Justification of a Pattern for Detecting Intellectual Property Theft by Departing Insiders

Andrew P. Moore
David McIntire
David Mundie
David Zubrow

March 2013

TECHNICAL NOTE
CMU/SEI-2013-TN-013

CERT® Program
Software Engineering Process Management Program

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

Copyright 2013 Carnegie Mellon University

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.

Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Department of Defense.

This report was prepared for the
SEI Administrative Agent
AFLCMC/PZE
20 Schilling Circle, Bldg 1305, 3rd floor
Hanscom AFB, MA 01731-2125

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

This material has been approved for public release and unlimited distribution except as restricted below.

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

External use:* This material 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.

* These restrictions do not apply to U.S. government entities.

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

DM-0000112

Table of Contents

Acknowledgments	v
Abstract	vii
1 Introduction	1
2 Background	3
3 Summary of the Subject Pattern	5
3.1 Context	5
3.2 Problem	5
3.3 Solution	5
3.4 Expected Benefits	6
4 Research Questions and Hypothesis	7
4.1 Research Questions	7
4.2 Hypothesis	7
5 Case Selection and Coding Procedures	10
5.1 Case Selection Criteria	10
5.2 Coding Procedures	11
6 Preliminary Analysis Results	12
7 Conclusion	15
Appendix Structure of the Solution Described by the Pattern	17
References	19

List of Figures

Figure 1:	Possible Form for Distribution of the Last Theft Times Across the Cases	8
Figure 2:	Breakdown of Insider IP Theft Cases	10
Figure 3:	The Results of the Crystal Ball Analysis	12
Figure 4:	The Cumulative Probability Function for the Resampled Data Set	13
Figure 5:	Sequence Diagram for Increased Review of Departing Insider Actions	17

Acknowledgments

The authors would like to thank our Conference on Pattern Languages for Programs (PLoP) 2012 shepherd, Lise Hvatum (Schlumberger, USA), the reviewers in our Pattern Writer's Workshop at PLoP, and our SEI colleague Lori Flynn for their valuable insights and recommendations for improving this paper; Robert Stoddard of the SEI for his insights into quantitative analysis methods and tools; and Paul Ruggiero for his excellent technical editing of this paper.

Abstract

This paper describes an analysis that justifies applying the pattern “Increased Review for Intellectual Property (IP) Theft by Departing Insiders.” The pattern helps organizations plan, prepare, and implement a strategy to mitigate the risk of insider theft of IP. The analysis shows that organizations can reduce their risk of insider theft of IP through increased review of departing insiders’ actions during a relatively small window of time prior to their departure. Preliminary research results show that approximately 70% of insider IP thieves can be caught by following the pattern’s recommendation of reviewing insiders’ actions for theft events during only the last two months of their employment. These results provide practical guidance for practitioners wishing to fine tune the application of the pattern for their organizations. “Increased Review for IP Theft by Departing Insiders” is part of the CERT® Insider Threat Center’s evolving library of enterprise architectural patterns for mitigating the insider threat, based on the Center’s collected data. The Center’s larger goal is to foster greater organizational resilience to insider threat, using repeated application of patterns from the library.

1 Introduction

This paper describes results of an investigation to justify a previously published pattern, “Increased Review for Intellectual Property (IP) Theft by Departing Insiders” [Moore 2011a].¹ Clear exposition of this pattern and our associated data analysis depends on the definition of several key terms:

- **insider:** an employee, contractor, or other trusted business partner of an organization
- **intellectual property (IP):** any information owned by the organization that the organization wishes to protect (that is, keep secret)
- **theft of IP:** any unauthorized exfiltration (copying or removal) of IP from the organization that owns that IP

Case data from the CERT® Insider Threat Center, part of Carnegie Mellon University’s Software Engineering Institute, indicate that many insiders who stole their organization’s information stole at least some of it fairly close to their date of departure. Based on this insight, the pattern calls for increased review of insiders’ actions as they leave the employment of an organization, in order to mitigate the risk of theft of IP.²

The design pattern community generally advocates that a pattern should be successfully used in a significant development context at least three times to show its efficacy and gain the pattern community’s acceptance. These uses are to be documented in the “Known Uses” section of the pattern write-up. We therefore refer to this view of patterns as the known-use view. Pattern mining in the known-use view involves examining how people have built systems in the past and capturing the essence of successful approaches in the pattern format. Patterns are not created—they are discovered.

