# Effectiveness of a Pattern for Detecting Intellectual Property Theft by Departing Insiders

ANDREW P. MOORE, CERT® Program, Software Engineering Institute
DAVID MCINTIRE, CERT Program, Software Engineering Institute
DAVID MUNDIE, CERT Program, Software Engineering Institute
DAVID ZUBROW, Software Engineering Measurement and Analysis, Software Engineering Institute

**Abstract.** This paper describes the effectiveness of the pattern "Increased Review for Intellectual Property (IP) Theft by Departing Insiders," which helps organizations plan, prepare, and implement a strategy to mitigate the risk of insider theft of intellectual property (IP). The CERT® Insider Threat Center's case data shows that many insiders who stole IP did so fairly close to their termination. Based on this insight, this pattern helps reduce the risk of insider theft of IP through increased review of departing insiders' actions during their last 60 days of employment. Preliminary research results show that approximately 70 percent 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 an evolving library of enterprise architectural patterns for mitigating the insider threat, based on our collected data. Our 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 determine the effectiveness of a previously published pattern, "Increased Review for Intellectual Property (IP) Theft by Departing Insiders," which helps organizations plan, prepare, and implement a strategy to mitigate the risk of insider theft of IP [Moore et al. 2011].[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 (i.e., keep secret)
- **theft of IP:** any exfiltration (copying or removal) of IP that harms the owner of the information, could harm the owner, or could benefit some party in a way that harms the owner

® CERT and CERT Coordination Center are registered marks owned by Carnegie Mellon University. The CERT Program is part of the Software Engineering Institute, a U.S. Department of Defense federally-funded research and development center.
© Copyright 2012 Carnegie Mellon University

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

The CERT® Insider Threat Center's case data indicates that many insiders who stole their organization's information stole at least some of it within 60 days of their termination. Based on this insight, the pattern calls for increased review of insiders' actions as they leave the employment of an organization.

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 and Panke 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.

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 demonstrate its potential effectiveness in practice. This paper presents evidence of the pattern's efficacy and 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 hypotheses that drive our quantitative research.
- Section 5 provides a preliminary analysis of the data collected.
- Section 6 concludes with a summary of the paper and remaining work.
- Appendix A describes the detailed structure of the pattern solution, which may aid understanding of the pattern in more detail for this paper without having to go to the full pattern exposition in Moore et al. [2011].
- Appendix B describes the case selection criteria and database coding procedures for this research.

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.[2] 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 2008]. 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 in Creswell [2011], 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. 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 et al. 2008, Moore et al. 2011, Hanley 2010, Hanley et al. 2011] and has applied the multiple (or comparative) case study method described by Robert 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).

[2] 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.

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 defender.[3] 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 suggests 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 and Moore 2012] based on the data we have collected and our previous qualitative analyses and previous work documenting security patterns [Schumacher et al. 2006, Hafiz 2011]. 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.

## 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 et al. 2011].

Insider threat case data shows that risk of insider theft of IP is greatest at the point of employee termination. 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

[3] 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.

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 (e.g., resigning) or involuntarily (e.g., firing).

### 3.2 Problem

How can the organization cost-effectively mitigate the risk of losing its critical IP? Data on 48 cases of theft of IP, from our insider threat database, shows that over 50 percent of the insiders stole at least some of the information within 30 days of their termination. 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 termination, 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), critical 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 termination 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 termination 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 critical IP during the review window before termination 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 critical IP, and the data owners must be included in the response decision-making. The complete, formal structure of the proposed solution is described in Appendix A.

### 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 HYPOTHESES

Appendix B to this paper describes the quantitative research methods we used to investigate the cases of theft of IP. This section discusses the driving research questions and hypotheses.

### 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 termination?

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 termination, 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 greater the effectiveness of the mitigation pattern.

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

> **Secondary Question:** What makes an insider's last confirmed theft-of-IP events prior to termination stand out from other, legitimate employee activity?

This question is beyond the scope of our current investigation, but it will be the subject of future research.

