

Deriving Candidate Technical Controls and Indicators of Insider Attack from Socio-Technical Models and Data

Michael Hanley

January 2011

TECHNICAL NOTE
CMU/SEI-2011-TN-003

CERT® Program
Unlimited distribution subject to the copyright.

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



This report was prepared for the

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

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

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

Copyright 2011 Carnegie Mellon University.

NO WARRANTY

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

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

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

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

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

For information about SEI publications, please visit the library on the SEI website (www.sei.cmu.edu/library).

Table of Contents

Acknowledgments	vii
Abstract	ix
1 Introduction	1
2 Definitions	2
3 Discussion and Related Work	3
4 Case Categories	5
5 Method	6
6 Case Examples	9
7 Linking Vulnerability Metrics	13
8 Limitations	14
9 Conclusion	16
References	17

List of Figures

Figure 1:	The “Entitled Independent Model” of Insider Theft of IP	4
Figure 2:	Sample Lab Topology for Creating Demonstrations and Testing Candidate Technical Controls and Indicators	11
Figure 3:	Screenshot from Demo Video Showing an Alert from a Splunk Rule Derived from Models of Insider Theft of IP	12

List of Tables

Table 1: An Example Insider Threat Case Decomposed

6

Acknowledgments

This paper originally appeared in the proceedings of the NSA Center of Academic Excellence Workshop on Insider Threat, held in St. Louis, Missouri, in November 2010. The author would like to thank the program chairs, Dr. Aaron Ferguson and COL Ron Dodge, as well as the rest of the program committee for their helpful reviews and feedback on this paper.

Also, the author thanks fellow members of the Insider Threat Center at CERT[®] who provided valuable feedback on this paper, including Dawn M. Cappelli, Andrew P. Moore, and Randall F. Trzeciak.

Finally, a special thank you to two of our CERT technical editors, Ed Desautels and Melanie Thompson, for their assistance with preparing this paper for distribution.

Abstract

The insider threat continues to be one of the prime issues facing government entities and organizations across critical infrastructure sectors. Extensive catalogues of case material from actual insider events have been used by CERT[®], part of Carnegie Mellon University's Software Engineering Institute, to create socio-technical models of insider crime to help educate organizations on the risk of insider crime. Building upon this work, this paper seeks to demonstrate how a useful method for extracting technical information from previous insider crimes and mapping it to previous modeling work can create informed candidate technical controls and indicators. This paper also shows current examples of case material and candidate indicators that have been successfully converted into well-received insider threat training modules.

1 Introduction

Since the original joint study on insider threat conducted by the U.S. Secret Service and the CERT[®] Program at Carnegie Mellon University's Software Engineering Institute in 2001, CERT has catalogued over 400 cases of actual insider crimes [CERT 2008]. The CERT Program's nearly 10 years of studying insider threats has produced several interesting reports ranging from targeted examinations of individual critical infrastructure sectors to system dynamics models of the behavioral and technical aspects of insider crimes [MERIT 2008]. Other published works have examined either the technical or the behavioral aspects of insider threat research. However, work at the intersection of both the technical and the behavioral aspects of insider crime is sparse. The CERT Program's vision for the ideal insider threat detection tool is one based on a predictive model that includes both technical and non-technical indicator identification implemented as a series of detection algorithms. Each algorithm will consist of chronological, weighted technical and non-technical indicators.

This paper describes how CERT is building upon our previous work in the modeling arena and leveraging our understanding of insider behavior to begin to work toward that vision. We are currently undertaking a combination of metrics research to identify behavioral and technical indicators, relative weights, and applied research to create technical controls that map to case information as potentially effective countermeasures to insider threat. We are using a novel approach based on data from the existing CERT Insider Threat Database. This approach comprises a method for extracting technical details and mapping them to existing tools. In this paper, we describe an example of how early work in this area has turned into well-received instructional materials that help educate interested parties on new countermeasures against the threats posed by insiders.

[®] CERT is registered in the U.S. Patent and Trademark Office by Carnegie Mellon University.

2 Definitions

CERT defines insider threat as the following:

A malicious insider is 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.