Another view of patterns—the hypothesized-use view—claims that patterns “are specific kinds of theories and that the process of pattern mining is similar to scientific discovery” [Kohls 2009]. From this perspective, a pattern can be viewed as a hypothesis to test, rather than a tried-and-true method of solving a development problem. The advantage of this view is that pattern developers can hypothesize previously untried methods as patterns and collect evidence that the patterns would be successful in development contexts. The method does not need to have been applied. The power of the pattern construct can be brought to bear in relatively new problem domains generated by the fast pace of technological development. In these cases, patterns should be explicitly viewed as prototype patterns, or proto-patterns, until sufficient evidence of their efficacy is gathered. The sufficiency of that evidence can be judged in much the same way as patterns have been judged in the past: through social processes, perhaps as part of pattern writers’ workshops.

¹ The original pattern was published under the title *A Pattern for Increased Monitoring for Intellectual Property Theft by Departing Insiders*.

® CERT is a registered mark owned by Carnegie Mellon University.

² Future patterns will address theft of IP by insiders not departing the organization [Mundie 2012].

In this paper, we do not attempt to continue the epistemological debate about patterns. We accept as legitimate the view of patterns as testable hypotheses. To our knowledge, there has been no successful demonstration or attempted usage of the pattern “Increased Review for IP Theft by Departing Insiders.” We view this pattern as a testable hypothesis, and we have collected data and conducted analyses to better justify its application in practice. This paper presents the results of an analysis that justifies and fine-tunes the pattern’s application. The paper is organized as follows:

- Section 2 summarizes our overall, mixed methods (qualitative and quantitative) of research involving insider threat mitigation patterns.
- Section 3 provides an abstracted view of a specific pattern identified by our qualitative research.
- Section 4 presents the research question and hypothesis that drive our quantitative research.
- Section 5 describes the case selection criteria and database coding procedures for this research.
- Section 6 provides a preliminary analysis of the data collected.
- Section 7 concludes with a summary of the paper and remaining work.

The appendix describes the detailed structure of the pattern solution. It is not our goal to argue for or against the view of patterns as testable hypotheses. However, we hope that the example described in this paper convincingly demonstrates the potential expansion of patterns in this direction and encourages other researchers to similarly exploit the power of the pattern concept.

2 Background

Over the last decade, the CERT Program at the Software Engineering Institute, part of Carnegie Mellon University, has cataloged hundreds of cases of malicious insider crimes adjudicated in U.S. courts. Insiders include current or former employees, contractors, and other business partners—anyone with authorized access to an organization’s systems beyond that provided to the general public. Malicious insider threat is the potential harm from insiders intentionally using or exceeding their authorized access in a way that damages the organization. This definition includes individuals who do harm by misusing their legitimate access privileges or by taking advantage of their knowledge of the organization and its systems.³ Based on our system dynamics modeling work and our analysis of cases, we have found that different classes of insider crimes exhibit different patterns of problematic behavior and mitigating measures [Cappelli 2009]. We have identified four primary categories of malicious insider threat cases: IT sabotage, fraud, theft of IP, and national security espionage.

Our research at the CERT Insider Threat Center uses a mixed-methods approach, which involves both qualitative and quantitative research. As described by Creswell and Clark, mixed-methods research may be appropriate when investigators do not know the exact questions to ask, variables to measure, and theories to guide the study, possibly due to the novelty of the research topic [Creswell 2011]. This is the case for our research. Creswell states that “in these situations, it is best to explore qualitatively to learn what questions, variables, theories, and so forth need to be studied and then follow up with a quantitative study to generalize and test what was learned from the exploration.”

The bulk of our previous research on insider threat is qualitative and exploratory [Cappelli 2009, Moore 2011b, Hanley 2010, Hanley 2011] and has applied the multiple (or comparative) case study method described by Robert Yin [Yin 2009]. The cases we included in our database fit the above definition of malicious insider threat and meet the following criteria:

- The crime occurred in the United States.
- The subject of the crime was adjudicated in a U.S. court (includes cases where the subject admitted to aspects of the crime in a plea agreement).
- Sufficient quantities and quality of data were available to understand the nature of the case.

We identified these cases from public reporting and included primary source materials, such as court records in criminal justice databases (found through searches on LexisNexis court databases), and secondary source materials, such as media reports (found through searches on LexisNexis news databases and internet search engines such as Google).

