

Air Force Institute of Technology

AFIT Scholar

Theses and Dissertations

Student Graduate Works

3-2-2006

Formal Mitigation Strategies for the Insider Threat: A Security Model and Risk Analysis Framework

Jonathan W. Butts

Follow this and additional works at: <https://scholar.afit.edu/etd>



Part of the [Information Security Commons](#), and the [Risk Analysis Commons](#)

Recommended Citation

Butts, Jonathan W., "Formal Mitigation Strategies for the Insider Threat: A Security Model and Risk Analysis Framework" (2006). *Theses and Dissertations*. 3304.

<https://scholar.afit.edu/etd/3304>

This Thesis is brought to you for free and open access by the Student Graduate Works at AFIT Scholar. It has been accepted for inclusion in Theses and Dissertations by an authorized administrator of AFIT Scholar. For more information, please contact richard.mansfield@afit.edu.



FORMAL MITIGATION STRATEGIES FOR THE INSIDER THREAT:
A SECURITY MODEL AND RISK ANALYSIS FRAMEWORK

THESIS

Jonathan W. Butts, 1st Lt, USAF

AFIT/GIA/ENG/06-02

DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

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

AFIT/GIA/ENG/06-02

FORMAL MITIGATION STRATEGIES FOR THE INSIDER THREAT:
A SECURITY MODEL AND RISK ANALYSIS FRAMEWORK

THESIS

Presented to the Faculty
Department of Electrical and Computer Engineering
Graduate School of Engineering and Management
Air Force Institute of Technology
Air University
Air Education and Training Command
In Partial Fulfillment of the Requirements for the
Degree of Master of Science

Jonathan W. Butts, B.S.

1st Lt, USAF

March 2006

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

FORMAL MITIGATION STRATEGIES FOR THE INSIDER THREAT:
A SECURITY MODEL AND RISK ANALYSIS FRAMEWORK

Jonathan W. Butts, B.S.
1st Lt, USAF

Approved:

/signed/

2 Mar 2006

Robert F. Mills, PhD (Chairman)

date

/signed/

2 Mar 2006

Rusty O. Baldwin, PhD (Member)

date

/signed/

2 Mar 2006

Gilbert L. Peterson PhD (Member)

date

Abstract

The advancement of technology and reliance on information systems have fostered an environment of sharing and trust. The rapid growth and dependence on these systems, however, creates an increased risk associated with the insider threat. The insider threat is one of the most challenging problems facing the security of information systems because the insider already has capabilities within the system. Despite research efforts to prevent and detect insiders, organizations remain susceptible to this threat because of inadequate security policies and a willingness of some individuals to betray their organization. To investigate these issues, a formal security model and risk analysis framework are used to systematically analyze this threat and develop effective mitigation strategies.

This research extends the Schematic Protection Model to produce the first comprehensive security model capable of analyzing the safety of a system against the insider threat. The model is used to determine vulnerabilities in security policies and system implementation. Through analysis, mitigation strategies that effectively reduce the threat are identified. Furthermore, an action-based taxonomy that expresses the insider threat through measurable and definable actions is presented.

A risk analysis framework is also developed that identifies individuals within an organization that display characteristics indicative of a malicious insider. The framework uses a multidisciplinary process by combining behavior and technical attributes to produce a single threat level for each individual within the organization. Statistical analysis using the *t*-distribution and prediction interval on the threat levels reveal those individuals that are a potential threat to the organization. The effectiveness of the framework is illustrated using the case study of Robert Hanssen, demonstrating the process would likely have identified him as an insider threat.

Acknowledgements

First, I would like to express my gratitude and appreciation to my faculty advisor, Dr. Robert Mills, for his expertise and guidance throughout this research effort. The dialogue and support was invaluable in helping me reach the end goal. I would also like to thank my committee members, Dr. Rusty Baldwin and Dr. Gilbert Peterson. Their feedback and experience was a vital component to my success. Additionally, Maj (Dr.) Robert Neher's insight and recommendations on statistics was greatly appreciated. A special thanks goes to my fellow classmates for their encouragement and insightful discussions. Finally and most importantly, I thank my family – my wife, son, and daughter. Their unconditional love and support throughout this endeavor helped make one of the more challenging tasks also one of the most rewarding experiences.

