A Multidiscipline Approach to Mitigating the Insider Threat

Jonathan W. Butts
Robert F. Mills
Gilbert L. Peterson
Air Force Institute of Technology*, Wright-Patterson AFB, OH, USA
jonathan.butts@afit.edu
robert.mills@afit.edu
gilbert.peterson@afit.edu

Abstract: Preventing and detecting the malicious insider is an inherently difficult problem that expands across many areas of expertise such as social, behavioral and technical disciplines. Unfortunately, current methodologies to combat the insider threat have had limited success primarily because techniques have focused on these areas in isolation. The technology community is searching for technical solutions such as anomaly detection systems, data mining and honeypots. The law enforcement and counterintelligence communities, however, have tended to focus on human behavioral characteristics to identify suspicious activities. These independent methods have limited effectiveness because of the unique dynamics associated with the insider threat. The solution requires a multidisciplinary approach with a clearly defined methodology that attacks the problem in an organized and consistent manner. The purpose of this paper is to present a framework that provides a systematic way to identify the malicious insider and describe a methodology to counter the threat. Our model, the Multidiscipline Approach to Mitigating the Insider Threat (MAMIT), introduces a novel process for addressing this challenge. MAMIT focuses on the collaboration of information from the relative disciplines and uses indicators to produce a consolidated matrix demonstrating the likelihood of an individual being a malicious insider. The well-known espionage case study involving Robert Hanssen is used to illustrate the effectiveness of the framework.

Keywords: Insider Threat, Counter Espionage, Information Security, Behavior Analysis

1. Introduction

The malicious insider is any authorized user that utilizes inherent trusts to intentionally harm or compromise information within a system. The majority of security professionals agree this threat poses the greatest risk to an organization and is the most difficult to detect (Schneier 2000). The threat is receiving significant attention in the business sector because of major losses in revenue and for governments because of the possible compromise of national secrets. According to the SANS Institute, the most serious security breaches resulting in financial losses occurred through unauthorized access by insiders (Kratt 2004). Their research showed an average cost of $57,000 for an attack originating from outside an organization and an average of $2.7 million dollars for damages from an insider attack. As a whole, even though there have been great strides in protecting systems from the outside threat, only modest work has been formalized for defending against the insider.

In August, 2000 a workshop of leading security professionals met to discuss the insider threat. Their findings determined a specific need for a comprehensive model identifying the level of risk posed by system users (Anderson 2000). To date, however, there has been little advancement in effectively achieving this requirement. Previous work has focused on certain aspects of the problem but has not successfully led to a methodology that produces a meaningful risk level. This article addresses these issues by developing a framework that brings information together to produce one consolidated matrix for assessing the threat level for each user in the system.

2. Attack Cycle

To effectively develop mitigation techniques and prevention methods, it is important to first examine the attributes associated with the insider threat. Two primary elements are required for a person to become a malicious insider: opportunity and motive (Heuer 2001). If either characteristic is missing, then the individual does not pose a serious threat to the organization. For example, if an individual has high motivation but no access to the system, then their current likelihood of being a credible insider threat is low. Additionally, if an executive within the organization has high access but no motivation, then their current likelihood is also low. The motive and opportunity factors together provide an indication to the threat level existing for an individual. Before an insider acts, however, their threat level must be great enough to overcome the natural inhibitions to commit criminal behavior. Some of the reasons that typically prevent people from acting include moral values, loyalty

* The views expressed in this article are those of the authors and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government.

to employer or co-workers, or fear of being caught. Once the threat level increases enough to overcome the inhibitions, an event or occurrence usually takes place that pushes the individual over the edge and leads to the actual betrayal. The activity triggering the attack can be a work related incident, personal crisis, threat of force, or other event in an individual's life. The following four conditions, taken directly from (Heuer 2001), are the characteristics identified as generally required before an individual betrays their organization and commits a malicious act:

- An opportunity to commit the crime
- A motive or need for satisfying themselves through the crime
- An ability to overcome natural inhibitions
- A trigger that sets the betrayal in motion

Once the attack occurs, the individual then evaluates how well the compromise went. This step forms the final phase of the attack cycle. Figure 1, the attack cycle for the insider threat, depicts the anatomy of the attack. Each area is now discussed in more detail.