The definition is clear and creates a scope that is easy to manage from an analytic perspective. Throughout the remainder of this paper, this definition accommodates all references to insiders. The most recent modification to this definition, the addition of language referencing business partners, is addressed in a related work that shows the danger of insider threats from non-employees who have a unique business relationship with an organization [Weiland 2010].

3 Discussion and Related Work

Any discussion of insider crime, particularly of the technical details associated with these crimes, should be tempered by considerations of the types of data available to researchers. In general, studying insider crime begins with the collection of open-source case information. At CERT, a significant portion of our case source material for non-national-security espionage cases is gathered from freely accessible internet news sites, blogs, and other such postings. Where possible and practical, we also collect court documents, including affidavits, transcripts, sentencing information, and other relevant items that contain information about the crime. From both media sources and court documents, we learn a great deal about the insider and what the insider stole or damaged at the victim organization. We also gain a general understanding of the sequence of events associated with the case. However, in many cases, the technical details of how the crime was committed are not available.

This lack of empirical data about what technical vulnerabilities were exploited makes it difficult to develop technical controls against malicious activity. This paper seeks to provide some guidance on extracting relevant technical information from previously catalogued cases of insider threat to develop technical controls at the intersection of the tool space and previous work on behavioral models. However, we believe that, with a reasonable degree of confidence, in even vague technical data, a skilled analyst can determine what technical countermeasures would be effective against insider crimes, even if all an analyst knows is what was included in a media report. Frequently, we find that knowing the exfiltration method and the source and destination of the stolen asset is sufficient to do a great deal of useful analysis. This is what we intend to demonstrate in the following few pages.

Consider also that a large number of cases detected by victim organizations are not reported to law enforcement. Last year, a survey conducted by *CSO* magazine in cooperation with the U.S. Secret Service, CERT, and Deloitte revealed that 72 percent of respondent organizations that experienced at least one malicious insider incident in the previous year handled the incident internally without involving law enforcement [CSO 2010]. More interesting, however, are the reasons for not reporting the incident, which range from lack of evidence (35 percent) to an inability to attribute the event to any individual malicious actor (29 percent) to the fear of negative publicity (15 percent). These are concerning figures, and while somewhat problematic for the purposes of research, they also motivate work in this area to help improve the ability of organizations to prevent, detect, and respond to malicious insiders.

Another important consideration here is that while the method, technical controls, and indicators discussed in this paper are all rooted in extensive case files and previous modeling work, they have not been operationally tested. We are interested in testing some of our proposed indicators for effectiveness on operational networks; however, at this time, the candidate indicators and controls remain untested, but well-informed, measures for insider threat defense.

It should also be noted that this work relies heavily on previous work developing socio-technical models of insider crime using system dynamics methods. CERT has produced several of these models of insider crime based on a focused analysis of cases from each of the types of crimes stu-

died at CERT, combined with extensive research and feedback from behavioral psychologists, security researchers, and management experts. These works have been well received and serve as an interesting mechanism for breaking insider crimes down into manageable feedback loops and common trends. The statistics and conclusions derived from these works form the foundation for our initial work in control development; however, this is not intended to be the exclusive input for control development. A sample output from the recent work on intellectual property (IP) theft modeling is shown below in Figure 1 [Moore 2009].




Figure 1: The "Entitled Independent Model" of Insider Theft of IP

It is also worth noting that there are other pieces of significant insider threat research that can be used as platforms for deriving candidate controls and indicators. Models and studies that are of interest include behavioral models from Shaw, Ruby, and Post [Shaw 1998] and extensive work done by Herbig and Wiskoff at the Department of Defense's (DoD) Defense Personnel Security Research Center [Herbig 2002], particularly for the study of national security espionage. There are still other works in related areas, such as misuse detection algorithms, that can also serve as an informative source when creating candidate controls [Cathey 2003]. Also, there are many robust commercial suites, particularly in the data loss prevention (DLP) space that can perform a wide range of data collection that make for an excellent set upon which to apply candidate controls and indicators to test for effectiveness. We feel this work complements other, more focused work in the tool space by directing research toward high-impact insiders' exploits observed in our database.

