NOT IT* actions include interpersonal actions (hiring, firing, conflict between coworkers, etc.), criminal justice actions (investigations, arrests, court proceeding), or psychological or medical issues.

Our final working set of event descriptors contained 4 top-level actor/target options with 10 second- and third-level options as well as 6 top-level action options with 83 second- and third-level options. Not every category required three levels of detail to achieve the desired descriptive granularity. Table 2 illustrates each level of the triad structure as it would represent the sample chronology in Table 1. A complete list of all triad categories is provided in Appendix A.

3.3.3 Developing and Training with the Triad Method

After developing the components of the triad, we agreed on an initial operational definition for each component, which was captured in coding guidelines for the coders. Additional coding guidance was also developed to define the coding process. For instance, the case was to be coded exactly as it appeared in the chronology, including maintaining the subject/object of the sentence. We required that the chronology explicitly state the involvement of IT systems for it to be coded as an *IT* actor or target. More specific action descriptors from the text in the chronology were to be coded if at all possible (e.g. contracted rather than hired). Additional guidance included:

- *NOT IT* to *NULL* should only be used for behaviors where the insider is the actor, including side businesses and feelings. This does not include actions involving or targeting the Victim Organization.
- The default value for vague chronology items is:

  Actor *NOT IT - Unknown*

  Target *NOT IT - Unknown*

  Action *Interpersonal - Other*
- Logic bombs should be coded as NOT IT to IT (both the insertion of the malicious code and the eventual execution of the bomb).
- If a customer detects a problem with the victim organization's IT systems, then IT is the target of that event.
- An organization fearing technical sabotage should be coded as a NOT IT to NOT IT (e.g. "The Victim Organization feared sabotage during the union negotiations").

3.3.4 Inter-Rater Reliability Testing of the Triad Method

Initially, three analysts independently examined the same seven cases, applying triad categories to each event, to test inter-rater reliability (IRR). Because the Fleiss' Kappa agreement (κ) [14] was not as high as desired ($\kappa \geq 0.8$), additional training was provided, and a second IRR was undertaken with four cases and an additional analyst. After reaching the desired level of agreement among analysts, the remaining cases were examined. Additional analysts were added throughout the project, initially examining the same cases as the analysts in the second IRR test. Inter-rater agreement was then calculated across all analysts to ensure a consistently high level of agreement.

During this process, we made a few minor changes to the analysis guidelines, to add or clarify specific actions. However, previously examined cases were not modified at the time. After analysis was complete, a senior analyst reviewed the initial round of results to determine if any changes should be made to earlier cases, based on the minor changes made during the analysis process. Very few event descriptors were changed. Overall, the coding process took four analysts about one month to complete 47 cases due to meetings with senior analysts and additional training. With training completed, the same tasks could easily have been done in about one week.

We noted several lessons learned throughout the analysis processes:

- Project leaders found it helpful to meet often with the analysts to discuss problem areas in the analysis and develop solutions for future analysis.
- Analysts found it useful to develop a default method for addressing chronology items they felt were unclear. To prevent default methods from becoming a catch-all, events in these categories were carefully watched by senior analysts to see if new categories needed to be created or existing operational definitions refined.
- Analysts felt it would be useful to develop an automated processes for entering commonly observed events, both technical and nontechnical (e.g., logged on to system, insider hired).
- Emphasis on avoiding assumptions during the analysis process was crucial (e.g., if the chronology description did not explicitly state that IT systems were involved, the analyst should not assume IT systems were part of the action).

3.3.5 Adding Categories

During IRR testing, analysts were instructed to select "Other" for unknown or unclear categories and asked to provide feedback about why "Other" was chosen. In addition to refining the operational definitions of the codebook, we wanted to make sure we included any other types of actors, targets, or actions not initially identified. Based on that feedback, several additional categories were added to the coding process, including the following:

Conspirator An associate or collaborator in the incident, known to be a fellow employee (regardless of supervisor, co-worker, etc.)