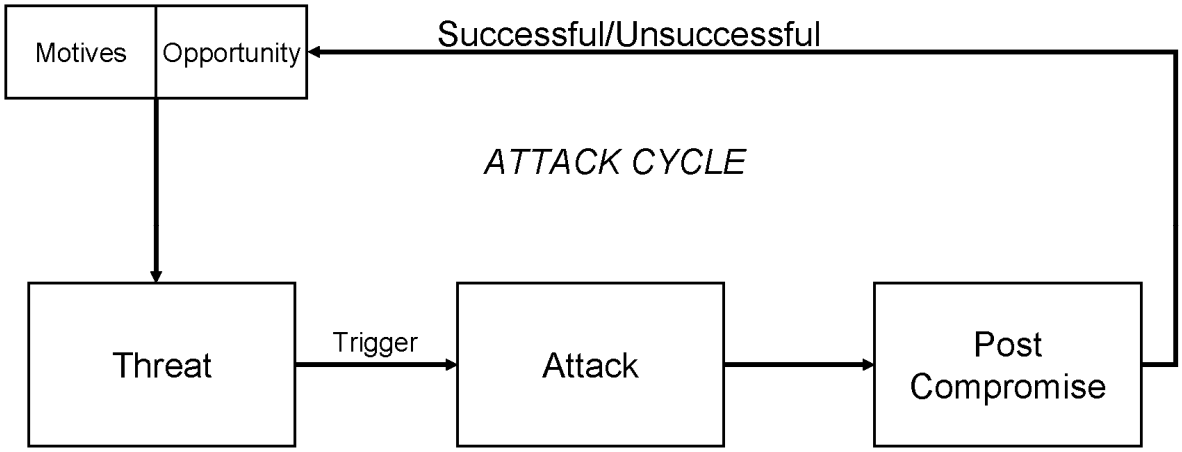

Figure 1: Attack cycle for the insider threat

2.1 Opportunity

The difficulty associated with the malicious insider is the fact that they are the very same person you trust (Schneier 2000). Opportunity for the insider can present itself through granted permissions, compromise of the system, or inadequate enforcement of organizational policies. The amount of privileges assigned or gained within the system directly correlates with the amount of damage possible. For example, the system administrator poses a significant risk to the system because of their level of access and opportunity. In some compromises, a malicious insider attempts to gain access to a more privileged account. In other cases, the person simply uses permissions granted them to carry out the compromise. These types of attacks emphasize the importance of enforcing sound security polices for users as well as maintaining systems (and their security mechanisms) up-to-date and patched.

2.2 Motives

In 2001, the Defense Personnel Security Research Center (PERSEREC) performed a study using open source information on 150 espionage cases to determine trends and patterns associated with the malicious insider (Herbig 2002). Their findings detailed the motivating factors associated with the criminals. The number one motivator was money. Also cited as motives were divided loyalties, disgruntlement with the employer, desire to please someone else, coercion, thrill seeking, and recognition.

Additionally, Dr. Mike Gelles of the Naval Criminal Investigative Service has classified motivators for insiders into two personality disorders commonly found in spies: antisocial personality disorder and narcissism (Gelles 2001). Individuals with antisocial personality disorder lack remorse or guilt when they do something wrong. These individuals reject established rules, are manipulative, self-serving, and seek immediate gratification of their desires. They typically have no interest for the future and are

more concerned with immediate gains. People with narcissism usually suffer from excess self importance or preoccupation and have difficulty living up to their own expectations. These individuals normally feel underappreciated by their supervisors and are unable to accept criticism or failure, because it threatens their inflated self-image. The characteristics for both of these disorders can produce a high threat level for someone to commit a malicious activity and are a serious security concern.

2.3 Threat, Trigger, and Attack

As mentioned previously, the combination of opportunity and motives results in an individual's threat level. Although most individuals within an organization have an opportunity and a financial or personal motive to attack, betrayal is relatively rare because the threat level is not high enough to overcome a person's natural inhibition (Heuer 2001). In those instances where a person performs a malicious act, their individual threat level becomes so elevated that the inhibitions no longer prevent the attack from occurring. Upon reaching this level, typically some event in their personal or professional life finally triggers the act of betrayal. Herbig *et al.* discovered that in one-fourth of the cases reviewed an attacker experiences a life crisis such as divorce, death of a loved one, or failed love affair in the months preceding the attacks (Herbig 2002). A serious financial loss or political event also provides a possible trigger ultimately causing a person to act on their threat level. The type of attacks occurring range from destructive actions to stealing information.