4 Case Categories

At the time of this writing, the CERT Insider Threat Database, henceforth referred to as “the database,” includes over 400 cases of insider crime that have been tried in the court system and that produced a conviction or guilty plea. The three core types of cases include IT sabotage, insider fraud, and insider theft of IP. A fourth, but smaller, category exists for miscellaneous insider crimes that do not fit in one of the three core categories. CERT also studies national security espionage involving classified information stolen by malicious insiders; however, that content is out of scope for the work described in this paper.

This paper focuses on cases of IP theft. CERT defines this type of crime as an insider’s use of information technology (IT) to steal IP from an organization. We also include critical and sensitive organizational information in our use of IP. This category includes industrial espionage involving insiders [Cappelli 2009]. A forthcoming paper describes deeper analysis of this case set; however, this work describes the underlying methodology and some preliminary findings.

5 Method

For a first pass at indicator development, we considered extracting relevant technical details about how the insider event was perpetrated, what assets were targeted, the mode of exfiltration, and any items of interest that could help to provide context for our analysis. This created a useful framework for quick comparison across cases. An example is included in the table below.

Table 1: An Example Insider Threat Case Decomposed

Attribute	Sample Case
Incident identifier	Sample ID
Incident summary	Insider engaged in discussions with the chair of a foreign competing firm regarding employment opportunities for the insider at the competing firm. Insider agreed to steal information from the victim firm and bring it to the competing firm in exchange for a job.
Asset attacked/target	Business plans, trade secrets, engineering and design specs
Source	Electronic documents, email attachments
Method of exfiltration	Remote network access, electronic theft of documents
Exfiltration comments	Insider compressed files and sent via corporate mail to competing firm.
Candidate Controls	
Prevention	Clarify ownership of IP, disallow remote access, disallow sending of sensitive materials and messages for competing domains, restrict access to sensitive information, improve filtering, employ digital rights management.
Detection	Monitor behavior between resignation and termination, monitor remote access, monitor access after normal working hours, monitor user network activity and downloads, monitor traffic to/from competitor domains.
Incident response	Disallow remote access, audit access controls, audit user activity, audit remote access logs, audit email logs, audit traffic to/from competitor domains.

For each attribute, we allow several standard responses. We developed the standard responses based on prior experience with the data set, though responses could be modified in the future to accommodate changes in tactics by insiders or by types of assets targeted in other insider crimes, particularly IT sabotage, since the goals of the malicious insider clearly differ in sabotage as opposed to theft of IP.

The only free-text field in the database is the Incident Summary for each incident. This allows an analyst working with this data set to quickly obtain context for the remaining material without having to work through the raw data used to code an insider threat case in our database. This is not intended to provide an analyst with the background required to become an expert on the case. Rather, it is intended to serve as a quick refresher on critical details when sifting through multiple cases from a sampled data set. Since source material could easily number in the tens, if not hundreds, of pages, a good summary is critical.

Next, when addressing the type of asset attacked, we allow different types of assets at roughly the same level of abstraction, such as trade secrets, customer information, source code, business plans, internal business information, and proprietary software. The attribute allows selection of one or more asset types, with the understanding that there may be overlap between two or more asset types given the information stolen or the critical infrastructure sector in which the case originated. As will be shown in a forthcoming work by Hanley and associates, we find that trade secrets are the most frequently attacked asset type.¹ This finding is based on a pool of 50 IP theft cases documented in the database. In this pool, trade secret theft represented over half of the cases. In roughly one-fourth of those cases, more than one type of asset was targeted by the insider.

We also categorize the source material associated with data stolen in various states, such as electronic documents, databases, printed documents, and so on. Since some data-loss-prevention tool suites have capabilities associated with detection and prevention of sensitive print jobs, access to certain document types, or the disallowance of certain database queries, we felt this was a useful distinction to include in the database for future study. Also, it is helpful to understand the stolen asset's format, particularly when discussing whether or not the data was physically exfiltrated from an organization (digital media or printed documents) or exfiltrated over the network (email with a sensitive attachment) in the following fields.