Bragged Made a boastful statement

Stole To take (another person's property) without permission or legal right and without intending to return it

Discovered To become aware of

Investigated To carry out a formal inquiry

Restricted Access To limit a person's ability to enter an IT system. (note: Use "Disabled", if access was completely removed (as opposed to limited)).

3.3.6 Applying Coding Methodology to All Cases

After we tested the triad method, we applied it to 49 chronologies of insider threat sabotage from our database. As we have previously described, chronologies are a list of discrete events organized from earliest to latest occurrence. While the chronologies already existed in the database, a few changes were required in order to apply the triad more effectively. For example, certain chronology items actually contained more than one event, which would make it difficult to code. Standardization was also required to frame similar events with the same subject and object focus. Therefore, prior to coding, two trained analysts examined them existing chronologies, modifying them if necessary, to ensure that each event contained only one action. Cases were coded up to the 'zero hour' which we define as the time when cyber damage begins [3].

4 Results

Events were determined to be behavioral or technical depending on the top-level actor or target categories as defined by the descriptor for each event. Events involving IT systems were denoted as technical, and those between individuals and/or organizations were denoted as behavioral. Of the 449 events analyzed, 352 (78.3%) were categorized as behavioral and 97 (21.6%) were categorized as technical.

In cases with observable insider-initiated actions prior to attack zero-hour, including four cases with only technical observable events prior to attack zero-hour, the insider demonstrated a behavioral action prior to any technical action in 34 of 44 (78%) cases. Of the cases with both behavioral and technical observable events prior to attack zero-hour, the insider demonstrated a behavioral action prior to a technical action in 34 of 40 (85%) cases, with an average of 2.85 behavioral events prior to the first observable technical event. Also, the first technical event occurred nearly two-thirds of the way through the timeline from first observable event to the point of attack.

All cases contained at least one technical observable event prior to, or including, the point of attack. Table 2 summarizes the results.

Cases with the following order of event precedence:	Of 44 cases with observable events prior to zero hour	Of 40 cases with both behavioral and technical observable events prior to zero hour
Behavioral Prior to Technical	34	34
Technical Prior to Behavioral	10	6

Table 3: Event Precedence (Behavioral vs. Technical)

The most common event descriptors found were related to organizational actions (hiring, firing, termination), insider feelings (almost exclusively negative), insider requests of the victim organization (often unmet), and insider access of victim organization IT systems. For the purposes of this study, we distinguished "Fired," "Terminated," and "Resigned," as follows:

- Fired: The point in time when the insider is notified his or her employment is being involuntarily terminated.
- Resigned: The point in time when the insider notifies the organization of a voluntary termination of employment.
- Terminated: The point in time when the insider's employment actually ends. This could be a point in time after notice of firing or resignation is actually given.

5 Analysis

Our findings reinforce previous literature on insider sabotage suggesting that insider behavioral actions often occur before technical actions [6, 3]. This suggests that organizational groups should receive further training or resources about detecting and reporting behavioral indicators. In fact, one of the top most common descriptors is 'Insider felt', often was an indicator disgruntlement. As we have indicated previously, disgruntlement often occurs well before other technical indicators of an attack. It is also notable that so many of the top most coded behavioral actions prior to the zero hour surround termination and resignation, which can be a time of potentially higher risk.

With respect to technical actions, the second most frequent technical action event found after system access was installation of software on organization IT systems, which frequently enabled malicious actions to occur. This underscores the importance of situational awareness of an organization's IT systems, including hardware and software, for detecting unauthorized changes. Given the technology available today to monitor and manage organizational IT systems, further research could be performed to better understand this technical behavior as a potential early indicator.

