The Human Element: Investigating the Insider Threat

Abner Mendoza and Lei Chen
Department of Computer Science, Sam Houston State University, Huntsville, Texas, U.S.A.

Abstract - *Statistics show that most crimes are committed by individuals known to the victims. In the computing industry, we can draw a parallel to that and show that most computer crimes involve trusted insiders. There have been numerous models proposed for analyzing and mitigating the insider threat problem. However, there has been relatively little effort to collect viable scientific data to support detection and prevention models that can be uniformly applied across different organization. The dynamic nature of the problem and the fact that it involves a human element makes it an immensely difficult problem to solve. One of the immediate problems with researching this issue is that there is no single concise definition of what exactly constitutes the insider problem. In this paper, we investigate the prior work in this field and discuss some additional thoughts that are relative to current research in terms of defining the insider threat.*

Keywords: insider threat, organizational security, human element

1 Introduction

The insider threat has proven to be one of the more complex issues in information security. The complexity starts with the industry's inability to define and agree on a clear and consistent definition of who is considered an insider, and what exactly constitutes an insider threat. The lack of a standard and precise definition hinders research in this field, and makes it difficult for organizations to implement solutions that were devised using a definition that differs from the definition that the organization itself uses. Additionally, most research will stop short of using extensive scientific data to prove or disprove the efficacy of any one approach. There needs to be more extensive and long-term research done, in order to gather scientific data for proper analysis of proposed techniques. One of the problems, as pointed out in [1], is that such data is hard to come by without full cooperation of organizations that are vulnerable to the insider threat. While most organizations acknowledge the risks they face from insider attacks, most are also hesitant to disclose data on any such risks or incidents that may have occurred. It is difficult to assess the nature of insider threats without ongoing large-scale real-world case studies. An additional hindrance to research is that the definition of an insider differs across organizations, as does the risk posed by such people or entities that are deemed to be insiders. One of the underlying issues is that what is considered an insider is relative across different organization or industries. As such, there is no one ideal and consistent definition that is both concise and universally applicable. In general terms, most research will agree that an insider inherently has privileged access. For scientific research, we require a more precise definition that also encompasses the entire scope of the problem so that conceptual models can be uniformly applied across different organizations. That is an ideal situation, which most likely is not attainable if we can make a sufficient argument that the definition is relative.

Another important point to note is that most research efforts focus specifically on the technical aspects of the insider threat, as well as just malicious insiders. In defining the insider threat, more consideration must be given to non-malicious insiders, as well as the non-technical aspects of the problem. Non-malicious insiders would be those individuals who inadvertently disclose information. An example would be an individual whose company laptop containing proprietary information is stolen. While there is clearly no malicious intent in such an incident, as well as it not being a purely technical issue, the risk is as severe as a disgruntled former employee who steals the same laptop. Credible data on the occurrence of insider attacks, if it was readily available, would likely show a high percentage of insider attacks being non-technical in nature, such as the stolen laptop example.

The insider threat vulnerability has been known for decades, and yet there still exists a huge disparity in the research efforts in this field. The lack of a universally accepted definition is one of the root causes of this problem. Nonetheless, there have also been huge strides made in developing solutions that aim to prevent, and mitigate the issue, such as intrusion and extrusion detection systems. The ultimate goal of research in this area is aimed at prevention, detection, and response. Both technical and non-technical measures must be employed to effectively combat the insider threat. Prior research has developed many conceptual models thus far, but few practical models that are proven to be effective. The efficacy of conceptual models has been difficult to measure due largely to lack of data and lack of cooperation from the business world.