Our research has not shown insider threats to be technically sophisticated attacks traversing strategically layered countermeasures or a complex back-and-forth between attacker and

³ This definition does not include individuals who damage the organization unintentionally. While the inadvertent insider threat is important, it is beyond the scope of this work.

defender.⁴ However, insider threat defense needs to be broad because of the insider's authorized physical and logical access to the organization's systems and intimate knowledge of the organization. Current solutions to insider threat are largely reactive and tactical, and they do not address the architectural needs demanded by the holistic nature of the problem. As a result, the sensitive and possibly classified information stored on organizations' information systems is highly vulnerable to disgruntled employees, who may seek revenge for a perceived injustice, or greedy employees, who may take advantage of organizational information for their own personal gain.

Our analysis of the insider threat case data has identified more than 100 categories of weaknesses in systems, processes, people, or technologies that allowed insider attacks to occur. Many of these weaknesses are due to failures during the system or software development lifecycle that are then perpetuated by failures associated with people, processes, and technology. Research at the CERT Program is identifying enterprise architectural patterns that protect against the insider threat to organizational systems. Enterprise architectural patterns are organizational patterns that involve the full scope of enterprise architecture concerns, including people, processes, technology, and facilities. This broad scope is necessary because insiders have authorized online and physical access to systems. In addition, our data suggest that malicious insiders have exploited vulnerabilities in organizational business processes as often as they have exploited technical vulnerabilities.

The CERT Program is developing a library of insider threat enterprise architectural patterns [Mundie 2012] based on the data we have collected and our previous qualitative analyses and previous work documenting security patterns [Schumacher 2006, Hafiz 2012]. Our previous research has generated strong, specific hypotheses for a follow-on quantitative and explanatory investigation of the pattern "Increased Review for IP Theft by Departing Insiders," which is the subject of this paper.

⁴ Insider IT sabotage events are among the most technically sophisticated malicious insider threats. However, this paper focuses on insider theft, which is typically not very technically sophisticated.

3 Summary of the Subject Pattern

The “Increased Review for IP Theft by Departing Insiders” pattern helps an organization plan, prepare, and implement a strategy to mitigate the risk of insider theft of IP. This section provides a summary of that pattern, the full details of which are provided in our PLoP 2011 paper [Moore 2011a].

Insider threat case data show that risk of insider theft of IP is greatest at the point of employee departure. This pattern helps reduce that risk through increased review of insiders’ actions as they leave the employment of an organization. This increased review is above and beyond what might be required for an organization’s baseline detection of potentially malicious insider actions. The intended audience of this pattern is data owners within an organization—those who make decisions about the protection requirements for certain data, including who has access to it—as well as managers of departments across the organization: information technology, human resources, physical security, and legal. The pattern applies to organizations large enough to have these distinct departments and roles. However, smaller organizations may also benefit from this pattern if they can identify individuals with the associated responsibilities.

3.1 Context

The context for this problem is an organization that has valuable IP at risk of insider theft. IP includes any of an organization’s sensitive or confidential information that it would like to protect. An insider of an organization includes any employee, contractor, or other business partner of that organization. The organization’s critical point of action is when an insider is being terminated, either voluntarily (resigning) or involuntarily (firing).

3.2 Problem

How can the organization cost-effectively mitigate the risk of losing its IP? Data on 48 cases of theft of IP, from our insider threat database, show that over 50% of the insiders stole at least some of the information within 30 days of their departure. Current case trends suggest that organizations regularly fail to detect theft of IP by insiders, and even when theft is detected, organizations find it difficult to attribute the crime to any specific individual.

The solution to this pattern is affected by the following forces: cost of insider action review, employee privacy, IP ownership rights, employee productivity during the period between resignation and departure, and legal propriety of insider action review.

3.3 Solution

To deal adequately with the risk that departing insiders might take valuable IP with them, the organization must ensure that the necessary agreements are in place (IP ownership and consent to review), IP is identified, the activities of key departing insiders are reviewed, and the necessary communication among departments takes place. When an insider resigns, the organization should increase its scrutiny of that employee’s activities within a well-defined window before the insider’s departure date. Computer audit logs of employee online actions must be kept for at least

the length of the review window so that those logs may be scrutinized even if an insider terminates employment immediately. Actions taken upon and before employee departure are vital to ensuring that IP is not compromised and the organization preserves its legal options.