In addition, we observed several benefits of using the new triad methodology including:

- Ease of separating technical and behavioral events, to more closely scrutinize for different mitigation strategies
- Ease of separating events in which the insider was the actor, to more closely scrutinize for indicators of insider threat
- Increased ability to compare events across cases, to view most common behaviors in sabotage

5.1 Initial Findings on Precursors of Insider IT Sabotage

While the purpose of the overall study was to examine patterns of insider behavior prior to attack, the scope of this paper is mainly to report findings in the development of our analysis methodology. However, we can report a couple of interesting results observed during the course of analysis:

- Insider IT sabotage is a fairly fast-acting crime, meaning there is a relatively short time span between when the insider clearly becomes disgruntled (to the point of planning an attack) and the point at which damage occurs, compared to other types of insider crimes [15]. Of the cases we studied, most had less than 1 week between the first observable action clearly indicating malicious intent and the actual moment of IT damage to the victim organization. Many of those cases actually had less than 1 day of warning prior to attack.

- Due to the nature of the cases we studied, limited detail was available on the insider's actions prior to attack. 19 of the 49 cases had between two and five observable events prior to attack. 36 of the 49 cases had between 2 and 10 observable events prior to attack.

- Of the observable events prior to attack, most involved administrative actions (i.e. hiring, firing, etc.).

- Of the non-administrative events prior to attack, the most common were the following (observed in six cases each):

  - Insider was tardy/absent
  - Insider installed software on victim org IT systems
  - Insider bullied, intimidated, or threatened others

- Of the most common victim organization initiated events, we noted that someone expressed concern about the behavior of the insider in five of the cases.

One other interesting finding that will perhaps require more analysis to flush out is that the insider made some sort of request of the organization in 9 of the 49 cases we studied. In 6 of those cases, the request was denied by the victim organization. Because we chose to focus more on semantics rather than specific details, we cannot at this time comment on the details of those requests, though they did happen to include requests for transfer, permanent employment status, etc. Further details of the analysis of the results of this project can be found in our related work [16].

6 Conclusion

In this study we focused on developing a new methodology to streamline the coding of actual insider threat cases in a chronology. After finding an older codebook to be too unrefined for the present study, we created a finite set of descriptors to apply to each event. We tested the method, and after reaching a sufficient level of inter-rater reliability, we coded 49 actual cases of insider threat sabotage. We discovered that a three-component structure, dubbed the "triad," allowed for an easy characterization of behavior or technical events. Future work may include the expansion of the triad for use with other types of insider threat cases (e.g., theft of IP or fraud).

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] C. Huth, D. Mundie, and S. Perl, "Toward an Ontology for Insider Threat Research: Varieties of Insider Threat Definition," in *Proc. of the 3rd Workshop on Socio-Technical Aspects in Security and Trust (STAST'13), Tulane University, New Orleans, LA, USA*, June 2013.

[2] R. Yin, *Case Study Research: Design and Methods*, 4th ed. Sage, 2009.

[3] W. R. Claycomb, C. L. Huth, L. Flynn, D. M. McIntire, and T. B. Lewellen, "Chronological Examination of Insider Threat Sabotage: Preliminary Observations," *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications (JoWUA)*, vol. 3, no. 4, pp. 4–20, December 2012.

[4] S. R. Band, D. M. Cappelli, L. F. Fischer, A. P. Moore, E. D. Shaw, and R. F. Trzeciak, "Comparing insider IT sabotage and espionage: A model-based analysis," Carnegie Mellon University/Software Engineering Institute, Tech. Rep. CMU/SEI-2006-TR-026, December 2006.

[5] M. Keeney, E. Kowalski, D. Cappelli, A. Moore, T. Shimeall, and S. Rogers, "Insider Threat Study: Computer System Sabotage in Critical Infrastructure Sectors," Carnegie Mellon University/Software Engineering Institute, Tech. Rep., May 2005.

[6] A. Moore, D. Cappelli, and R. Trzeciak, "The "big picture" of insider IT sabotage across U.S. critical infrastructures," http://www.cert.org/archive/pdf/08tr009.pdf, Carnegie Mellon University/Software Engineering Institute, Tech. Rep. CMU/SEI-2008-TR-009, May 2008.