Lastly, we look at the methods used to exfiltrate the data. We do recognize that multiple network protocols and physical assets can be used to exfiltrate information. However, we try to identify the primary technical methods or protocols used to facilitate the crime, such as a large-volume download over VPN, email to a direct competitor, and so on. This is one of the more important technical items we catalogue because physical exfiltration and networked exfiltration appear to have very different implications when considering tools and countermeasures. In the previously mentioned publication, from our 50-case sample, we found 32 cases involved exfiltration of at least one stolen asset via the network, with email and remote file transfers over VPN being the most frequent protocols used to move the stolen data. More concerning, roughly one-third of these cases involved remote network access after normal business hours. Where insiders used physical exfiltra-

¹ Hanley, M. P., et al. "An Analysis of Technical Observations in Insider Theft of Intellectual Property Cases." Software Engineering Institute, Carnegie Mellon University, Forthcoming.

tion for stealing the information, their two most frequently used tools were a work-issued laptop and removable media, such as a USB drive or writable CD.

We also use previously developed CERT best practices [Cappelli 2009] to describe new control strategies in three areas: prevention of the crime, detection of the crime, and mitigation or response measures to the crime. There are tool strategies that cover all three areas; however, we delineate them specifically with the intention of providing a multifaceted technical approach to this problem. For example, an organization with a well-instrumented network that is looking to bolster its incident response capabilities can derive benefit from our suggested insider-threat-specific strategies for mitigation and response. This allows for a piecemeal approach to insider threat defense by picking and choosing the suggested controls that represent low-hanging fruit to an organization's IT department. Once these strategies have been enumerated in abstract terms, we focus on translating these controls to a set of real-world tools and policy measures packaged for quick deployment in an organization as an operational capability.

Lastly, we are looking to identify cases of insider incidents that could have been prevented through the use of specific best-in-class commercial or open-source tool suites designed to prevent data loss or specifically marketed as insider threat defensive tools. Note that our suggested mitigation strategies often consist of technical measures combined with policies and processes based on patterns in the crimes identified in our previous modeling work. In a forthcoming work, we will examine this more closely by reviewing several best-in-class tools and their capabilities and how those tools could have been successful at preventing, detecting, and responding to the actions committed by insiders in various cases. For the purposes of this work, we show how providing standardized generic controls provide the basis for mapping these ideas back to specific operational tools.

6 Case Examples

Recently, the Insider Threat Center at CERT selected a set of insider theft of IP cases for deeper study. An example of how this framework is applied to an actual, but anonymized, case demonstrates the usefulness of this method for identifying new countermeasures. The example employs tools already being used by many organizations and captures them as instructional assets to convey to analysts, managers, and other persons with an interest in insider threat defense.

The case example used included several interesting factors and aligned closely with the recent IP theft system dynamics models produced for this type of crime. We know that insiders who steal IP are typically scientists, engineers, or programmers. They steal assets they created and to which they have authorized access. The insiders usually steal the information within 30 days of announcing their resignation [Moore 2009]. Common exfiltration methods include sending email to competitors or foreign organizations, using personal email accounts from work, and downloading files to removable media or to laptops. In the case selected, the insider, employed by a firm manufacturing primarily electronic devices and microprocessors, used inside knowledge and privileged access to steal proprietary product information and send it to a competing firm in a foreign country. After communicating back and forth with a high-level official at the foreign competitor, the insider submitted his resignation to his employer with no mention of the foreign competitor. Following his notice of resignation and prior to his last day of work, the insider proceeded to email several compressed sets of confidential files off the network directly to a contact at the competing firm. The case detail also suggests the insider had emailed sensitive information from the corporate network before, specifically, to a personal email address.

Using the aforementioned framework method, we begin by breaking out key components of the case into technical areas of interest. First, we consider the above introduction to the case to be a summary of sufficient length and detail to provide the analyst a clear picture of what happened. Second, we identify the target asset: stolen trade secrets. Next, we consider the source of the asset, which appears to have been a repository of sensitive documents, likely a file server. Lastly, the medium used to exfiltrate the data was the corporate network, specifically the standard corporate email environment, from which the insider sent an email directed to an individual at a foreign competitor firm.

Next, when considering control strategies, we examine what may have prevented the crime, led to its detection via monitoring, or allowed for more efficient and effective incident response after the crime occurred. Of the three outcomes we consider, prevention is preferable. However, this is not always possible, especially in organizations that move millions of email messages across their network every day.