### 4.2 Hypotheses

Critical to the statement of our research hypotheses 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 last theft events in cases will conform to a geometric distribution of chronology, as illustrated in Figure 1. This makes sense: the approach of termination day accelerates the insider's decision-making process in a nonlinear, geometric progression. Therefore, our primary hypothesis is the following:

> ***Primary Hypothesis:*** *The distribution of the times between the following two dates conforms to a geometric distribution:*

- the date of the last confirmed theft-of-IP event by an insider prior to that insider's termination
- the date of the insider's termination

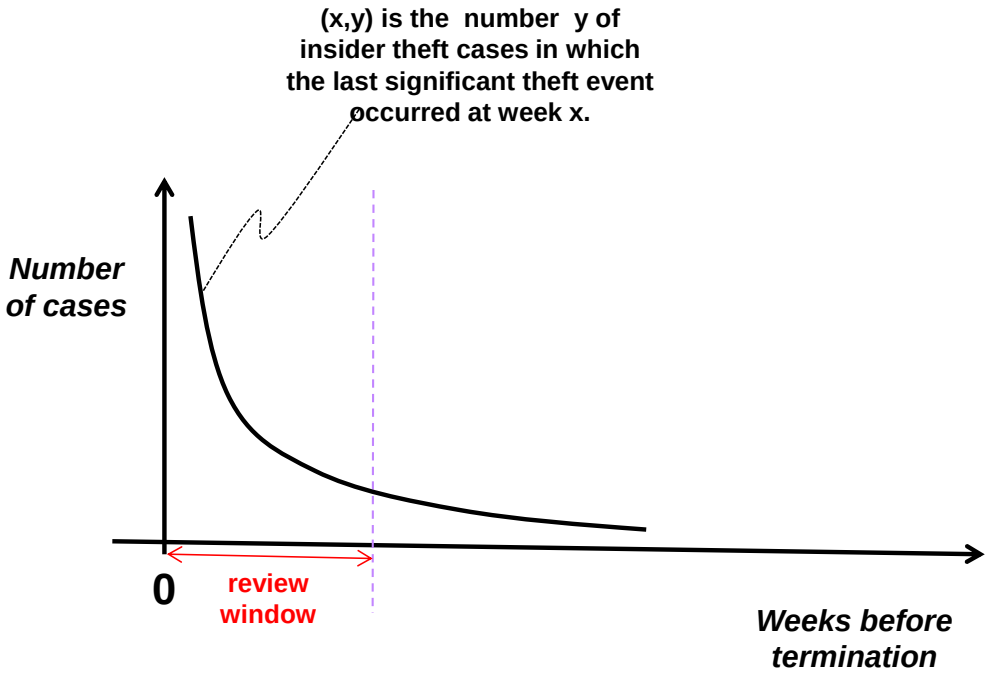


**Fig. 1.** Expected distribution of the last theft times across the cases

The geometric distribution of times between an insider IP thief's last confirmed theft-of-IP event before termination and the date of the insider's termination is of more than of abstract, theoretic interest. The precise nature of the geometric distribution will help characterize the effectiveness 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.

We have also formulated a secondary hypothesis to test that would demonstrate the concrete utility of the pattern to an organization:

> ***Secondary Hypothesis:*** *At least 70 percent of insider IP theft cases can be caught by reviewing for significant theft events by the insider during his or her last two months of employment.*

The above hypotheses make several assumptions.

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

We are primarily concerned with catching the theft at any point prior to termination rather than at its earliest occurrence. This is an assumption about the application context, because if the harm

occurs before termination, 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 (e.g., heightened scrutiny during a merger) or on risk assessment (e.g., heightened scrutiny of an employee who has displayed negative workplace behavior).

**Assumption 3:** There is advanced 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 termination. By termination 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 advanced notice of an insider's termination that the organization takes additional actions to more intensively review the insider's behavior, possibly both retrospectively, in the time leading up to the advanced notice, and prospectively, between the advanced notice and final termination.

## 5. PRELIMINARY ANALYSIS RESULTS

To evaluate the primary hypothesis, we entered our data into the Crystal Ball software package to find the best-fit distribution to the observed data.[4] The software considered binomial, discrete uniform, Poisson, negative binomial, and geometric distributions. As shown in Figure 2, the best fit was a geometric distribution with a parameter 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.