[7] 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 *Proc. of the 2005 International Conference on Intelligence Analysis*, May 2005.

[8] J. Hunker and C. Probst, "Insiders and insider threats - an overview of definitions and mitigation techniques," *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications (JoWUA)*, vol. 2, no. 1, pp. 4–27, March 2011.

[9] D. M. Cappelli, A. G. Desai, A. M. Moore, T. J. Shimeall, E. A. Weaver, and B. J. Willke, "Management and education of the risk of insider threat (MERIT): System dynamics modeling of computer system sabotage," http://www.cert.org/archive/pdf/merit.pdf, Carnegie Mellon University/Software Engineering Institute, Tech. Rep., 2006.

[10] R. Trzeciak, "The CERT Insider Threat Database," http://www.cert.org/blogs/insider_threat/2011/08/the_cert_insider_threat_database.html, August 2011.

[11] E. Shaw and L. Fischer, "Ten Tales of Betrayal: The Threat to Corporate Infrastructure by Information Technology Insiders," Defense Personnel Security Research Center (PERSEREC), Tech. Rep. 05–13, 2005.

[12] J. Sweller, J. J. G. van Merrienboer, and F. G. W. C. Paas, "Cognitive Architecture and Instructional Design," *Educational Psychology Review*, vol. 10, no. 3, pp. 251–296, 1998.

[13] M. East and M. Herman, "Targets in foreign policy behavior," in *Describing Foreign Policy Behavior*, P. T. Callahan, Ed. SAGE Publications, Inc., 1982, pp. 115–132.

[14] J. L. Fleiss, "Measuring Nominal Scale Agreement Among Many Raters," *Psychological Bulletin*, vol. 76, no. 5, pp. 378–382, November 1971.

[15] A. P. Moore, D. M. Cappelli, T. C. Caron, E. Shaw, D. Spooner, and R. F. Trzeciak, "A Preliminary Model of Insider Theft of Intellectual Property," Carnegie Mellon University/Software Engineering Institute, Tech. Rep. CMU/SEI-2011-TN-013, June 2011.

[16] W. R. Claycomb, C. Huth, L. Flynn, and D. McIntire, "Early Warning Indicators of Malicious Insider IT Sabotage," in *Proc. of the 47th Annual IEEE International Carnahan Conference on Security Technology (ICCST'13), Medellin, Colombia*. IEEE, October 2013.

Author Biography

William R. Claycomb is the Lead Research Scientist for the CERT Enterprise Threat and Vulnerability Management program at Carnegie Mellon University's Software Engineering Institute. His primary research topic is the insider threat; current work includes discovery of insider threat behavioral patterns and corresponding sociotechnical countermeasures. Dr. Claycomb is also involved in other efforts at CERT exploring cloud computing, incident response, systems modeling, and vulnerability analysis. Prior to joining CMU, he was a Member of Technical Staff at Sandia National Laboratories, focusing on enterprise systems security research, including insider threats, malware detection, and data protection. Bill is currently an adjunct faculty member at CMU's Heinz College, teaching in the School of Information Systems and Management.

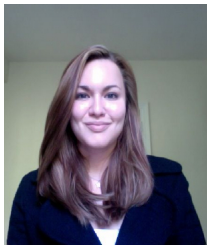

Carly L. Huth is an insider threat researcher with the CERT program at the Software Engineering Institute. Huth's current areas of research include the intersections of privacy and technology as well as the effects of the current regulatory regime on insider threat prevention practices. She holds a Juris Doctor from the University of Pittsburgh School of Law.