Moving to the next control type, detection, we find the biggest opportunity for improvement. An important link between the behavior of the insider and the technical countermeasures exists when we consider whether or not the insider had been subjected to additional monitoring as a result of his pending resignation. In this case, we know that the insider submitted his resignation in advance of ending his employment with the victim firm and that the data was stolen during this period. Also, we know, through our work in system dynamics modeling, that 65 percent of insiders

steal information within a month of resignation [Moore 2009]. This compelling statistic on insider behavior creates an interesting opportunity from which we derive the following sample of a candidate technical control in our instructional demonstration. Also, while the 65-percent measure is by no means sufficient to be considered an assured way of detecting insider data exfiltration, it remains a substantial finding that allows for the creation of informed rules from empirical findings. In concert with other strategies, for example, targeting this rule toward insiders who have exhibited other behaviors that make them more likely to commit a crime against the organization should be far more effective than the 65 percent statistic suggests.

To create the example control for a demonstration video, we considered the primary exfiltration method (email) and ignored other media the insider may have used to conduct the crime. In a workshop setting, where these demonstration videos are most frequently used, the intent is to drive home the “big picture” view of insider crime rather than focusing overly on any one behavioral or technical detail. We also considered which tools would potentially provide visibility into the insider’s movement of the data off the network via email. While there are several security appliances and points between the client and the gateway where the message traffic could be inspected, we were interested in an approach that focused on our suggested best practices related to auditing and monitoring [Cathey 2003]. Specifically, we are interested in how we can utilize centralized logging, or a centralized log querying mechanism, to tie in the known technical indicators of insider crime with known behavioral aspects. The technical indicator in this case is the email to a direct competitor’s domain containing an attachment. The key behavioral aspect of this type of crime is the finding that 65 percent of insiders steal within the one-month window surrounding resignation.

To demonstrate this, we created a small virtual environment, designed as a microcosm of an enterprise network. The environment used for this demonstration consisted of tools and appliances ranging from net flow collectors to a Microsoft Exchange infrastructure for handling corporate email. With these services configured, we added a Splunk appliance, chosen primarily for ease of configuration and our prior experience with the tool on other operational networks [Splunk 2010]. Splunk served as a central query system for accessing both the Exchange logs and the domain controller event logs. The complete lab network topology is shown in Figure 2.