[4] See http://www.oracle.com/us/products/applications/crystalball/crystalball-066563.html.

.

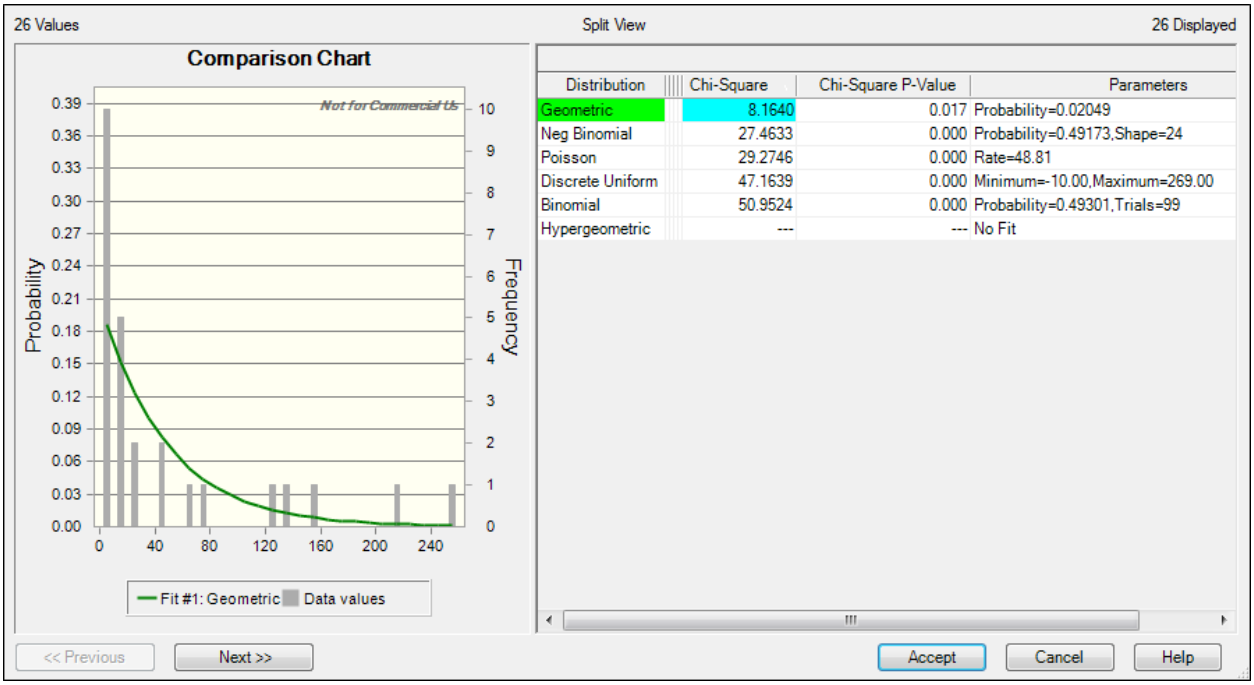


**Fig. 2.** The results of the Crystal Ball analysis. On the left, the gray bars represent our data, while the green line is a geometric distribution with $p$=0.02049

To evaluate the secondary hypothesis, 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 termination could occur. From those results, we graphed the cumulative probability function shown in Figure 3. 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 percent of last theft events could be detected by targeted review of the insider's activities during the last 60 days of employment.