The HR department needs to track insiders who have access to the IP so that when the insider resigns, HR can ask IT staff or systems to review that insider's online behavior for signs of suspicious exfiltration of IP. IT staff or systems need to closely review the insider's access to IP during the review window before departure because many IP thieves have stolen information within this window. Although the organization may decide to begin review before the review window, restricting review to this period may help the organization balance the review costs with the risks of losing the IP. IT staff or systems must inform the data owners of any suspicious access to IP, and the data owners must be included in the decision-making on how to respond. The structure of the proposed solution is described in more detail in the appendix.

3.4 Expected Benefits

The primary expected benefit of the "Increased Review for IP Theft by Departing Insiders" pattern is that review of departing insiders is tailored to ensure a good cost-benefit ratio, while keeping insiders productive during their final days at work.

4 Research Questions and Hypothesis

While the next section describes the research methods we used to investigate the cases of theft of IP, this section discusses the driving research questions and hypothesis.

4.1 Research Questions

The primary question that the “Increased Review for IP Theft by Departing Insiders” pattern addresses is the following:

Primary Question: *What percentage of insider IP thefts could an organization expect to detect by increased review of insider actions during a relatively small period prior to departure?*

The higher the percentage of insider IP thefts detected when reviewing activity in the review window in the insider threat cases in our study, the better the prospects are for the subject pattern to detect insider IP theft incidents in a particular organizational context. If those incidents can be detected prior to departure, the organization has a better chance of addressing them, as described by the pattern, before any damage occurs. Thus, the higher the percentage of thefts detected, the better the justification of the mitigation pattern.

A secondary question allows better understanding of how organizations can detect the last theft event:

Secondary Question: *How can organizations distinguish between insider theft activity and legitimate employee activity?*

An answer to this question would help an organization use the mitigation pattern cost effectively by reducing the chance of false positives during the review process. This question is beyond the scope of our current investigation, but it will be the subject of future research.

4.2 Hypothesis

Critical to the statement of our research hypothesis is the notion of a theft-of-IP event (also referred to as a theft event): a single unauthorized transfer of IP off the organization’s network systems or outside the organization’s physical boundaries. Of course, a case may include a series of related theft events, but it is the last such theft event in the case that is of specific interest in our study.

Past qualitative analyses of our insider threat data have suggested that the timing of the last theft events prior to departure will conform to a nonlinear distribution, as exemplified in Figure 1. This makes sense: the approach of the date of departure accelerates the insider’s decision-making process in a nonlinear progression. Therefore, our hypothesis is the following:

Hypothesis: *The distribution of the times between the following two dates conforms to a nonlinear distribution:*

- *the date of the last confirmed theft-of-IP event by an insider prior to that insider’s departure*
- *the date of the insider’s departure*

Figure 1: Possible Form for Distribution of the Last Theft Times Across the Cases

The distribution of times between an insider IP thief's last confirmed theft-of-IP event before departure and the date of the insider's departure is of more than abstract, theoretic interest. The precise nature of the distribution will help characterize the efficacy of the pattern by allowing an estimate of the percentage of last confirmed insider IP thefts that are detectable by the pattern. Section 5 demonstrates how this is accomplished.

The above hypothesis make several assumptions.

Assumption 1: The primary harm associated with theft of IP occurs after departure.

We are primarily concerned with catching the theft at any point prior to departure rather than at its earliest occurrence. This is an assumption about the application context, because if the harm occurs before departure, the pattern will do little to prevent the harm, though it could still aid recovery.

Assumption 2: Insider actions are logged continuously throughout their employment, but those logs cannot be practically reviewed for suspicious behavior of all employees with full intensity all the time.

Organizations can log insider actions automatically at relatively little cost. Review of those logs, which involves analyzing logs for suspicious activities, is often primarily a human activity, so it is more expensive and requires a higher level of discretion. While continuous review involves some baseline level of scrutiny of audit logs performed all the time, targeted review involves a heightened level of scrutiny performed on a subset of organizational logs. The selection of the subset may be based on chronology (for example, heightened scrutiny during a merger) or on risk

assessment (for example, heightened scrutiny of an employee who has displayed negative workplace behavior).

Assumption 3: There is advance notice that an insider is departing the organization, either from the insider's resignation or the organization's plan to fire the insider.