2 Related Work

2.1 Defining the Insider Threat

The insider threat has existed since the dawn of the computing age, and will be with us for the foreseeable future. However, it has only been in the last two decades that the seriousness of the issue has been recognized in relation to the ever growing computing infrastructure. As technology continues to advance at a rapid pace, so does our dependency on the technology. Because our information infrastructure has become such a critical and valuable asset, people have inevitably determined that the systems could be compromised for some sort of vested interest. Insiders have the broadest reach into an organization's system, and so they naturally possess a very real and elevated risk of violating security policies as they relate to information systems. In 1999, the RAND Corporation initiated a research effort to develop an agenda to enact policy changes with regard to the insider threat problem. This also coincided with a similar effort by the Department of Defense to direct research efforts with regards to the insider threat problem. In its research paper, the RAND Corporation cites work by CERT together with the US Secret Service to investigate the motivations of convicted insiders. The general principle was that to attempt to predict and prevent the problem, a thorough understanding of the human element of such crimes, as well as what motivated the crime in the first place, must be attained. The far-reaching spectrum of the insider threat problem presents a huge challenge in doing comprehensive research on every aspect of the problem. Additionally, the data collected by these research initiatives cast some level of doubt on the significance of the problem. Although it can be surmised that the problem is very real and very significant, the supporting data is lacking in many aspects. In the RAND research, there were doubts cast about the credibility of data supporting the significance of the threat. Such doubts were fueled by results of a survey that showed different results in two successive years, and suggesting that the data may have been skewed by influence from vendors pushing solutions that are aimed at the insider threat, as well as by the unwillingness of organization to reveal details of breaches from insiders. The goal of the approach explored by the RAND research was to examine the nature and magnitude of threats, and to use technology to suggest appropriate responses. The paper introduces a taxonomy framework to distinguish different aspect the insider threat problem, and then describes articles from three research groups that explore methods of responding to threats.

The first major hurdle in researching the insider threat problem is that there is no clear and consistent definition of the problem. In a 2008 paper, Matt Bishop and Carrie Gates, set out to do just that. Their paper aims to define exactly what an insider is in a concise and consistent manner that is universally applicable. The argument is made that the lack of a consistent definition makes it difficult to perform appropriate research on this topic. Further, it is argued that the lack of a consistent definition results in varying definitions devised by researchers to conform to their research data sets and goals, rather than the other way around where the research is focused on an unbiased definition of the problem. As a result, it is often the case that the results of different research on this issue are not easily applied across different domains simply because of the inconsistency in defining the issue. The definition presented by Bishop considers an insider with reference to a set of security policies, and violation of either those policies or access controls that enforce those policies. The definition proposed presents a non-binary approach where recognition is given to different degrees of "insiderness" with respect to the access they have to certain resources. The authors contend that most research is done on the premise of a specialized definition, rather than a more broad and far-reaching definition of the problem. For example, in [3], the authors give a definition of the insider as individuals who currently or at one time had trusted privileges within a secured system. This definition is more or less binary, which is contrary to what is proposed in [5]. The threat in [3] is further described as the harmful consequences of activities carried out by insiders, whether by malicious intent or disregard for policies. In other words, attackers are either insiders or non-insiders. In [4], the insider is defined as an individual with legitimate access and knowledge of an information system, a broad but vague definition at best. In [9], the implied definition is of a person who intentionally violates security policies of computer systems within their organization. In [8] the authors contend that while the definition of an insider is not agreed upon, there is more conformity as to what an "insider threat" is. This paper subsequently defines the insider threat as encompassing both malicious and non-malicious, and altogether undesirable risks from people with privileged access. So there is a subtle but important distinction regarding the definition of an "insider" as opposed to the "insider threat".

[3] Proposes that there is no clear line that should be drawn to determine insiders versus outsiders. In [7], Matt Bishop further expounds on the definition of the insider threat. In this paper, the authors present the idea of an Attribute Based Group Access Control (ABGAC) mechanism, which sounds conceptually similar to the Group Based Attributed Access Control presented in [5], except with a different name. As in [5], the ABGAC is a generalization of Role-Based Access Control (RBAC), and assigns rights based on general attributes rather than defined work roles. In determining the structure of the ABGAC, and namely to define the different attributes, and groups, the authors present the concept of different layers of a unifying policy hierarchy. This hierarchy allows analysis of a security policy based on what is ideal, feasible, configured, and instantiated. The argument is that while an all-knowing "ideal oracle" can understand abstract concepts such as the intent of a user, and control access based on such information, a real configured and instantiated system has technical limitations that prevent such in-depth decision making when granting or denying

access. The authors content that the discrepancies between these layers are areas that give rise to vulnerabilities in processes and protocols that insiders can exploit. In identifying access attributes, and grouping protection domains, this model aims to give an assessment of the degrees of risk associated with each group of insiders. The basic premise is that not all insiders are alike, and based on these groupings, and organization can analyze and prioritize the insider threat so that implemented controls can focus on those insider groups that pose the most risk. Additionally, it comes down to an analysis of where the most risk is based on the associated costs not to mitigate those risks.