2.4 Post Compromise

After attack completion, a period of post compromise follows. During this time, the malicious insider may attempt to cover their tracks, sell the information, or create back doors in the system for future compromises. The individual also evaluates the success of the attack. A successful attack may lead to increased confidence and the reassurance of the ability to get away with their actions. An unsuccessful attack does not necessarily mean they were discovered, but could simply be the failure to obtain the appropriate information or finish the attack. Either of these results produces a change in the malicious insider's motives and/or opportunity. With success, motivation may increase because of a pay off or self-satisfaction. A failed attempt may also enhance motivation due to a sense of desperation or desire to complete the attack. Opportunity may increase if the system was compromised and privileges were elevated. The effect of the successful or unsuccessful attack on an individual's motives and opportunity effectively changes their threat level, thus creating an attack cycle. The goal of the security community is to detect the malicious insider as early in this cycle as possible.

2.5 Indicators

Throughout the attack cycle, the malicious insider produces characteristics that are capable of being observed. These characteristics in the form of behavior attributes or technical activities are the indicators that can identify a potential insider threat.

Traditionally, law enforcement and counterintelligence communities look for behavior indicators when identifying suspicious activities. These indicators typically relate to an insider's actions or motives, such as sudden increase in spending, suspicious travel plans, withdrawal from co-workers and changes in personal life. In 2004, the Secret Service National Threat Assessment Center and the CERT Coordination Center of Carnegie Mellon University's Software Engineering Institute published a study on insider incidents (Randazzo 2004). Their findings showed that no common profile for the attackers could be established. The attributes of the individuals ranged from 18 to 59 years of age with 42 percent being females. The report also stated the insiders were from a variety of ethnic and racial backgrounds with 54 percent single and 31 percent married. Because of the diversity of the malicious insiders, profiling individuals based solely on attributes provides insufficient indicators about a person's threat level. The focus instead should be on changes in behaviors, which can indicate an individual that poses a threat (Shaw 2002).

On the other hand, the technology community is searching for methods for identifying technical indicators. Technical indicators are independent of characteristics and instead focus on the capabilities, or opportunities, within the system. A few examples include an attempt to compromise an administrator's password, bypassing security mechanisms to access secure documents, or emailing documents to an unauthorized individual. Because of the vast amount of information and data on

systems, tools for detecting attacks focus on the development and use of automated processes. Some current techniques being deployed as countermeasures are Anomaly Detections Systems, data mining event logs, and false production systems used to lure malicious insiders (Honeypots). These techniques attempt to either determine when an attack occurs or make security controls so effective that a compromise is unfeasible.

The law enforcement and counterintelligence agencies as well as the technical community are making great strides in the development of processes for countering the malicious insider. Studies such as (Gelles 2001, Herbig 2002, Shaw 2002) have provided great insight into the behavior characteristics of these individuals and have shaped how law enforcement and counterintelligence agencies gather their information. Additionally, the rapid advancement of technology has led to the growth and expansion of automated mitigation techniques. The problem, however, is there has been minimal success at actually mitigating the threat. Unfortunately, current methodologies to combat the malicious insider are not effective primarily because the strategies have been stove-piped into individual areas of expertise. The law enforcement and counterintelligence communities are effectively gathering information; however, their efforts are in isolation from the technical community and vice versa. These independent methods have failed to be effective because the insider threat is a problem that transcends functional areas. The solution requires a methodology that leverages the indicators from each discipline and collaborates them in a cohesive manner that can be used to identify the malicious insiders.

3. MAMIT

The goal of the Multidisciplinary Approach to Mitigating the Insider Threat (MAMIT) is identification of suspicious individuals within an organization that display a credible amount of threat so follow-up action can be taken. Figure 2 demonstrates the proposed cohesive process for countering the malicious insider.