The study described in this report involves reviewing insider actions during the final weeks before departure. By departure we mean the end of an insider's employment by an organization, either voluntarily (an individual quitting) or not (an individual being forcibly removed). We refer to resignation as an insider's notification to the employer of his or her intention to leave his or her job at a designated point in the future. Resignations are not always strictly voluntary, as in a "forced resignation."

It is at the point of advance notice of an insider's departure that the organization takes additional actions to more intensively review the insider's behavior, possibly both retrospectively, in the time leading up to the advance notice, and prospectively, between the advance notice and final departure.

5 Case Selection and Coding Procedures

5.1 Case Selection Criteria

The “Increased Review for IP Theft by Departing Insiders” pattern involves only cases of IP theft that are in the CERT insider threat database. At the time the analysis was conducted, the database contained a total of 93 insider theft of IP cases, 26 of which we identified as appropriate for studying the subject pattern. The criteria for selecting these cases were the following:

- The case had to involve theft of IP. We do not include cases of national security espionage—which involve the theft of U.S. Government classified or controlled information—because these cases are typically much more sensitive.
- The theft of IP had to take place before the insider departed. Departure is defined as the end of an insider’s employment at an organization. Departure is either voluntary (an individual quitting) or involuntary (an individual being forcibly removed).
- The insider had to have departed before the theft was discovered.

Figure 2 shows the breakdown of cases into the following categories:

- **IP Theft Cases Not Suitable for Coding:** These cases of IP theft either did not have the required sequence of events or the event sequence was not recorded in the case file.
- **IP Theft Cases with Time Data Not Present:** These cases fit our criteria for analysis, but date information was not present, was incomplete, or was inexact.
- **Codeable Cases:** These cases met all prerequisites for coding and had time data with a margin of error of less than one day.

Figure 2: Breakdown of Insider IP Theft Cases

After selecting this subset of cases from the database, we searched for cases that had the following data points available either in the insider threat database or in case data:

- date of last observable theft event
- date of departure

We used currently available date information in the insider threat database, court documents, and media reports. In most cases, the dates were listed in only one source and could not be checked against other data. If any margin of error greater than one day was indicated in documentation, the case was not used in further research.

5.2 Coding Procedures

Case coding is a critical process in which information gathered through review of case file documents is entered into the case database according to a prescribed methodology documented in a codebook.⁵ The codebook provides operational definitions and examples of all the required items. The codebook for our current research asks the following questions about IP theft cases:

- On what date did the last theft of IP occur?
- What is the time period (in number of days) between the last observable theft event and departure?
- What is the overall margin of error we can assert regarding the dates listed?⁶

At the beginning of the coding process, all date information was contained in the Chronology section of the database. For this study, the CERT Program created a new section in the Theft of IP codebook that records the above data points and makes them easy to query.

Because reliability is important for all types of data collection, we develop, test, and follow specific procedures to ensure that the data is collected and coded consistently and predictably. To address inter-rater reliability, coders are briefed on the codebook's conceptual framework and typology. Coders are given one or more of the same cases to code independently, and their coding is compared to a master coding established by the research team. We use attribute agreement analysis to quantify the consistency of judgment within and between different coders. Attribute agreement analysis constructs and implements a brief experiment in which coders participate in a short case-coding exercise.⁷

⁵ McIntire, D.; Moore, A. P.; Mundie, D. A.; & Flynn, L. *Coding Guide – Multiple Case Study of Increased Review for Intellectual Property Theft by Departing Insiders* (Special Report). Software Engineering Institute, Carnegie Mellon University, forthcoming.

⁶ Margin of error is defined as the estimated error of a date specification. The margin of error is expressed as a number of days by which the date may be incorrect. Dates representing an exact time on a specific day have a margin of error of 0.

⁷ Personal communication with Robert Stoddard, member of the Software Engineering Measurement and Analysis department of the Software Engineering Institute.

6 Preliminary Analysis Results

To evaluate the hypothesis, we entered our data into the Crystal Ball software package to find the best-fit distribution to the observed data.⁸ The software considered binomial, discrete uniform, Poisson, negative binomial, and geometric distributions. As shown in Figure 3, the best fit was a geometric distribution with a probability (p) value of 0.02049. However, the chi square goodness-of-fit statistic had a p -value of 0.017, indicating that the distribution of observed data was statistically different from the theoretical distribution. Nonetheless, given the small number of cases available, we chose to use this best-fit distribution as the basis for the following resampling method.⁹

Note: On the left, the gray bars represent our data, while the green line is a geometric distribution with $p=0.02049$.