2.2 Detecting and Preventing the Insider Threat

Detection has been one of the major issues of the insider threat problem. If the threat is not detectable, then there is effectively no means of preventing it. In a 2011 paper [2], Jung-Ho et al, describes a framework for a defense system to prevent the insider threat. Most preventative strategies suggested by research are based on methods of detecting the malicious behavior of insiders, whether intentional or unintentional. In [2] the authors explore a framework for a defense system consisting of three modules whose main goals are prevention by the use of monitoring techniques. The authors propose the use of an attack tree model whereby they can identify the paths of potential risks to monitor, and make comparisons with normal activity to discern abnormal activity that may be indicative of misuse by an insider. The authors acknowledge the difficulty in defining what constitutes an insider, and they approach the insider as having malicious intent, thus the title of the paper. No consideration is given to the threats posed by insiders without malicious intent, however. The methods described first establish a baseline of what could be considered normal behavior within the confines of the security policy. Usage activity is monitored at different layers of the framework, and deviations from normal activity are then flagged as potential malicious threats. The first module in the framework is the information checker which grants connection based on the user profile as well as logs of previous or current activity which may have been flagged as malicious. The second module is the behavior checker which uses the attack-tree to monitor behavior and flag any potential malicious activities which may pose risks or deviate from normal activity. The third module is the resource misuse monitor which monitors the pattern of usage for different processes that the user may normally execute, and makes a decision on whether the pattern is malicious if it deviates from the norm. This framework basically comes down to behavioral analysis. The problem, of course, is that behavioral analysis in itself is not enough to effectively detect and prevent insider attacks.

In [8], the author again takes a behavioral approach to detecting insider attacks. The author starts off by rightfully acknowledging the issues abound in this field, as it relates to our understanding of the issues and subsequently our ability to handle the issues. The author also proposes user profiling as a means of predicting insider threat, while also pointing out the potential legal implications that such practices could cause. As for detecting insider threats, he describes methods that use behavioral analysis to alert on insider attacks based on deviation from normal activity of a user. Finally, the author presents several suggestions on what steps should be considered when responding to actual insider attacks.

In [4] the author presents a similarly layered approach for prevention, detection, and recovery from insider attacks. The goal is to achieve comprehensive security by adopting a holistic approach to the problem where all aspects of a system are analyzed, rather than just the technical aspects. An attack classification scheme is developed by extending an existing taxonomy that largely focuses on computer attacks. The author further describes the idea of the Attack Surface, which describes the various means by which someone could access the system, and the Impact Zone, which describes boundaries and constraints for levels of access. These ideas are put together to formulate a strategy of defense against insider attacks. The resulting approach proposed is a three-layer approach consisting of Social, Logical, and Physical layers. By analyzing the progression of attacks through these layers, defenses can be employed to mitigate the problems of insider threats.

CERT published a technical note in October 2011 which presents an analysis of findings from studies on 86 particular insider threat cases that found a pattern of insiders stealing information within 30 days before being terminated. At the time of the writing, the CERT Insider Threat Center database contained more than 600 documented cases of insider attacks from which researchers showed that fraud was the motive behind most attacks. This paper was specifically focused on theft of intellectual property and in determining a pattern to develop a rule that could be used in a log indexing application that is used to detect insider attacks. The goal was to determine a signature for such insider attacks that could be used to detect, and possibly prevent future attacks from being completed. This is similar to the concept in malware detection where known bit patterns, or signatures, are used to detect malware. An example of such an analysis is described in the publication using the Splunk log indexing application.

As opposed to other research efforts cited in this paper, the CERT research uses real world data to assert their findings and to develop a robust rule that is shown to work. This is unlike the theoretical detection frameworks explored thus far from prior research. In 2009 CERT also published what was the third edition of a comprehensive guide for security professionals with practical guidance on measures that can be implemented to prevent and detect insider threats [12]. The guidance draws on insight from research on real world data gathered in the CERT's insider threat incident database, and also builds on guidance from previous releases

of the same publication. The guidance presented is a result of research on these documented incidents, and shows that most incidents are a result of insiders seeking financial gain, business advantage, or simply sabotage. Such examples show that there continues to be progress in this field as more real world data is collected and analyzed in an effort to detect and ultimately prevent the insider threat.