Jonathan W. Butts

Table of Contents

	Page
Abstract	iv
Acknowledgements	v
List of Figures	ix
List of Tables	x
I. Introduction	1
1.1 Overview	1
1.2 Background	1
1.3 Research Focus	2
1.3.1 Objectives	2
1.4 Thesis Organization	3
II. Literature Review	4
2.1 Malicious Insider Profile	4
2.1.1 Insider Threat Studies	4
2.1.2 Trends	5
2.2 Case Studies	8
2.3 Security Models	11
2.3.1 Security Models for the Insider Threat	11
2.3.2 Schematic Protection Model	12
2.3.3 Model Specifics	12
2.3.4 Example	16
2.3.5 Extensions and Implementation	16
2.4 Risk Analysis	16
2.4.1 Workshops	17
2.4.2 Frameworks for Identifying the Insider Threat	19
2.5 Summary	22
III. Methodolgy	23
3.1 Problem Definition	23
3.1.1 Goals and Hypothesis	23
3.2 Security Model for the Insider Threat	24
3.2.1 Approach	25
3.2.2 Model Boundaries	25
3.2.3 Features	27

	Page
3.2.4	Specifications 27
3.2.5	Analysis 28
3.3	Risk Analysis for the Insider Threat 28
3.3.1	Approach 29
3.3.2	Specifications 29
3.3.3	Attributes 29
3.3.4	Analysis 30
3.4	Evaluation 30
3.5	Summary 31
IV.	Developing a Security Model for the Insider Threat 32
4.1	Taxonomy Development 32
4.1.1	Approach 32
4.1.2	Methodology 33
4.1.3	Example 36
4.1.4	Decomposition 36
4.1.5	Taxonomy Attributes 38
4.2	SPM-IT 38
4.2.1	Rights 39
4.2.2	Actions 40
4.2.3	Ticket Use and Transfer 44
4.3	Analysis 45
4.3.1	Policy and Implementation 45
4.3.2	Alteration Threat 47
4.3.3	Snooping Threat 48
4.3.4	Distribution Threat 48
4.3.5	Maximal State 50
4.3.6	Mitigating the Threats 51
4.4	Summary 55
V.	Risk Analysis for Detecting Malicious Insiders 57
5.1	Attack Cycle 57
5.1.1	Opportunity 58
5.1.2	Motives 59
5.1.3	Threat, Trigger, and Attack 59
5.1.4	Post Compromise 60
5.1.5	Indicators 60
5.2	MAMIT 61
5.2.1	Likelihood Matrix 61
5.2.2	Centralized Human Analyst 65

	Page
5.3 Case Study	66
5.3.1 The Hanssen Attack Cycle	66
5.3.2 Indicators and Likelihood Matrix	66
5.3.3 Identifying a Spy	68
5.4 MAMIT Implementation Scheme	69
5.5 Summary	70
VI. Conclusions	71
6.1 Problem Summary	71
6.2 Conclusions of Research	71
6.2.1 Security Model	71
6.2.2 Risk Analysis	72
6.3 Significance of Research	72
6.4 Recommendations for Future Research	73
Appendix A. Security Policy with Mitigation Strategies	75
Bibliography	78
Index	1

List of Figures

Figure		Page
2.1.	Attack trends.	7
2.2.	Risk analysis framework presented at C3I/DARPA workshop.	18
2.3.	ARDA taxonomy of cyber events.	18
2.4.	ARDA insider threat risk analysis.	19
2.5.	Schultz insider threat framework	20
2.6.	Magklaras insider threat prediction tool.	21
3.1.	Access Control Matrix.	26
3.2.	Example of insider threat vulnerability	27
4.1.	Operations for an Access Control Matrix.	34
4.2.	The four actions in the functional decomposition tree.	35
4.3.	Decomposed tree representation.	37
4.4.	An example decomposing Distribution: file sharing.	38
4.5.	Distribution.	42
4.6.	Distribution of a copy.	43
4.7.	Distribution through association.	44
4.8.	Initial state.	47
4.9.	Distribution threat.	49
4.10.	Distribution threat using a copy.	49
4.11.	Distribution threat through association.	50
4.12.	Maximal state.	51
4.13.	Two-person integrity mitigation scheme.	53
4.14.	Separation of duty mitigation technique.	54
4.15.	Implementation of mitigation techniques.	55
5.1.	Attack cycle for the insider threat.	58
5.2.	MAMIT framework for mitigating the insider threat.	62

List of Tables

Table		Page
2.1.	Number of successful attacks reported by organizations.	6
2.2.	Revenue losses reported by organizations	8
5.1.	Indicator values for Robert Hanssen.	68

FORMAL MITIGATION STRATEGIES FOR THE INSIDER THREAT: A SECURITY MODEL AND RISK ANALYSIS FRAMEWORK

I. Introduction

1.1 Overview

The Department of Defense (DoD) and the Air Force rely heavily on information systems to accomplish their missions. The rapid growth in technology continues to bolster communication, information sharing, and asset management—to name just a few benefits. It is impossible to imagine any scenario where the DoD or Air Force would wage a war or defend the nation’s interest without leveraging technology and information systems.

Unfortunately, with these benefits also comes an increased risk for attack. The reliance on and use of information systems have made information progressively easier to obtain and exploit. For the DoD and Air Force, this makes the insider threat one of the greatest risks to information systems and the resources they store.