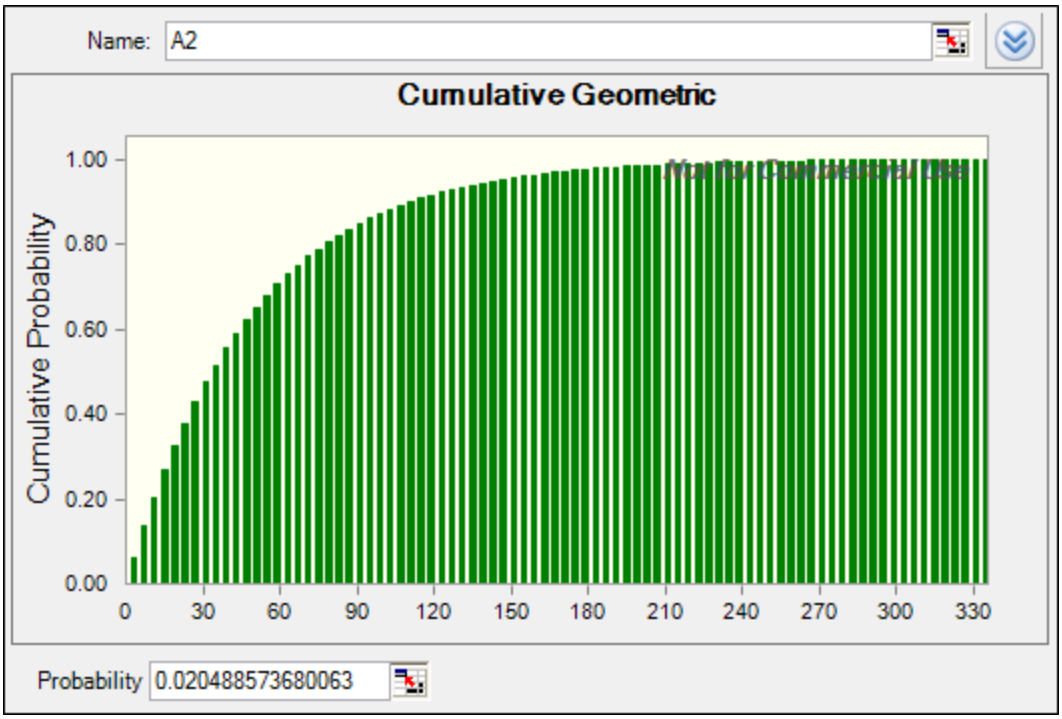


**Fig. 3.** The cumulative probability function for the resampled data set

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 3 shows that slightly more than 30 percent 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 termination. 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 termination 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.

## 6. CONCLUSION

This paper provides partial evidence for the effectiveness of 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 hypotheses for investigation, and focused on our quantitative data collection that partially supports the validity of these hypotheses. 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 termination. Insiders stealing IP did so within a period of 60 days before termination 70 percent 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;[5] one might expect that the pattern is less effective in the latter case
- directly test the effectiveness 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 effectiveness 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 as much evidence of the pattern's effectiveness as possible 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 3—to fine tune their application of the subject mitigation pattern. Depending on their organization's tolerance to the

[5] Note that this pattern deals with terminations 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.

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 enterprise systems to protect against insider threat [Mundie and Moore 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.

## ACKNOWLEDGEMENTS

The authors would like to thank our PLoP shepherd, Lise Hvatum (Schlumberger, USA) 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.

## REFERENCES

Buschmann, F.; Henney, K.; & Schmidt, D.C. *Pattern-Oriented Software Architecture Volume 5: On Patterns and Pattern Languages*. Wiley, 2007.

Cappelli, D.M.; Moore, A.P.; Trzeciak, R.F.; & Shimeall, T.J. *Common Sense Guide to Prevention and Detection of Insider Threats (3rd ed.).* Software Engineering Institute and CyLab, Carnegie Mellon University, 2008.

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

Ellison, R.J.; Moore, A.P.; Bass, L.; Klein, M.; & Bachmann, F. *Security and Survivability Reasoning Frameworks and Architectural Design Tactics* (CMU/SEI-2004-TN-022). Software Engineering Institute, Carnegie Mellon University, 2004.

---

This material is based upon work funded and supported by the United States 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.

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.

Hafiz, M.; Adamczyk, P.; & Johnson, R. “Growing a Pattern Language (for Security).” *Proceedings of the 2011 Conference on Pattern Languages for Programs*, 18th Conference on Pattern Languages of Programs (PLoP ’11), October 21-23, 2011.

Hanley, M. “Candidate Technical Controls and Indicators of Insider Attack from Socio-Technical Models and Data.” *Proceedings of the 2010 NSA CAE Workshop on Insider Threat*, November 2010.

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.

Kohls, C. & Panke, S. “‘Is that true …?’ Thoughts on the Epistemology of Patterns.” *Proceedings of the 2009 Conference on Pattern Languages for Programs*, 16th Conference on Pattern Languages of Programs (PLoP ’09), August 28-30, 2009.