In [6] we also see an example of real world data being used to develop detection mechanisms that employ simple application of best practices as a first step in preventing security vulnerabilities from being exploited by an insider. The article describes real-world cases of insiders subverting the security policies at their place of employment and causing damaging consequences. The authors propose that insiders can be stopped with stringent layered security composed of policies, procedures, and technical controls. The latter part of the article presents a list of 13 best practices that can be implemented in an effort to mitigate threats from insiders. This is similar to the guidance given in [12]. While there are obvious limitations to the list of security practices, it is certainly not intended to be an exhaustive and comprehensive list. Instead, it identifies the problem and presents solutions to start analyzing and dealing with the problem.

Of course, part of the problem in collecting real world data is that organizations are often hesitant to disclose such data or allow collection of such data in fear of backlash from customers or shareholders if the data collection reveals previously unknown exploits by insiders. In [6] extent of the insider threat problem is even put into question. This is partly because of conflicting statistics and reports that sometimes even show a decline in insider threat. Part of this doubt is that statistics are sometimes made to fit into a research or the agenda of a security vendor, so they may massage the data in their favor and perhaps curtail the true extent of the problem.

One aspect of the problem that is not given enough attention is the need for more effective and proper response to insider security violations. After detecting and hopefully preventing an attack, there needs to be responsive measures taken to ensure that the threat is effectively mitigated. In [8] the author makes a point to describe the importance of responding effectively to an insider threat based on the nature of the threat. Most models presented by research focus on prevention and detection. Accepting that this problem will most likely never go away, it is imperative that an organization is prepared to respond to an attack so that it is effectively contained and eliminated if possible, while considering all aspects of the attack including technical and non-technical aspects. One problem that could arise with the behavioral analysis detection method is that it can produce many false positives if not properly implemented.

3 The Human Element

It must be acknowledged that there is much to be desired in research into the insider threat problem. The difficulty of the problem is such that few advances have been made that point to a comprehensive solution. From our investigation into the research on this problem, an abundance of technical solutions have been found, but little in the way of solutions with regard to user education and other approaches that appeal to the human element of the problem. For example, one long term study could draw much conclusion from looking at insider issues between one company that pushes user education in this regard, and another company that does not put any priorities on user education. In [9] the authors specifically address the issue of insiders intentionally violating security policies of computer systems within their organization. The paper presents two case studies in their analysis of computer crimes committed by insiders. In one case study, an employee manipulated the accounting systems in order to produce false profits that would allow him to be awarded huge bonuses. In the other case, an employee was able to manipulate the accounting system as well, but his motives were quite different and he did not personally gain from those actions. Both cases are presented to illustrate the different motives for insider attacks, but the root of the problem in most cases boils down to elevated and unmonitored access by trusted individuals. In analyzing the technical controls, the authors also prescribe that informal measures such as awareness training and education are essential in mitigating these issues.

In [10], the authors specifically explore the psychological aspects of the insider threat. The paper presents ideas on understanding the psychology of malicious insiders, and contends that such understanding is crucial in understanding how to mitigate the risks from insider attacks. They present summaries from several case studies where trusted insiders have violated security policies. This is done to demonstrate the critical nature of the problem, as well as to analyze the different motives behind insider attacks. The authors suggest that many of these incidents boil down to human problems that cannot simply be solved with technological solutions. The research described in this paper is an effort to formalize the psychological profiles of insiders who violate security policies, in an effort to understand the personalities, motives, and circumstances that give rise to such actions. The goal is to establish user profiles that can be applied to business practices such as job application screenings, and management of prospective or current employees. This paper explores an important non-technical area which has not been given enough attention. While it may not be enough to predict future insider attacks based solely on a psychological profile, it is a tool that can be considered in the larger scheme of things, and in addition to other technical and non-technical methods of measuring threats. One of the drawbacks, as pointed out in [8] may be the legal ramifications of actually using such methods. While it may be

more acceptable in governmental/military setting, as was the focus of the research described, the dynamics of such methods may prove much different in corporate settings.