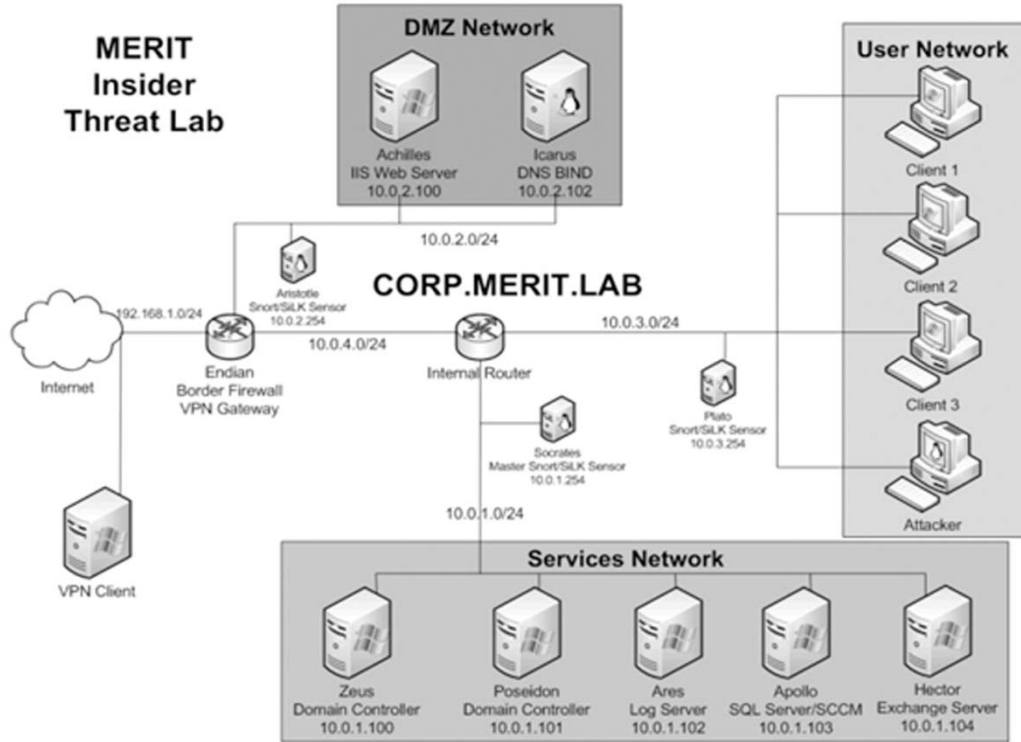


Figure 2: Sample Lab Topology for Creating Demonstrations and Testing Candidate Technical Controls and Indicators

The demonstration shows two methods of using Splunk, or an equivalent tool set, to implement a query based on tracking email by volume and destination from employees who have accounts set to expire on a certain date, as well as queries that retrieve the prior 30 days' worth of email traffic for an insider whose account is disabled, as shown in Figure 3. The demonstration goes on to show how, through creating simple queries in the tool based on information derived from prior modeling work, we can dramatically narrow the scope of our investigation to a handful of email messages sent in a short period to a set of undesirable message recipients. This immediately narrows the security operators' search space from potentially millions of email messages to a much more manageable set associated with an individual who is likely to steal information within a very specific timeframe with a high degree of confidence. The demo also allows an instructor to pause at various points during the demonstration to engage in discussion with the audience about various ways of implementing the suggested queries and how they can be modified to operate in alternate tool environments or use different organizational security policies. The important lesson is that the suggested controls are easily tailored for varying environments, rooted firmly in real case data, and tied closely to peer-reviewed articles on behavioral models associated with each type of crime.

Reply | Reply All | Forward | [Icons] | [Icons] | [Icons] | [Icons] | [Icons]

Splunk Alert: Last 30 days of email

splunk-alerting@corp.merit.lab

Sent: Tuesday, July 20, 2010 7:30 AM

To: Security Admin

Attachments: [Icon] splunk-results.csv (2 KB)

Saved search results.

Name: 'Last 30 days of email'

Query Terms: 'host=HECTOR [search host="ZEUS.corp.merit.lab" Message="A user account was disabled.**" | eval Account_Name=mvindex(Account_Name, -1) | fields Account_Name | streat Account_Name "@corp.merit.lab" sender_address | fields - Account_Name] total_bytes > 50000 AND recipient_address!="*corp.merit.lab" startdaysago=30 | fields client_ip, sender_address, recipient_address, message_subject, total_bytes'

Link to results: <http://ARES:8000/app/search>

[/@go?sid=scheduler_admin_search_TGFzdCAzMkYXZlIG9mIGVtYWls_at_1279636200_1060779988](#)

Alert was triggered because of: 'Saved Search [Last 30 days of email]: number of events(6)'

Search results attached:

Figure 3: Screenshot from Demo Video Showing an Alert from a Splunk Rule Derived from Models of Insider Theft of IP

7 Linking Vulnerability Metrics

A key component to a successful insider threat strategy includes asking organizations to identify critical assets and, therefore, where controls such as those we are interested in creating should be applied. Complementary work at CERT aims to quantify an organization's vulnerability to insider threat based on an assessment of the organization's security posture and known information about how insiders have previously exploited organizational weaknesses. This work benefits from the fact that the over 400 existing cases of insider crime have been analyzed individually to create a set of over 4,000 observed insider exploits and organizational vulnerabilities that contribute to insider crime. Current work here involves creating a firm taxonomy that can be used to accurately classify the vulnerabilities and exploits into a useful hierarchy.

We suspect there are strong patterns in this data, both in the types of technical indicators that appear in pairs or triplets as part of an insider attack and in behavioral indicators that appear in patterns that non-IT security departments should be able to identify and report to concerned parties. Our intent for our control development and measurement work is for each to provide input to one another from which we can create controls focused on areas of particular weakness and vulnerability. We also intend for this work to help organizations demonstrate, in a quantifiable way, how they are effectively using organizational resources to measurably reduce their vulnerability to insider attack. This work is also fundamental to creating intelligent analysis tools for automated detection of indicators potentially signaling an imminent or in-progress attack by an insider.

