Human Behavior, Insider Threat, and Awareness

An Empirical Study of Insider Threat Behavior [i]

Deanna D. Caputo, Greg Stephens, Brad Stephenson, and Minna Kim
The MITRE Corporation
31 July, 2009

Introduction

According to the FBI[ii], the increasing value of proprietary information and spread of enabling technologies has increased the threat posed by malicious insiders stealing information. They estimate that billions of dollars are lost by companies every year to global competitors. The problems that insiders can pose to an organization have become of greater concern and a focus of research. We were interested in better understanding the information-use behavior of malicious insiders, whom we defined as those insiders who leverage their assigned privileges to gather sensitive or proprietary information. Keep mind that an insider may be an employee, contractor, vendor, partner, or even a visitor who is provided internal access privileges. We have all seen in the media an abundance of anecdotes from health care, automobile, and financial institutions revealing some details of information theft by insiders. Although these anecdotes are interesting, they do not help us understand how the targeted organizations could have spotted the unlawful insiders before they were able to do harm.

One of the real challenges in developing technology to help us tackle this cyber challenge is that malicious insiders usually do not need to engage in rule breaking behavior. They can use their legitimate access to gather and steal sensitive information. Their actions remain largely unseen using traditional cyber-detection methods such as log auditing and intrusion detection, which largely focus on detecting attempted or actual rule-breaking behavior. We therefore aimed to study how malicious insiders operate within their privileges to misuse information. The second real challenge is getting access to field data for testing and evaluation of new technologies or methods. Understandably, organizations that have been affected by malicious insider actions are hesitant to share the details of the violations for fear that their security processes could be further manipulated. And although there are a fair amount of post-mortem cases studies of insider attacks, few of them contain enough detail regarding actual computer usage for further analysis. Since getting data has been difficult, we sought to design an experimental framework showing researchers how to generate and analyze their own data.

The purpose of this research was to identify differences between malicious and benign insiders to help organizations spot suspicious behavior. In pursuing that objective we used rigorous experimental methods to generate good malicious insider data and put together practical guidance for the early detection of malicious activity. To accomplish the first goal, it was necessary to combine the skill sets of computer scientists and social scientists onto one multi-disciplinary team. This was the only way we could successfully integrate understanding of both human behavior and information-use patterns from a cyber perspective.

To accomplish our research objectives, we designed and executed an experiment using our organization's employees. These participants used a monitored laptop to complete a scenario that varied their intent for searching the organization's Intranet and the Internet. We then analyzed their behaviors to determine if users with malicious intent showed patterns of behavior that differed from the possible behavioral patterns of users with benign intent. The experimental methods, study design, data analysis, and lessons learned are summarized below.

Experimental Methods

The social sciences have developed rigorous methods for studying human behavior that added significant value to this research. Foremost, great importance was placed on including a baseline data set with which to compare our malicious activity, more formally called a "control group." The value of a controlled baseline of participants is that it allows us to directly compare their behavior to the malicious behavior to determine if only the variable of interest affected how they completed the information-use task. In order to design a proper control group, we had to recruit participants under the cover story that we were testing anti-keystroke logging software. This specific cover story was created so that legitimate consent to be monitored was provided by participants, even if they did not understand why they were being monitoring. The deception is necessary because, although it may be easy to tell participants to behave maliciously, it is impossible to overtly tell control participants not to behave maliciously. We added further rigor to the design by following double-blind procedures which reduce experimenter and participant bias. Double-blind means that neither the experimenter nor the participants knew the hypotheses of the study, ensuring the proctor does not leak the goals of the study and preventing participants from altering their behavior to aid the experimenter. The cover story controlled for participant bias while the hiring of a proctor without knowledge of the project goals controlled for experimenter bias.

In order to get as representative a sample of the organization as possible, we used a stratified sampling method. With this method, the organization population is divided into subgroups and we randomly sampled the correct proportion from each subgroup to get a good distribution. For this study, we stratified our sample by cumulative job experience so that we got participants including administrative, technical, and managerial staff. After participants completed the study task, we used what is referred to in the social sciences as "funnel debriefing." The proctor was trained to carefully gather specific information from the participants without triggering suspicion. We were interested in determining if any of the participants had discovered the study's hypotheses, as well as gauging their level of identification with their assigned role and task. The funnel debriefing aided us in eliminating a few participants who did not properly complete the study and a few other participants that had deduced the hypotheses of the research.

Study Design

We successfully recruited 61 participants through advertisements and word of mouth. Participants were randomly assigned to one of two conditions: malicious user (experimental group) or benign user (control group). Each condition was given a scenario that described the same task to be completed but provided a role that manipulated their "intent." In the scenario all participants had fallen on the same hard financial times but the malicious role required participants to accept a new, higher paying job that required them bringing inside information about a sensitive contract from their old company. The benign role required participants to research a specific topic for a high profile tiger team where exceptional performance would get

them a promotion and pay increase. Thus, the scenarios were completely balanced except for their intentions while completing the information searching task.