McIntire, D.; Moore, A.P.; Mundie, D.A.; & Flynn, L. *Coding Guide – Study of the Effectiveness of Increased Review for Intellectual Property Theft by Departing Insiders*. Software Engineering Institute Special Report, forthcoming.

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

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, D.A. & Moore, A.P. “A Pattern for Trust Trap Mitigation.” *Proceedings of the 2011 Conference on Pattern Languages for Programs*, 18th Conference on Pattern Languages of Programs (PLoP ’11), October 21-23, 2011.

Mundie, D.A. & Moore, A.P. “Multi-Dimensional Pattern Languages: Growing a Pattern Language for Insider Threat.” *Proceedings of the 2011 Conference on Pattern Languages for Programs*, submitted to the 19th Conference on Pattern Languages of Programs (PLoP ’12), 2012.

NIST. *Special Publication 800-53, Recommended Security Controls for Federal Information Systems and Organizations, Revision 3*. August 2009. http://csrc.nist.gov/publications/nistpubs/800-53-Rev3/sp800-53-rev3-final.pdf

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

Kaarbo, J. & Beasley, R. “A Practical Guide to the Comparative Case Study Method in Political Psychology.” *Political Psychology 20*, 2 (June 1999): 369-391.

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

## APPENDIX A: STRUCTURE OF THE SOLUTION DESCRIBED BY THE PATTERN

Figure 4 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 main organizational departments or groups involved in the solution appear along the top middle of the diagram: 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 insider's involvement appears along the left side of the diagram, and the review of critical IP appears along the right side.

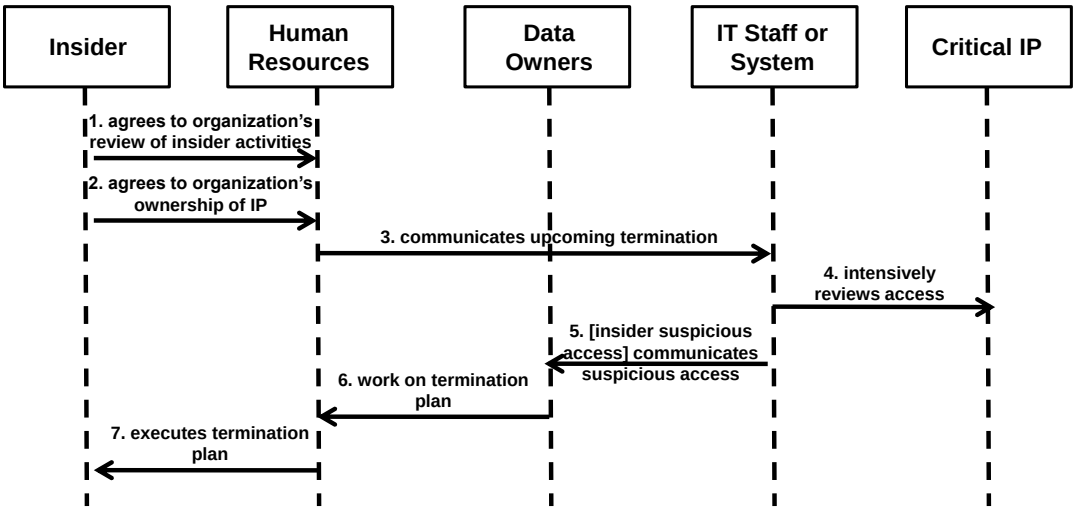


**Fig. 4.** Sequence diagram for increased review of departing insider actions[7]

The solution dynamics are portrayed in the interactions among the actors with the roles and responsibilities listed along the top of Figure 4. 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 critical 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 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 critical IP during the review window before termination (Interaction 4) because

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

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 critical 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 termination, HR and the data owners need to formulate an appropriate response as part of the termination plan (Interaction 6). The organization can then confront the employee with that response during the exit (termination) 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.

## APPENDIX B: CASE SELECTION CRITERIA AND CODING PROCEDURES

Our research method for investigating the hypotheses of this study involves the case selection criteria and the data coding procedures.