8 Limitations

While the preceding discussion and proposed method for creating candidate technical controls are based on actual case data and on well-received socio-technical models of varying types of insider crime, it is worth restating that the controls and indicators that are in development are considered to be candidates in that they are, as yet, untested. Possible concerns that should be considered when piloting a candidate control or indicator include several important issues. First, the number of false positives generated by an alert created using a proposed control or indicator may be high, creating additional work for operators. While a valid concern, this can be mitigated by the suggestion that any insider threat defensive strategy should be multi-faceted, relying on no single alert from a monitoring tool or a single report for a human being in the organization. In concert with other indicators, reports, and concerns, a candidate indicator can be more useful and likely report to an operator with more accuracy than it would on its own.

An additional point of concern involves a discussion of the source material itself. While the CERT Program's insider threat case library includes over 400 cases of actual insider crimes, it is important to consider that this library by no means represents all insider threat cases. As already noted, insider crimes go unreported and undetected for a variety of reasons, not the least of which is lack of sufficient evidence to attribute the crime to an individual and prosecute. Further, an argument can be made that these 400 cases were identified and prosecuted because the insiders involved were somehow not as effective as those insiders who are not prosecuted due to lack of evidence or those who go undetected entirely. While these are problems associated with case collection, the 400 cases nonetheless represent a significant set of insider crimes and provide the best mechanism available to us for studying insider crime.

Candidate indicators and controls may also fail to act as a preventative measure entirely and serve only as a passive alerting mechanism to an attack in progress. While prevention is likely the preferred avenue for an organization considering these candidate controls and indicators, it should be noted that this work is not attempting to design forecasting routines for predicting insider attacks. While some rules may be preventative in that they alert an operator and allow them time to engage and stop an insider attack in progress, a passive alert may, if nothing else, enable incident response activities. In particular, the added monitoring and alerting from an informed set of candidate controls could lead to improved ability to attribute a crime to an individual and prosecute, if so desired. While future work in prevention and prediction is interesting, this work is not intended to accomplish the latter directly, and the former is likely a product of a mature insider threat strategy across the enterprise and not just from any single indicator or control.

Lastly, any discussion of a monitoring strategy should be accompanied by an equally important discussion with the organization's human resources, legal, and senior leadership. Since monitoring strategies and technical controls designed to detect malicious human behavior may tread into areas relating to employee privacy, strategies may be tempered by the laws, regulatory requirements, and other governing practices and policies at an organization. Further, the choice of controls and indicators will obviously vary depending on who has operational responsibilities for responding to the controls and indicators. For example, a counterintelligence organization may choose controls and respond to alerts differently than a typical security operations center.

Over the course of the next year, CERT plans to develop a suite of candidate indicators and controls for testing in both lab and operational environments to determine effectiveness of the controls themselves and the method by which they are created. Further testing could also lead to improved educational materials that convey the importance of a blended strategy of technical tools, organizational behavior, and policy.

9 Conclusion

While the behavioral modeling of insider crime has matured steadily in the last several years, there has been a growing need to link these findings to the creation of informed and useful technical controls for combating insider crimes. In this paper, we have discussed a simple method for extracting candidate information for technical controls from real cases of insider crime. We have also shown how this method has led to the creation of useful instructional materials in the form of demonstration videos. Finally, we discussed future work linking vulnerability metrics to technical controls to provide even more granular information to an organization deciding where to allocate resources to stop malicious insiders. These items are critical to the discussion surrounding the development of improved insider threat tools, specifically concerning development of informed indicators, triggers, and alerts in a way that does not overwhelm the organization with false positives, but rather works through alerts rooted in real case information and which are genuinely cause for concern.

We also believe successful application of the principles discussed in this paper (those concerning theft of IP and sabotage cases) could lead to interesting work in the insider threat space specific to combating national security espionage. CERT has catalogued over 120 cases of espionage to study the methods used by spies against the United States government. Further work in this area could greatly benefit counterintelligence analysts and information security personnel who are defending national security information.