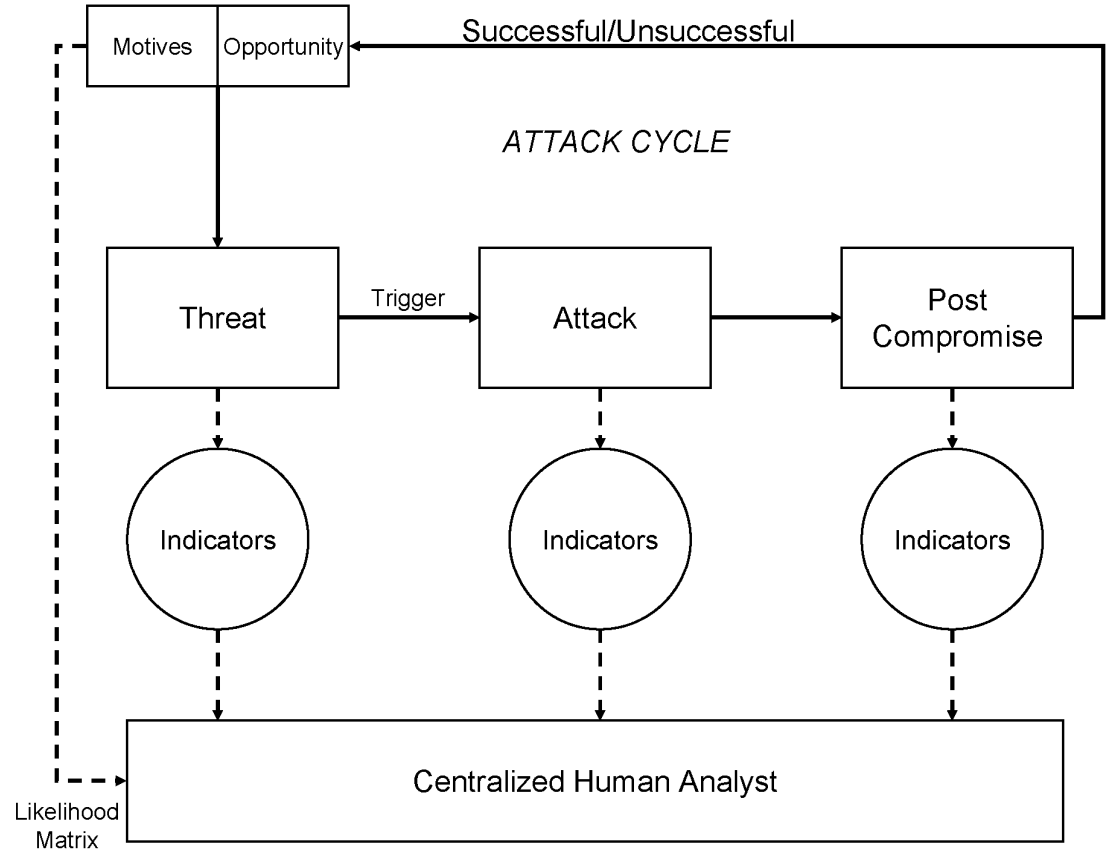

Figure 2: Multidisciplinary process for mitigating the insider threat

3.1 Likelihood Matrix

The MAMIT process requires a methodology for combining the different indicators to form one specific threat level that identifies an individual's risk of being a malicious insider. The Likelihood Matrix performs this function by leveraging the independent findings from each area and merging them together. This process involves developing a method for quantifying the behavior and technical indicators.

3.1.1 Behavior Indicators

Herbig *et al.* performed a study of past American spies and determined that 80% exhibited one or more conditions of security concern defined in the Guidelines for Security Clearance (Herbig 2002). The Guidelines for Security Clearance is a United States directive outlining areas of interest when considering if an individual should be granted a clearance for accessing classified information (U.S. Policy 1997). The adjudicative process examines a number of behavior areas in a person's life to decide if they can be trusted with sensitive information. The authors are currently in the process of examining the 150 case studies in the PERSEREC database (PERSEREC 2004) to determine which of the thirteen specified guidelines are most common to the malicious insiders and provide the best indicators. Early trend analysis showed the following areas were prevalent in a significant number of studies:

- Guideline A: Allegiance to the United States
- Guideline B: Foreign influence
- Guideline E: Personal conduct
- Guideline F: Financial considerations
- Guideline I: Emotional, mental, and personality disorders
- Guideline J: Criminal conduct
- Guideline K: Security violations
- Guideline M: Misuse of information technology systems

These guidelines provide an adequate method for quantifying the motives and behavior characteristics of the insider threat. Guideline A, naturally, is modified to specify allegiance to the organization and Guideline B refers to an outside or competing organization. Each area is independently evaluated and assigned a numeric value based on the determined risk for that area.