Figure 3: The Results of the Crystal Ball Analysis

To better understand how an organization can use information regarding the above distribution, we used the Crystal Ball software to run a Monte Carlo simulation that generated 1,000 resampled data sets from the geometric distribution with $p=0.02049$. Using the geometric distribution instead of the observed data allowed for the possibility that unobserved values of durations from last IP theft to departure could occur. From those results, we graphed the cumulative probability function shown in Figure 4. This graph makes it straightforward to predict what percentage of last theft events would be detected as a function of the duration of the targeted review. These preliminary results confirm our secondary hypothesis: on average almost exactly 70% of last theft events

⁸ See <http://www.oracle.com/us/products/applications/crystalball/crystalball-066563.html>.

⁹ This paper describes preliminary research results that characterize what we expect to see in our final analysis. Future research will add cases to better identify the underlying distribution, and we will refine our analysis accordingly.

could be detected by targeted review of the insider's activities during the last 60 days of employment.

Figure 4: The Cumulative Probability Function for the Resampled Data Set

Our preliminary analysis allows organizations to explore tradeoffs between cost of review and probability of discovering IP theft by departing insiders. For instance, the graph in Figure 4 shows that slightly more than 30% of last theft events could be detected with just three weeks of targeted review. The graph also highlights the dramatic decrease in the return on review investment per theft detected when review is conducted more than about 130 days before departure. In sum, our analysis supports our position that targeted review can improve the efficiency of detecting theft of IP.

It is important to emphasize the limitations of our data analysis to date. The cases we are analyzing were prosecuted in U.S. courts. Due to different national laws and norms, our results may not be representative of what happens in other countries where, for example, the period between resignation and departure could be somewhat longer. In addition, our data analysis and results are preliminary partly because of the small number of cases in our data set. While the best-fit distribution was the geometric distribution (as compared to a wide variety of other distributions), the fit was statistically different from the theoretical distribution. While future research will continue to work to add additional cases to better identify the underlying distribution and refine our analysis, the resampling approach described above allowed us to use the data that we had to greatest effect.

The result was a more conservative (larger) estimate of the size of the review window needed to detect a certain percentage of the last theft events associated with the theft of IP cases. This more conservative estimate was justified given the uncertainties associated with the small sample size.

Could the best-fit distribution be wrong? Yes, it could. However, we contend that the guidance coming from the resampling is surely better and more robust as compared to that based on the raw data alone. We also contend that the geometric distribution being the best fit for the data provides at least *prima facie* evidence that the subject mitigation pattern will be effective in fighting insider theft of IP. Continuing research will strive to bolster this evidence.

7 Conclusion

This paper provides partial justification for the “Increased Review for IP Theft by Departing Insiders” pattern. We have summarized our overall mixed-methods research approach, described how our qualitative research has identified specific questions for investigation, and focused on our quantitative data collection to test the validity of our hypothesis. In particular, our analysis confirms our initial assumptions that a significant percentage of last confirmed IP theft events takes place in a narrow window before insider departure. Insiders stealing IP did so within a period of 60 days before departure 70% of the time. All of the studied cases produced technical observable events, the review of which by an organization could help stop or remediate the theft.

Future work will seek to

- better understand the detectability of the last theft events (see focus question 2)
- increase the volume of data available for research not in the number of cases but also in the characteristics being coded for each case
- conduct a more extensive attribute agreement analysis to assess and improve case coding consistency
- bring findings up to date with new IP theft cases from the CERT insider threat database
- analyze differences in insider behaviors when insiders voluntarily resign compared to when they are fired or forced to resign;¹⁰ one might expect that the pattern is less effective in the latter case
- directly test the efficacy of the mitigation pattern within organizations

Additional attribute agreement analysis will involve multiple experienced coders, a larger sampling of coded cases undergoing analysis, and a more refined Pass-Fail process to determine whether a case can be used further. Testing the efficacy of the mitigation pattern directly within organizations is a necessary step in the process of validating the pattern, but it is difficult partly due to the relatively low frequency of insider attacks and partly due to the potential disturbance that can occur, e.g., the increased scrutiny and potential false accusation of innocent insiders. The intent of this paper is to gather evidence of the pattern’s efficacy prior to testing within organizations.