All participants completed consent forms and a pre-study questionnaire and then received a study laptop and their respective scenario. They had 10 hours over 7 days to play their role and complete the task. The study laptop was configured to monitor both their network and host-based behavior at all times. When participants were finished they returned the laptop and a CD containing their research findings to the proctor. Lastly, all participants completed a post-study questionnaire and were debriefed.

High Level Findings

This research was focused on producing guidance for security practitioners. In assessing the study design we found that most participants were fully engaged with their assigned roles and completed the task to the best of their abilities within the time allotted. In monitoring their computer usage behavior, we were interested in both network and endpoint cyber events. The data analysis goals included developing indicators or groups of indicators that highlight suspicious users and reporting interesting data which might generate ideas for future human behavior and cyber security studies.

The data analyses produced a number of interesting findings. Malicious users demonstrated a consistent "quantity over quality" approach to their information search task. They tended to perform more search queries and download as many documents about the topic as they could find. After acquiring the information, malicious participants were more likely to keep entire documents and deliver most of what they found on CD. On the other hand, benign users were significantly more likely to synthesize the information they found. These participants gathered less information and spent time organizing that information in ways that made sense to them. The average total information delivery in megabytes was less for benign users (M=69) then malicious users (M=375). Both the malicious "grab and go" and benign "synthesize and organize" behaviors could be identified through the monitored computer use events.

In addition to these general approaches, the malicious users were more likely to break up their task completion into more information gathering sessions (i.e., more logons and logoffs) and access information indirectly (i.e., use of search engine cache queries), indications that they were trying to avoid the monitoring software by using evasive tactics. Malicious users also avoided directly searching for detailed information by going through project sites and relevant people's shared information folders (i.e., people lookups, and project lookups). Benign users did not show any significant signs of evading the monitoring system.

Lessons Learned.

We learned many lessons about experimental design and user monitoring during the course of this study. First, it takes a significant amount of effort and repetition to make participants understand the importance of their benign or malicious role. This is not unexpected but it highlights the importance of scripting out the entire experimental process and piloting the design with a few people (i.e., dry runs) before beginning formal data collection. In it is also important to carefully consider the analysis restrictions that come from certain design decisions. For example, by providing a study-specific laptop we were unable to capture the behaviors of individuals interleaving their work projects and personal activities. This was not a mistake in the design but a necessary component in order to conduct the experiment since participants would not permit us to put monitoring software on their individual laptops.

As for monitoring, a key lesson was discovering how difficult it can be to separate machine behavior from user behavior. There are many machine-generated events (e.g. auto-save

and browser cookie deletion) that had to be filtered manually before we could perform analysis of the user behavior. To reduce this noise, we recommend focusing monitoring tools on the application layer as much as possible, possibly even at the user interface level. In addition, we discovered a number of monitoring gaps that prevented us from looking at application-specific behaviors such as bookmarking, searches within a web page, and use of tabbed browsing. Researchers should ensure that the software they use is able to differentiate the behaviors of interest. We also suggest they report any monitoring gaps to the software vendors so they can fill in the gaps with future releases.

Conclusions

To our knowledge, this study is the first of its kind looking at the insider threat problem and may also be a model for other human factors and cyber security research questions. We have found the multi-disciplinary team approach to be very rewarding, combining the best practices of social and computer sciences to study the insider threat problem. This study demonstrates the significant value in using a controlled baseline for making direct comparisons between user groups. We understand that experimental designs of this type do have limitations of generalizability and interpretation of findings but much can be learned about this problem set by carefully analyzing specific subsets of it. The benefits of studying one aspect of a problem at a time include better understanding the cause and effect relationship between variables, which is lost in a study that tries to measure everything. We believe that other researchers can and should use this experimental model to study additional aspects of the insider threat problem. There is also significant value in study replication and we encourage researchers to do so whenever possible. Finally, we believe that our finding that malicious users take a "grab and go" as opposed to an organized and methodical approach to information gathering is valuable information which can be used by today's practitioners to focus their monitoring efforts.

Related Publications

Caputo, D.D., Maloof, M.A., & Stephens, G.D (in press). Detecting the Theft of Trade Secrets by Insiders: A Summary of MITRE Insider Threat Research, *IEEE Security & Privacy, Fall 2009.*

Caputo, D.D, Stephens, G.D., Stephenson, B., Cormier, M. & Kim, M (2008). An Empirical Approach to Identify Information Misuse by Insiders," *Recent Advances in Intrusion Detection, Lecture Notes in Computer Science,* Volume 5230, Springer, pp. 402–403.

Maloof, M.A., & Stephens, G.D. (2007). ELICIT: A System for Detecting Insiders Who Violate Need-to-Know," *Recent Advances in Intrusion Detection, Lecture Notes in Computer Science,* Volume 4637, Springer, pp. 146–166.

[i] This material is based upon work supported by the U.S. Department of Homeland Security under Grant Award Number 2006-CS-001-000001, under the auspices of the Institute for Information Infrastructure Protection (I3P) research program. The I3P is managed by Dartmouth College. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Department of Homeland Security, the I3P, or Dartmouth College.

[ii] Schramm, J. H., "FBI's Focus on Economic Espionage", 2006 American Society of Criminology (ASC) Annual Meeting. Los Angeles, CA..