3.1.2 Technical Indicators

Technical indicators focus on specifying the opportunity within the system. Quantifying these attributes can be a straight-forward process. The first indicator focuses on characterizing position and system access within the organization. For example, mid-level management may receive a rating somewhere near the middle of the scale, whereas a high ranking official or a system administrator with full access receives a high rating. The second indicator concerns the technical ability of the individual. Someone that is technically savvy with an extensive understanding of the system's inner workings possesses a higher capability than a person with little technical ability. This rating is based solely on a person's skill and knowledge, not their role within the organization. These two ratings, coupled with the behavior indicators, produces ten total variables used in determining an individual's threat level.

3.1.3 Putting it together

The first step in building the Likelihood Matrix is performing an initial analysis on an individual by assigning a numerical value for each of the ten variables. The values are then averaged, producing an individual threat level. This process is accomplished for each person within the organization. When everyone has been assigned their threat level, one overall mean threat level for the organization is calculated by averaging the individual threat levels together. Statistical analysis using a *T* Distribution is then performed to determine a prediction interval for the mean threat level of the organization. If an individual's threat level falls above the interval, then they are identified as a potential insider threat. The *T* Distribution is chosen because it estimates the mean without necessarily knowing the population variance (Devore 2004). The prediction interval is used because the objective is to determine if a single individual's threat level lies outside the statistical norm for the organization's population. The following calculation is used for determining the interval for the mean threat level:

$$\bar{X} \pm t_{\alpha/2} \cdot S\sqrt{1+\frac{1}{n}}$$

The prediction level is 100(1-α)%
$\bar{X}$ is the overall mean threat level of the organization
t represents the *T* Distribution
S is the calculated standard deviation
n is the number of individuals within the organization

When determining the prediction level to choose, recall from statistics that a larger percentage results in a greater range. Setting the value to a smaller percentage can be beneficial for an organization concerned with strict security, such as a highly classified government agency. The trade-off, however, is an increased likelihood of false positives.

3.1.4 Updates

The evaluation of personnel and their threat levels is an ongoing process that requires constant updates. For example, a promotion or modification within the system may change an individual's opportunity. Additionally, a divorce or financial loss may increase the behavior risk. It is important to continually monitor and update these factors to maintain current threat levels. A significant part of this process involves active participation by immediate supervisors. Engagement by supervision is critical because they are typically in the best position to notice changes when they first occur (CSO 2005).

3.2 Centralized Human Analyst

Ultimately, the purpose of MAMIT is to identify suspicious individuals by leveraging inputs gathered by the different professional areas. To effectively implement this scheme, there must be a central analyst to funnel the information to and maintain the Likelihood Matrix. This responsibility falls on the role of the Centralized Human Analyst (CHA).

The CHA is a section within the organization that compiles the intelligence received from each of the different areas and updates the matrix. This concept provides one entity that has full scope of the problem and can maintain two-way communication to each area involved in the process. The CHA should be composed of trusted agent(s) and limited to the number of people that can effectively monitor the organization. Using the Likelihood Matrix, the CHA can identify the individuals that require further observation. Maintaining one central oversight should provide earlier detection and less compromise to information systems. These techniques are demonstrated in the next section using a high profile case study.

4. Case Study

The case involving Robert Hanssen is perhaps one of the most well-known and damaging incidents involving a malicious insider. The details of his case study are used to demonstrate the application and effectiveness of the MAMIT process. The information gathered for this analysis is a collection of documented testimony and reports from (Davey 2002, PERSEREC 2004, Rodriguez 2001).

4.1 The Hanssen Attack Cycle

Robert Hanssen was a 27-year veteran of the FBI who was caught spying for Russia for more than 15 years. During this time he had a high level of opportunity through clearance and access as high as almost anyone else in the government. He participated in operational security and counterintelligence efforts for programs involving some of the most sensitive projects within the intelligence community. Throughout his ordeal he demonstrated motivation factors of money, job dissatisfaction and a feeling of superiority. Hanssen demonstrated the personality traits identified by Dr. Gelles consistent with narcissism. It is possible that these traits stem from claims of negative experiences early in life, especially abuse by his father. These factors resulted in a significant threat level that was ultimately triggered by financial circumstances. The complete success of his attacks for such an extended period of time led to an extensive amount of compromise. Court documents revealed that Hanssen

divulged some of the most highly compartmented information regarding intelligence projects, including US nuclear war defenses. It was determined that he was responsible for providing over 6,000 pages of classified documents and the identities of three Russian agents working for the United States. The Russians paid Hanssen $1.4 million for the information he divulged to them.