A Appendix - Triad Coding Guidelines

A.1 Overall coding guidelines

1. Please code the case exactly as it appears in the event chronology (e.g. keep the actor/target focus the same).

2. Do not categorize an actor or target as IT unless the chronology explicitly includes IT.

3. Enter in at least Level 1 of both Actor and Target prior to entering Action Level 2.

4. Please be as specific as possible given the known information (e.g. choose 'deleted' rather than 'attacked').

5. Please note definitions/events that are troublesome in the comments section.

6. Logic bombs should be coded as NOT IT to IT (both the insertion of the malicious code and the eventual execution of the bomb).

7. NOT IT to NULL should only be used for behaviors where the Insider is the actor, including side businesses and feelings. This does not include actions involving or targeting the Victim Organization.

8. NOT IT to NOT IT, Interpersonal, Other is the default for vague chronology items.

9. lease choose the word closest to the one used in the chronology if possible (e.g. keep 'contracted' instead of changing to 'hired').

10. Stealing passwords the target is Victim Org property not the owner of the password.

11. If a customer detects a problem with the Victim Org's IT, then IT is the target.

12. Organization fearing technical sabotage should be coded as a NOT IT to NOT IT (e.g. The Victim Organization feared sabotage during the union negotiations).

13. Victim Org owned systems should be coded as Victim Org IT even if the Insider uses this system.

A.2 Operational Definitions

The following contains definitions and examples for the three pieces of the triad: *Actor, Target,* and *Action*. There are three level: Level 1, Level 2, and Level 3.

Actor The subject of the event; the person, organization, or system that takes action on the target.

Target The object of the event; person, organization, or system that is being acted upon.

Action The subject's behavior that affects the target.

A.3 Actor/Target Categories

IT any information system that initiates or is the target of the action. If IT is the actor, than this is automated process (e.g. an intrusion detection system notifying an admin.)

Victim Org IT The Victim Organization is the organization that is harmed by the Insider's action. This category is only used for the IT system itself, not for IT staff.

Other Org IT any organization's IT that is not the Victim Organization (e.g. new employer that may be the source of the attack, but was not harmed).

Internet public worldwide computer network system (e.g. actions would include posting).

Insider's IT the perpetrator of the incident's own computer, device, etc.

Unknown this category should only be used when the actor or target is known to be an information system but the owner of the system is unknown.

Other any IT system not described above (please describe this actor/target on spreadsheet).

Not IT any other known actor or target should be coded as Not IT, includes people (including IT staff), organizations, court, life events, etc.

Insider the perpetrator of the incident

Insider his or herself same as above.

Insider's property any property (such as a home) that is not an IT system.

Victim Organization The Victim Organization is the organization that is harmed or intended to be harmed by the Insider's action.

Supervisor the Insider's immediate boss only. This may only be used for the Insider's direct, immediate supervisor, if other supervisors are present in the chronology please code them as management.

Management any other people in the organization that are higher on an organizational chart than the insider. This does not include the Insider's immediate boss, who should be coded as supervisor.

Co-worker any people in the organization that are at the same level as the insider in the organization.

Subordinate any person who directly or indirectly reports to the insider.

Organization Group (e.g. HR/IT) this should be used when someone acts on behalf of the organization (e.g. hire, fire, a member of IT staff sends out a policy violation notification).

Property any non IT property that may be involved (usually as a target)

Conspirator associate or collaborator in the incident, known to be a fellow employee (regardless of supervisor, co-worker, etc.)

Unknown this should be used when it is known that the victim organization acted, but it is unknown who or what within the organization acted. This category should be used when the relationship between the actor and the insider is not known (e.g. conspirator who also worked for the victim organization).

Criminal Justice an actor or target involving practices related to law enforcement or criminal or civil court proceedings.

Court an assembly of persons for the administration of justice.

Authorities anyone involved with law enforcement.

Outside any other actor or target that is not a member of the Victim Organization, Criminal Justice, or the Insider themselves.

People/Property any other person or piece of property that is part of the incident, either as an actor or target.

Organizations any other organization that is part of the incident, either as an actor or target.