For those responsible for an organization’s information system security, this paper has provided information—by way of the cumulative probability function in Figure 4—to fine tune their application of the subject mitigation pattern. Depending on their organization’s tolerance to the risk of insider theft of IP, they can choose a review window that detects an acceptable percentage of insider thieves, thus balancing the risk with the costs of review of insider activities. In conclusion, we expect that architectural patterns and pattern systems developed through this research will enable coherent reasoning about how to design and, to a lesser extent, implement

¹⁰ Note that this pattern deals with departures in which the organization does not escort insiders off the premises after resignation, keeping insiders productive as long as possible. Other mitigation patterns will address organizations removing the insider’s access immediately after resignation to reduce insider threat risk.

enterprise systems to protect against insider threat [Mundie 2012]. Instead of working with vague security requirements and inadequate security technologies, system designers will have a coherent set of architectural patterns. So armed, they will be able to develop and implement effective strategies against the insider threat more quickly and with greater confidence.

Appendix Structure of the Solution Described by the Pattern

Figure 5 illustrates the structure and dynamics of the solution as a sequence diagram for the pattern “Increased Review for Intellectual Property (IP) Theft by Departing Insiders.” The actors in the solution appear along the top of the diagram, starting on the left with the insider. Next come the main organizational departments or groups involved: human resources (HR), data owners, and information technology (IT) staff or systems. The IT systems could play the IT-related role if the review role is automated, or IT staff could if the review is manual. The IP itself appears on the far right of the diagram.

Figure 5: Sequence Diagram for Increased Review of Departing Insider Actions¹¹

The solution dynamics are portrayed in the interactions among the actors with the roles and responsibilities listed along the top of Figure 5. An organization needs to make sure its employees, as a condition of employment, consent to review (Interaction 1) and agree that the organization owns the IP (Interaction 2). The employee’s clear and formal acceptance of the organization’s IP ownership helps ensure that the organization’s right to ownership will stand up in court. Consulting with the organization’s legal counsel will help ensure the organization is on firm legal ground. The organization can convey ownership to employees through devices such as nondisclosure agreements, IP ownership policies, and references to IP ownership in a network-acceptable-use policy.

We assume that data owners identify and properly label their IP. HR needs to track insiders who have access to the IP so that when the insider resigns, HR can ask IT staff or systems to review

¹¹ Interaction 5: If the insider engages in a suspicious access, then the IT staff or system communicates that access to the data owners.

that insider's online behavior for signs of suspicious exfiltration of IP (Interaction 3). Data in our insider threat database shows that among insider IP thieves, scientists, engineers, programmers, and salespeople are especially likely to steal IP. IT staff or systems need to closely review the insider's access to IP during the review window before departure (Interaction 4) because many IP thieves have stolen information within this window. Although the organization may decide to begin review before the review window, restricting review to this period may help the organization balance the review costs with the risks of losing the IP. No matter what level of review organizations use, they must ensure that insiders are treated consistently and fairly.

IT staff or systems must inform the data owners of any suspicious access to IP, and the data owners must be included in the response decision-making (Interaction 5). The organization must be able to either block exfiltration or detect it and confront the employee. If the suspicious activity occurs prior to departure, HR and the data owners need to formulate an appropriate response as part of the departure plan (Interaction 6). The organization can then confront the employee with that response during the exit (departure) interview (Interaction 7). If the insider has violated an agreement regarding IP, the organization may wish to pursue legal remediation, with advice from legal counsel.

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 Threats, 3rd Edition – Version 3.1*. Software Engineering Institute and CyLab, Carnegie Mellon University, 2009.
<http://www.cert.org/archive/pdf/CSG-V3.pdf>

[Creswell 2011]

Creswell, J. W. & Clark, V. L. P. *Designing and Conducting Mixed Methods Research, 2nd Edition*. Sage Publications, 2011.

[Hafiz 2012]

Hafiz, M.; Adamczyk, P.; & Johnson, R. “Growing a Pattern Language (for Security),” 139–158. *Proceedings of the ACM International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software (Onward! 2012)*. Tucson, AZ, Oct. 19–26, 2012. ACM, 2012.

[Hanley 2010]

Hanley, M. “Candidate Technical Controls and Indicators of Insider Attack from Socio-Technical Models and Data.” *2010 NSA CAE Workshop on Insider Threat*. St. Louis, MO, Nov. 14, 2010.