4.2 Indicators and Likelihood Matrix

For analysis on Hanssen's threat level, each indicator variable is set to a value representative of his rating for the majority of the compromise period. A range of 1 to 10 is used in increments of 0.5, with the value 10 representing the highest threat. When assigning the values, indicators were used from the case studies that were revealed during the compromise timeframe to co-workers, supervisors, and other individuals in a position to report the incidents to a CHA. If something was revealed after his capture, it was not included in the assessment of his ratings.

Externally, Hanssen appeared to have a high level of allegiance to his organization. He was a loyal member of a conservative Catholic group, Opus Dei, which takes a stance on rejecting communism. Although he appeared loyal to the U.S., there were some allegiance indicators through notes sent to supervisors that he was unhappy with the administration within the FBI and was dissatisfied with his sudden lack of promotions. His foreign influence was reasonably high. His 27-year career in the FBI involved travel to different countries and introduced him to many foreign contacts. He even learned to speak Russian fluently. Hanssen's personal conduct displayed a strong dislike for co-workers and a sense of isolation or not fitting in. His detached personality was a significant factor in the pass over for promotions late in his career. He also displayed strong financial indicators by living more extravagantly than expected of his salary. His expenses included tuition for six children in private schools and colleges, along with a house payment in the New York suburbs. His brother-in-law, also an FBI agent, grew suspicious about his money and reported Hanssen to supervisors. His mental and emotional aspect consistently displayed signs of narcissism through an inflated self-image and lack of empathy for others. Hanssen did not appear to have a criminal record. However, some of his security violations seem to indicate a tendency towards criminal behavior. He had been identified a number of times with highly classified information in his possession that he should not have been viewing and would often boast about insight into stories that were not relevant to his work. There are several instances where Hanssen was counseled for his misuse of information systems including browsing through computer systems and hacking into his boss's computer. As far as technical indicators, Hanssen had an extreme amount of access to essentially all information within the organization and was a recognized expert in computer systems and technology. Table 1 summarizes the values assigned for the indicators based on the relevant information from the case studies.

Table 1: Indicator values for Robert Hanssen

Indicators	Assigned Values
Allegiance to the organization	4.0
Foreign influence	8.5
Personal conduct	8.0
Financial considerations	10.0
Emotional, mental, and personality disorders	9.0
Criminal conduct	4.0
Security Violations	9.0
Misuse of information technology systems	9.5
Access	10.0
Ability	10.0

4.3 Identifying a Spy

Averaging the assigned indicator values, the threat level for Hanssen equates to a value of 8.2. Unfortunately, information on Hanssen's co-workers is unavailable and must be approximated. For demonstration purposes of MAMIT, assume Hanssen's co-workers consisted of 30 other employees and the organization's overall mean threat level of all individuals equated to 4.0 with a variance of 2.5. These assignments for an FBI organization do not appear unreasonable and would indicate a somewhat moderate risk level by the employees. Using the *T* Distribution over a 90% prediction

interval, the range of acceptable threat levels is calculated to fall between 1.3 and 6.70. Because Hanssen's threat level falls above this range, he is identified as a potential threat. Using 4.0 as an acceptable mean, Hanssen would be flagged as a threat for any level of variance up to 6.0. This is a somewhat significant variance and in all likelihood is greater than what would be expected.

When Hanssen was finally detected, he was caught mainly by chance. A Russian double agent provided the U.S. a file that had been transferred by Hanssen to the Russians. Hanssen was ultimately identified by a clear set of fingerprints located on a garbage bag that had contained the file. It is highly probable that if the MAMIT process had been used and information was correctly funneled to a CHA for analysis, Hanssen would have been identified as a possible threat and discovered much earlier.

5. Implementation Scheme

A significant consequence that arises from the MAMIT process is the need to closely exam how typical organizational structures are aligned. Current organization structures typically isolate the different areas of security, with system administrators usually responsible for network security and separate divisions for intelligence and physical security. A study performed by CSO found that 81% of companies separate information and physical security (Fitzgerald 2003). The separation of these areas does not allow for a cohesive flow of information. Additionally performing dual tasks, such as using the system administrator to perform network management and security, does not provide responsible oversight.