Conspirator associate or collaborator in the incident, known to be outside the organization

Other any other actor or target that is not listed.

Unknown this category of actor/target should be used only when the chronology does not make clear if an information system was the actor/target or not.

Insider the perpetrator of the incident.

Victim Org The Victim Organization is the organization that is harmed by the Insider's action.

Outside any other actor or target that is not a member of the Victim Organization, Criminal Justice, or the Insider themselves.

Conspirator associate or collaborator in the incident, unknown whether or a fellow employee or outsider

Other any other actor or target that is not listed.

Unknown nothing is known about the actor/target.

Null this category should only be used to code life events (see Actions section for further explanation). For these events, please code NOT IT, INSIDER as the actor and NULL as the target.

A.4 Action Categories

NOT IT TO NOT IT Interpersonal actions which take place between at least two individuals (or individual and other property), including actions on behalf of the organization.

Made Agreement at least two parties entered into an arrangement accepted by all parties. e.g. Insider agreed with Supervisor to document all code.

Broke Agreement at least one party did not do what they had agreed to do e.g. Insider did not document all code, despite making an agreement with his supervisor that he would do so.

Bullied/Intimidated/Threatened to frighten, argue, or state intention to take hostile action e.g. Insider told his supervisor that he was going to take down the system if he did not get paid.

Argued exchange or express diverging or opposite views, typically in a heated or angry way e.g. Insider argued with co-worker about a project.

Bragged a boastful statement e.g. Insider bragged to co-workers about his hacking skills.

Stole take (another person's property) without permission or legal right and without intending to return it e.g. Insider stole office supplies.

Denied refuse to give or grant, including denial of physical access. e.g. Supervisor denied the Insider's request to work from home.

Approved officially agree or accept. e.g. Supervisor approved the Insider's request to obtain additional resources for his project.

Complained express dissatisfaction or annoyance with. e.g. Insider complained to co-worker about supervisor.

Requested asking for something. e.g. Insider requested additional resources from his Supervisor.

Blamed to assign responsibility for a fault or wrong. e.g. Supervisor blamed Insider for the slow movement of the project.

Mitigated make less severe or serious (includes recovering property). e.g. Supervisor mitigated the situation by offering the Insider time off; the IT staff mitigated the Insider's attack by taking the server offline.

Physically Contacted includes all forms of related to touching the body, positive or negative. e.g. Insider hugged co-worker; Insider hit co-worker.

Reported give a spoken or written account of what one observed e.g. Co-worker reported physical contact to Supervisor.

Flirted behave as though attracted or trying to attract someone. e.g. Insider propositioned co-worker.

Regretted to feel repentant or disappointed over. e.g. Insider told Supervisor he regretted attempting to delete files.

Expressed Concern to indicate worry or anxiety. e.g. Supervisor expressed concern over the Insider's sloppy working habits.

Lied/Concealed to make a deliberate false statement; to intentionally keep secret or hidden . e.g. Insider spoke with his Supervisor, telling him that he was documenting code, even though he was not.

Communicated to share or exchange information with; this should only be used when one of the more specific categories does not apply e.g. Supervisor called the Insider.

Hired to employ someone. e.g. the Victim Organization hired the Insider as a sys admin (the Victim Organization is always the actor in this case).

Contracted to enter into a formal agreement. e.g. the Victim Organization contracted the Insider to a temporary IT staff member.

Terminated this should be used when the Insider actually leaves the organization.

Resigned to give formal notice of leaving employment. e.g. the Insider gives notice to the Victim Organization that he or she will be leaving (when the Insider actually leaves please mark this as termination).

Fired to dismiss from the job; use this only when the Victim Organization gives notice to the Insider that he or she will no longer be working there due to some fault of the employees (when the Insider actually leaves please mark this as termination).

Laid Off use this only when the Victim Organization gives notice to the Insider that he or she will no longer be working for the organization, not because of the personal performance but due to changes in the organization (when the Insider actually leaves please mark this as termination).