References

URLs are valid as of the publication date of this document.

[Cappelli 2009]

Cappelli, D. M.; Moore, A. P.; Trzeciak, R. F.; & Shimeall, T. J. *Common Sense Guide to Prevention and Detection of Insider Threat 3rd Edition—Version 3.1*. www.cert.org/archive/pdf/CSG-V3.pdf (2009).

[Cathey 2003]

Cathey, R.; Ma, L.; Goharian, N.; & Grossman, D. “Misuse Detection for Information Retrieval Systems.” *Proceedings of the Twelfth International Conference on Information and Knowledge Management*. New Orleans, LA, March 2003. ACM Press, 2003.

[CERT 2008]

CERT & U.S. Secret Service. *Insider Threat Study*. http://www.cert.org/insider_threat/study.html (2008).

[CSO 2010]

CSO magazine; U.S. Secret Service; CERT; & Deloitte. *2010 CyberSecurity Watch Survey*. www.csoonline.com/documents/pdfs/2010CyberSecurityResults.pdf (2010).

[Herbig 2002]

Herbig, K. L. & Wiskoff, M. *Espionage Against the United States by American Citizens 1947–2001* (PERSEREC 02–5). TRW Systems, Defense Personnel Security Research Center, 2002.

[MERIT 2008]

MERIT Team. *Insider Threat Modeling at CERT*. http://www.cert.org/insider_threat/modeling.html (2008).

[Moore 2009]

Moore, A. P.; Cappelli, D. M.; Caron, T. C.; Shaw, E.; & Trzeciak, R. F. “Insider Theft of Intellectual Property for Business Advantage: A Preliminary Model.” *Proceedings of the 1st International Workshop on Managing Insider Security Threats*. West Lafayette, IN, June 2009. Springer, 2009.

[Shaw 1998]

Shaw, E.; Ruby, K. G.; & Post, J. M. “The Insider Threat to Information Systems: The Psychology of the Dangerous Insider.” *Security Awareness Bulletin* 2, 98 (1998): 1–10.

[Splunk 2010]

Splunk. *What is Splunk?* <http://splunk.com/product> (2010).

[Weiland 2010]

Weiland, R. M.; Moore, A. P.; Cappelli, D. M.; Trzeciak, R. F.; & Spooner, D. L. *Spotlight On: Insider Threat from Trusted Business Partners*.

<http://www.cert.org/archive/pdf/TrustedBusinessPartners0210.pdf> (2010).

REPORT DOCUMENTATION PAGE			<i>Form Approved OMB No. 0704-0188</i>	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE January 2011	3. REPORT TYPE AND DATES COVERED Final		
4. TITLE AND SUBTITLE Deriving Candidate Technical Controls and Indicators of Insider Attack from Socio-Technical Models and Data		5. FUNDING NUMBERS FA8721-05-C-0003		
6. AUTHOR(S) Michael Hanley				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Software Engineering Institute Carnegie Mellon University Pittsburgh, PA 15213			8. PERFORMING ORGANIZATION REPORT NUMBER CMU/SEI-2011-TN-003	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) HQ ESC/XPK 5 Eglin Street Hanscom AFB, MA 01731-2116			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12A DISTRIBUTION/AVAILABILITY STATEMENT Unclassified/Unlimited, DTIC, NTIS			12B DISTRIBUTION CODE	
13. ABSTRACT (MAXIMUM 200 WORDS) The insider threat continues to be one of the prime issues facing government entities and organizations across critical infrastructure sectors. Extensive catalogues of case material from actual insider events have been used by CERT®, part of Carnegie Mellon University's Software Engineering Institute, to create socio-technical models of insider crime to help educate organizations on the risk of insider crime. Building upon this work, this paper seeks to demonstrate how a useful method for extracting technical information from previous insider crimes and mapping it to previous modeling work can create informed candidate technical controls and indicators. This paper also shows current examples of case material and candidate indicators that have been successfully converted into well-received insider threat training modules.				
14. SUBJECT TERMS insider threat, information security, system dynamics, behavioral modeling, security controls, counterintelligence, security metrics			15. NUMBER OF PAGES 32	
16. PRICE CODE				
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	