4 Conclusion

We have explored the general issues and proposed solution to the insider threat problem. The most striking issue within the problem is that there is no clear definition of precisely what the insider threat problem really is and who is or is not considered an insider. The insider threat problem has been listed on the national Hard Problem List for many years as an issue that ought to take priority in current research. Despite this, there has been relatively little progress in devising universally applicable models for effectively and efficiently mitigating this risk. Reasons for this stem from the complexity of the problem as well as the lack of available scientific data. A universally applicable definition of what is encompassed in the insider threat and what constitutes an insider has been a major hurdle in solving the issue and making significant research progress. We have seen that CERT, in particular, has continued to collect and analyze data and the fruits of their labor is shown in the guidance and research publications they have produced on the topic. Still, there exist many areas of void as it relates to the research into the insider threat problem. Future work must also place emphasis on the non-technical aspects of the problem such as user education and public relations responsive measures. Additionally, mitigating measures must be properly implemented to optimally reduce, if not totally avoid, the risk from the insider threat.

5 References

[1] Pfleeger, Shari Lawrence; Stolfo, Salvatore J., "Addressing the Insider Threat," Security & Privacy, IEEE , vol.7, no.6, pp.10-13, Nov.-Dec. 2009

[2] Jung-Ho Eom; Min-Woo Park; Seon-Ho Park; Tai-Myoung Chung; , "A framework of defense system for prevention of insider's malicious behaviors," Advanced Communication Technology (ICACT), 2011 13th International Conference on , vol., no., pp.982-987, 13-16 Feb. 2011

[3] Greitzer, F.L.; Moore, A.P.; Cappelli, D.M.; Andrews, D.H.; Carroll, L.A.; Hull, T.D.; , "Combating the Insider Cyber Threat," Security & Privacy, IEEE , vol.6, no.1, pp.61-64, Jan.-Feb. 2008

[4] Clive Blackwell, "A security architecture to protect against the insider threat from damage, fraud and theft.", Proceedings of the 5th Annual Workshop on Cyber Security and Information Intelligence Research: Cyber Security and Information Intelligence Challenges and Strategies (CSIIRW '09), ACM, New York, NY, USA, , Article 45 , (2009)

[5] Matt Bishop; Carrie Gates. "Defining the insider threat." In Proceedings of the 4th annual workshop on Cyber security and information intelligence research: developing strategies to meet the cyber security and information intelligence challenges ahead (CSIIRW '08), (2008), ACM, New York, NY, USA, , Article 15, (2008)

[6] Dawn Cappelli (CERT), et. al., "Protecting Against Insider Threat" (2007), (Online). Available: http://www.sei.cmu.edu/library/abstracts/news-at-sei/securitymatters200702.cfm

[7] Matt Bishop, Sophie Engle, Sean Peisert, Sean Whalen, and Carrie Gates. "We have met the enemy and he is us." In Proceedings of the 2008 workshop on New security paradigms (NSPW '08). (2008), ACM, New York, NY, USA, 1-12.

[8] E.E. Schultz, "Predicting, Detecting, and Responding to Insider Attacks", In ISSA Journal - December 2008, (2008), (Online). Available: http://www.issa.org/Library/Journals/2008/December/Schultz-Predicting-Detecting-Responding%20to%20Insider%20Attacks.pdf

[9] Gurpreet Dhillon, Steve Moores, Computer crimes: theorizing about the enemy within, Computers & Security, Volume 20, Issue 8, 1 December 2001, Pages 715-723

[10] Eric Shaw, et. al., "The Insider Threat to Information Systems", In Security Awareness Bulletin, No. 2/98 (Online). Available: http://www.pol-psych.com/sab.pdf

[11] Michael Hanley, Joji Montelibano, "Insider Threat Control: Using Centralized Logging to Detect Data Exfiltration Near Insider Termination", CERT Program Technical Note, October 2011. (Online) Available: http://www.cert.org/archive/pdf/11tn024.pdf

[12] Dawn Cappelli, et. al., "Common Sense Guide to Prevention and Detection of Insider Threats 3rd Edition – Version 3.1", CERT, January 2009. (Online) Available: http://www.cert.org/archive/pdf/CSG-V3.pdf