Transferred to move from one location to another. e.g. the Victim Organization transferred the Insider to a different department.

Sanctioned to give a penalty for disobeying a rule. e.g. the Victim Organization sanctioned the Insider for coming into work late by cutting the Insider's pay.

Rewarded to give recognition for good behavior. e.g. the Victim Organization gave the Insider a raise because of his good performance.

Promoted to advance position in rank. e.g. After his supervisor left, the Insider was promoted to his old supervisor's position.

Demoted to give someone a lower rank. e.g. Because of the Insider's poor performance, the Insider was demoted to a lower level employee.

Discovered to become aware of. e.g. the Victim Organization discovered that the Insider had a criminal history.

Investigated to carry out a formal inquiry. e.g. the Victim Organization interviewed coworkers about the missing money.

Reorganized to change structure. e.g. the Victim Organization reorganize changing the Insider's group (e.g. adding or modifying departments). (Please code triad as Victim Organization, Organization Group as both the actor and the target)

Was Tardy/Absent to arrive to work late or not at all. e.g. Insider was late to work.

Other to be used only when none of the actions above fit.

Criminal Justice Actions Actions involving practices related to law enforcement or criminal (or in a few cases civil) court proceedings.

Arrested the act of taking someone into custody. (Actor must be LE) e.g. Law enforcement came to the office and arrested the Insider.

Pleaded to present and argue a position in court. (Target must be Court) e.g. Insider plead guilty, not guilty or no contest before the court.

Sentenced when the court assigns punishment to a guilty defendant. (Actor must be Court) e.g. the Court sentenced the Insider to four years in prison and 2,000 in restitution.

Reached Verdict (Actor must be Court) when the court decides whether the defendant is guilty or not guilty. e.g. the Court found the Insider not guilty.

Searched to look carefully inspect an area or person. (Actor must be LE) e.g. Law enforcement searched Insider's home.

Identified to establish or indicate who or what something is. (Actor must be LE) e.g. Law enforcement identified the Insider as the perpetrator through.

Sued to institute legal proceeding against. e.g. the Insider sued the Victim Organization for unpaid wages.

Seized to take through formal legal means. (Actor must be LE) e.g. Law enforcement seized the Insider's home laptop.

Investigated to carry out a formal inquiry. (Actor must be LE) e.g. the Victim Organization called in law enforcement to investigate the sabotaged system.

Analyzed to examine a piece of property methodically. (Actor must be LE) e.g. Law enforcement analyzed logs.

Evaded escape or avoid. e.g. Insider evaded arrest by hiding with relatives.

Indicted/Charged to bring formal charges against. e.g. The Court indicted the Insider.

Other to be used when none of the other actions fit.

Psych/Medical actions which relate to the physical, mental, or emotional health of the insider. (If the report states that a doctor diagnosed, please choose OUTSIDE, PERSON otherwise please choose OUTSIDE, ORGANIZATION.)

Diagnosed to formally identify an illness. e.g. a doctor diagnosed the Insider with depression.

Treated to provide medical care for. e.g. the Insider was treated with medication for depression.

Other to be used when none of the other actions fit.

NOT IT to IT **Logged in** to go through the procedures to begin use of a computer, database, or system. e.g. Insider logged into his co-workers work station using his own credentials.

Logged off to go through the procedures of exiting a computer, database, or system. e.g. Insider logged off of his work account. (If the Insider keeps the connection open do not code it as a logged off)

Accessed to retrieve or examine information. e.g. Insider accessed proprietary information stored on the server.

Executed to carry out an instruction or program. e.g. Insider executed a code sequence to delete data.

Moved to change the location of. e.g. Insider moved all files onto one server.

Copied to make another version of. e.g. Insider copied a file onto his laptop.

Exploited to take advantage of a vulnerability. e.g. Insider exploited a known vulnerability in order to enter the system after his access had been cut off.