This model demonstrates a need to shift away from the current culture and way of doing business. The MAMIT functionality establishes a requirement for one authority, as the role of the CHA, to maintain the full spectrum of security. The CHA is responsible for the security management of the organization, but relies heavily on inputs from the system administrators, direct supervisors and other related functions. These other areas are still a critical part of solving the problem and require individuals in these positions to be security conscious. The main difference, however, is that the accountability and consolidation of information focuses on one area, providing single oversight. Additionally, the CHA should be outside of the normal organization structure and report directly to top-level executives. The benefits of this methodology are that it brings all aspects of security within the organization together and creates a flow of information to one responsible authority for identifying security risks.

6. Discussion and Ongoing Research

The MAMIT methodology introduces a novel approach for addressing the insider threat. MAMIT consists of a framework that leads to effective identification of possible malicious insiders based on their threat level. The strategy focuses on collaboration of information using a multidisciplinary approach producing a single identifier for risk analysis. Because the indicators are provided to one central analyst, individuals with an elevated threat level can be identified earlier and techniques enacted to mitigate the threat. The effectiveness of MAMIT was illustrated through the case study of Robert Hanssen which demonstrated the process would likely have identified him as an insider threat. The methods discussed in this paper produce tangible results useable by organizations to effectively align their security areas for identifying and minimizing security risks.

Ongoing research includes data mining the PERSEREC case studies to determine the importance of each guideline. Based on the historical cases, a weighting scale can be developed to emphasize the significant areas. These results should add to the effectiveness of the model by accentuating the most typical attributes for a malicious insider. Additionally, methods are being investigated to determine an effective and straightforward way to rate each of the ten areas.

References

Anderson, R., Bozek, T., Longstaff, T., Meitzler, W., Skroch, M., and Van Wyk, K. (2000*)*, *Research on Mitigating the Insider Threat to Information Systems - #2 Proceedings of the Insider Workshop*, CF-163-DARPA, Arlington, VA.

CSO Magazine, U.S. Secret Service, and CERT Coordination Center (2004), "2004 E-Crime Watch Survey", *CSO Magazine*, May 2005.

Davey, M. (2002), "Secret Passage", *Chicago Tribune* [online], 21 Apr

http://www.cicentre.com/Documents/DOC_Hanssen_Tribunemag.htm
Devore, J. (2004), *Probability and Statistics for Engineers and Scientists*, 6th edition, California: Brooks/Cole.
Fitzgerald, M. (2003), "All Over the Map", *CSO Magazine*, June 2003.
Gelles, M. (2001), "Exploring the Mind of a Spy", *Defense Personnel Security Research Center*, [online], http://www.dss.mil/search-dir/training/csg/security/Treason/Mind.htm
Herbig, K. and Wiskoff, M. (2002), *Espionage Against the United States by American Citizens 1947-2001*, PERSEREC Technical Report 02-5.
Heuer, R. (2001), "The Insider Espionage Threat", *Defense Personnel Security Research Center*, [online], http://www.dss.mil/search-dir/training/csg/security/Treason/Insider.htm
Kratt, H. (2004), "The Inside Story: A Disgruntled Employee Gets His Revenge", SANS *Institute*, [online], http://www.sans.org/rr/whitepapers/engineering/1548.php
PERSEREC Unclassified Database (2004), "ESPIONAGE CASES 1975-2004", *Defense Personnel Security Research Center*, [online], http://www.dss.mil/training/espionage
Randazzo, M., Keeney, M., Kowalski, E., Cappelli, D., and Moore, A. (2004), *Insider Threat Study: Illicit Cyber Activity in the Banking and Finance Sector*, U.S. Secret Service and CERT Coordination Center/SEI.
Rodriguez, P. (2001), "Diary of a Spy", *INSIGHT Magazine*, [online], 16 Jul, http://www.mdep.org/DiaryOfASpy.html
Schneier, B. (2000), *Secrets and Lies*, Indianapolis: Wiley Publishing.
Shaw, E. (2002), "The Insider Threat to Information Systems", *Security Awareness Bulletin*, No. 2-98, pp 27-46.
U.S. Policy, (1997), *Adjudicative Guidelines for Determining Eligibility for Access to Classified Information*, [online], http://www.usaid.gov/policy/ads/500/adj_guidelines.pdf