[Hanley 2011]

Hanley, M.; Dean, T.; Schroeder, W.; Houy, M.; Trzeciak, R. F.; & Montelibano, J. *An Analysis of Technical Observations in Insider Theft of Intellectual Property Cases* (CMU/SEI-2011-TN-006). Software Engineering Institute, Carnegie Mellon University, 2011.
<http://www.sei.cmu.edu/library/abstracts/reports/11tn006.cfm>

[Kohls 2009]

Kohls, C. & Panke, S. ““Is That True ...?” Thoughts on the Epistemology of Patterns.” *Proceedings of the 16th Conference on Pattern Languages for Programs (PLoP ’09)*. Chicago, IL, Aug. 28–30, 2009. ACM, 2009.

[Moore 2011a]

Moore, A. P.; Hanley, M.; & Mundie, D. “A Pattern for Increased Monitoring for Intellectual Property Theft by Departing Insiders.” *Proceedings of the 18th Conference on Pattern Languages for Programs (PLoP ’11)*. Portland, OR, Oct. 21–23, 2011. ACM, 2011.

[Moore 2011b]

Moore, A. P.; Cappelli, D. M.; Caron, T. C.; Shaw, E.; Spooner, D.; & Trzeciak, R. F. “A Preliminary Model of Insider Theft of Intellectual Property.” *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications* 2, 1 (Special Issue Addressing Insider Threats and Information Leakage, 2011): 28–49.

[Mundie 2012]

Mundie, D. A. & Moore, A. P. “Multi-Dimensional Pattern Languages: Growing a Pattern Language for Insider Threat.” *Proceedings of the 19th Conference on Pattern Languages for Programs (PLoP '12)*. Tucson, AZ, Oct. 19–21, 2012. ACM, 2012.

[Schumacher 2006]

Schumacher, M.; Fernandez-Buglioni, E.; Hybertson, D.; Buschmann, F.; & Sommerlad, P. *Security Patterns: Integrating Security and Systems Engineering*. John Wiley & Sons, Ltd., 2006.

[Yin 2009]

Yin, R. K. *Case Study Research: Design and Methods, 4th Edition*. Sage Publications, 2009.

REPORT DOCUMENTATION PAGE			<i>Form Approved</i> OMB No. 0704-0188	
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 March 2013		3. REPORT TYPE AND DATES COVERED Final
4. TITLE AND SUBTITLE Justification of a Pattern for Detecting Intellectual Property Theft by Departing Insiders			5. FUNDING NUMBERS FA8721-05-C-0003	
6. AUTHOR(S) Andrew P. Moore, David McIntire, David Mundie, David Zubrow				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Software Engineering Institute Carnegie Mellon University Pittsburgh, PA 15213			8. PERFORMING ORGANIZATION REPORT NUMBER CMU/SEI-2013-TN-013	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFLCMC/PZE/Hanscom Enterprise Acquisition Division 20 Schilling Circle Building 1305 Hanscom AFB, MA 01731-2116			10. SPONSORING/MONITORING AGENCY REPORT NUMBER n/a	
11. SUPPLEMENTARY NOTES				
12A DISTRIBUTION/AVAILABILITY STATEMENT Unclassified/Unlimited, DTIC, NTIS			12B DISTRIBUTION CODE	
13. ABSTRACT (MAXIMUM 200 WORDS) This paper describes an analysis that justifies applying the pattern "Increased Review for Intellectual Property (IP) Theft by Departing Insiders." The pattern helps organizations plan, prepare, and implement a strategy to mitigate the risk of insider theft of IP. The analysis shows that organizations can reduce their risk of insider theft of IP through increased review of departing insiders' actions during a relatively small window of time prior to their departure. Preliminary research results show that approximately 70% of insider IP thieves can be caught by following the pattern's recommendation of reviewing insiders' actions for theft events during only the last two months of their employment. These results provide practical guidance for practitioners wishing to fine tune the application of the pattern for their organizations. "Increased Review for IP Theft by Departing Insiders" is part of the CERT® Insider Threat Center's evolving library of enterprise architectural patterns for mitigating the insider threat, based on the Center's collected data. The Center's larger goal is to foster greater organizational resilience to insider threat, using repeated application of patterns from the library.				
14. SUBJECT TERMS Security, insider threat, security pattern, organizational pattern, enterprise architecture, information security, information assurance, cyber security			15. NUMBER OF PAGES 31	
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	