### Case Selection Criteria

The "Increased Review for IP Theft by Departing Insiders" pattern involves only the theft of IP cases in the CERT insider threat database. There is a total of 93 insider theft of IP cases in the database, of which we identified 26 that were appropriate for studying the effectiveness of 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 was terminated. Termination is defined as the end of an insider's employment by an organization. Termination is either voluntary (an individual quitting) or involuntary (an individual being forcibly removed).
- The insider had to have been terminated before the theft was discovered.

Figure 5 shows the breakdown of cases:

- **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.

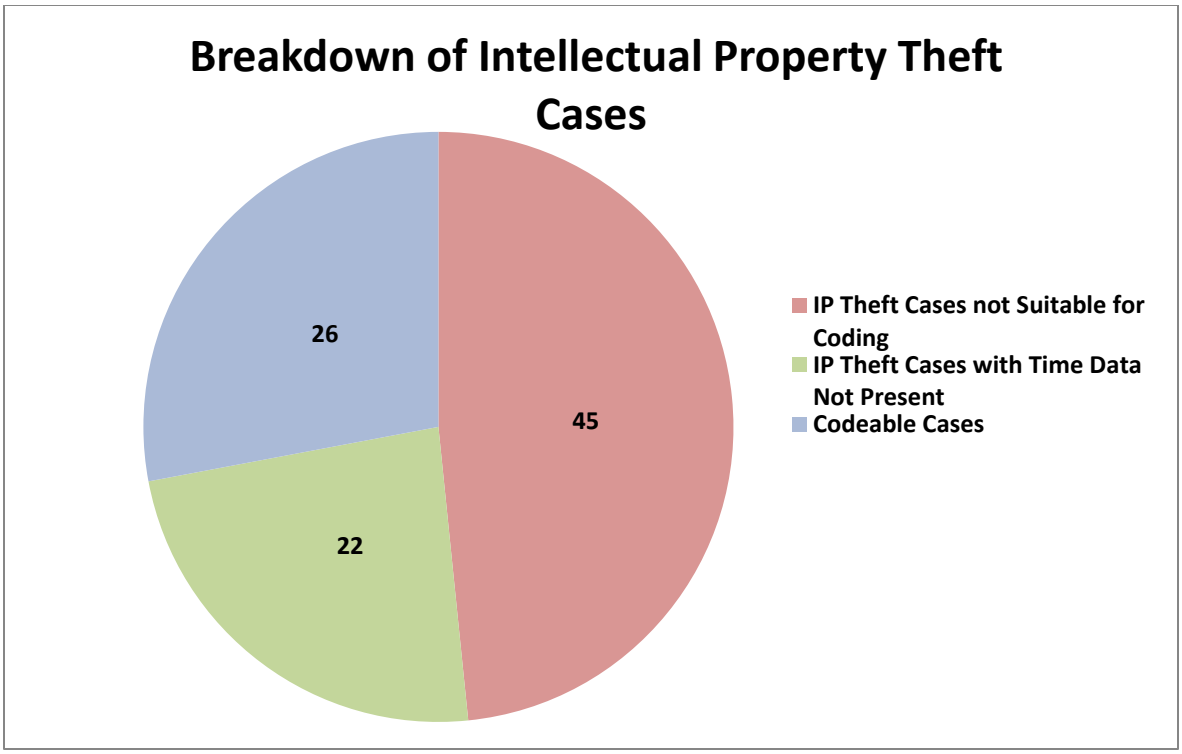


**Fig. 5.** Breakdown of insider IP theft cases

The CERT Program does track the quality of case data by using a scoring system associated with a number of important characteristics, but this score was not a factor in choosing whether a case was viable for study. Typical quality indicators like name and age were unimportant in this research due to its narrow focus.

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 termination

In collecting the above dates, 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.

Case coding is a critical process in which information gathered through case file document review is entered into the case database according to a prescribed methodology documented in a codebook (McIntire et al., forthcoming). The codebook provides operational definitions and examples of all the required items. The codebook for our current research, which is used to guide the coding process, asks the following questions (located in an IP Theft section that records information important to IP theft cases in particular):

- 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 termination?
- What is the overall margin of error we can assert regarding the dates listed?[8]

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 easily queriable.

### Coding Procedures

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.[9]

[8] 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.

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