Analyzed to scrutinize. e.g. Supervisor analyzed logs from Insider's computers.

Posted to publish in a public place. e.g. Insider posted information on how to cause a denial of service attack to the Victim Organization.

Downloaded (came from Internet) to copy data from one computer system to another or to a disk e.g. Insider downloaded a hacker tool.

Printed to produce a paper copy. e.g. Insider printed proprietary information.

Read to view on an IT system. e.g. Insider read emails on his boss's computer.

Delayed to make some part of the IT system late or slow. e.g. Insider delayed network access to co-workers. (Use only if known that that the attack was remediated)

Deleted to remove or obliterate part of an IT system or data stored on the system. e.g. Insider deleted code that was needed to allow the system to run properly.

Modified to change part of an IT system or data stored on the system. e.g. Insider modified code that was needed to allow the system to run properly.

Created to make or add something new to part of an IT system or data stored on the system. e.g. Insider created a backdoor to enter the Victim Organization's system.

Disabled to greatly impair the function of some part of an IT system or data stored on the system. e.g. Insider disabled access to Victim Organization's system; Victim Organization disabled Insider's access after his termination.

Recovered to restore or repair part of an IT system or data stored on the system. e.g. The Victim Organization recovered the files that the Insider attempted to delete.

Installed to load software on a computer. e.g. Insider installed a password cracker on the Victim Organization's system.

Physical Contact to touch or affect and IT system through physical means. e.g. Insider shot server.

Attacked use this only if you do not know any specific methods for how the Insider attacked.

Discovered to become aware of e.g. The Victim Organization discovered that the malicious software.

Investigated to carry out a formal inquiry e.g. The Victim Organization investigated the source of the malicious software.

Restricted Access to limit a person's ability to enter an IT system. Use "Disabled" instead, if access was completely removed (as opposed to limited). e.g. The Victim Organization restricted access to the IT systems during the union negotiations.

Other to be used when none of the other actions fit.

IT to IT Executed Schedule Programs automated IT processes.

Scripts to carry out automated series of instructions carried out in a specific order.

Backup Jobs the procedure for making extra copies of data in case the original is lost or damaged.

File Transfers to transmit files over a computer network.

Other to be used when none of the other actions fit.

IT to NOT IT Alerted to warn of a possible violation. e.g. the Victim Organization's IT system alerted the Victim Organization IT staff that there was a policy violation.

Other to be used when alerted does not fit.

NOT IT to NULL (life events) Changed locations to move from one place to another, temporarily or permanently.

Traveled to make a trip or journey. e.g. the Insider traveled frequently for work.

Moved please use this for any change of physical location not related to a transfer. e.g. Insider moved to a different neighborhood.

Felt to feel a specific emotion. e.g. Insider felt upset by his supervisor's criticism of his work.

Financial any action the Insider takes with respect to money.

Changed financial status a sudden increase or decrease in income. e.g. Insider had a financial problems.

Sold to hand over property in exchange for money. e.g. Insider sold office supplies he had stolen.

Purchased to acquire something by paying for it. e.g. Co-workers noticed that the Insider had made some expensive recent purchases.

Changed Health Status (Insider only) Insider became ill or well (but a doctor was not noted to be involved). e.g. Insider had ongoing medical problems.

Personal actions that the Insider takes on his or herself.

Drank to consume a beverage (typically alcohol). e.g. After being fired, the Insider drank heavily.

Did drugs to take an illegal substance. e.g. After being fired, the Insider began taking drugs.

Gambled to play games for money. e.g. Co-workers noticed Insider's frequent gambling trips.

Other please use this category for other actions that the Insider takes to his or her own person.

Family any actions that take place surrounding the Insider's relatives.

Had family problem issues surrounding the Insider's relatives. e.g. Insider's father was in poor health.

Changed family status divorce, marriage, adoption, etc. e.g. Insider was recently divorced.

Other to be used when none of the other actions fit.