The insider threat is characterized as authorized users performing unauthorized activities. The difficulty associated with the malicious insider is they are the very same person you trust, making them one of the hardest threats to detect. An insider is in position to cause significant damage because the individual already has access to the system and can usually ignore mechanisms designed to prevent an attack. Additionally, the insider typically knows where the target is and can exploit information without drawing attention to himself. This makes the insider threat one of the most challenging problems facing the security to information systems.

1.2 Background

The security of a system is defined by its confidentiality, integrity, and availability requirements. These aspects are protected by an organization’s security policy. A

security policy is a statement that dictates what is allowed and what is not allowed on the system. Once the policy is defined, strategies and techniques are devised to enforce the policy. The benefits of an effective security policy alone, however, are not enough to prevent attacks. Ultimately, security is a people problem and the behavioral characteristics of individuals can thwart even the strictest security policy. To implement a secure system, it is therefore necessary to identify the risks inherent with individuals as well as enforce a sound security policy.

Two important tools in computer and information security are security models and risk analysis. Security models can be used to determine the effectiveness of a policy and determine how to enforce security for a system. Risk analysis is a formal process that evaluates the likelihood of a compromise occurring.

The insider threat continues to be an elusive problem. There is currently no formal security model to analyze the safety of a policy against the insider threat and risk analysis methodologies are not sufficient. If a security model can determine the effectiveness of a security policy and risk analysis can identify potential malicious insiders, then the insider threat to an organization can be reduced. Thus, the need for a formal security model and risk analysis for the insider threat is clear.

1.3 Research Focus

The primary focus of this research is to develop methods to formally analyze the safety and security of a system for the insider threat. The problems are addressed through a formal security model that analyzes security policies and a formal method for identifying individuals that pose a risk. By addressing the insider threat using systematic and formal processes, this study encourages sound policy implementation and risk analysis that can reduce the insider threat to information systems.

1.3.1 Objectives. This research has two objectives. The first is to develop a security model for the insider threat. Security models represent an abstract policy in a systematic manner so analysis can be performed. The security model provides

appropriate assumptions about the system and generalizes what is possible or impossible, given the assumptions. Through logical analysis or mathematical proof the model determines if violations to the policy occur and where mitigation techniques can be applied. A security model provides a process to determine if a given protection schema can formally be proven safe.

The second objective is to develop a formal risk analysis process for the insider threat. When an individual betrays his organization, he produces characteristics that are capable of being observed. These characteristics take the form of behavior attributes or technical activities that can identify a potential insider threat. These indicators identify suspicious individuals that display a credible amount of threat so follow-up action can be taken.

These two objectives identify the scope of this research. Developing a security model and risk analysis technique provides a systematic and formal methodology for addressing the insider threat. The attributes of formal analysis leads to effective prevention and detection techniques.

1.4 Thesis Organization

The remaining document is organized as follows. Chapter II reviews the insider threat by examining trends, case studies, and related work on security models and risk analysis. Chapter III presents the methodology for accomplishing the goals of this research. Chapter IV introduces a security model for the insider threat and demonstrates its effectiveness. Chapter V discusses a risk analysis process for identifying potential malicious insiders. Chapter VI presents conclusions, significance, and areas for future research.

II. Literature Review

This chapter is an overview of the principles associated with the insider threat. Initially, the characteristics and attributes of the malicious insider are examined. Security models are reviewed to determine methodologies for representing the insider threat. A formalization for analyzing the safety of an information system is discussed using the Schematic Protection Model. Finally, risk analysis frameworks for the insider threat are presented.

2.1 *Malicious Insider Profile*

Examining trends in attacks and attributes of prior attackers provide a better understanding of threats and lead to more effective countermeasures. For this reason, analysis of the problem begins by examining the profile and characteristics of the malicious insider.

2.1.1 Insider Threat Studies. In August, 2004 the Secret Service National Threat Assessment Center (NTAC) and the CERT Coordination Center of Carnegie Mellon University's Software Engineering Institute (CERT/CC) published a study of insider incidents involving real-life case studies [29]. The study reviewed 23 incidents involving malicious insiders in the banking and finance sector. One major finding of the research was that most incidents required little technical sophistication. According to their analysis 87% of the malicious insiders performed simple, legitimate user commands to carry out their actions and only a small number of cases required more technical knowledge. This finding indicates it is important for organizations to realize the risks associated with conventional users and not focus solely on the privileged administrators or technically savvy.

Furthermore, malicious insiders typically planned their actions. In fact, 81% of the activities were planned in advanced. This suggests a possibility for prevention or detection of the attackers prior to the incident. Additionally, no common profile for the attackers could be established. The attributes of the individuals ranged from

18 to 59 years of age with 42% being females. Insiders were from a variety of ethnic and racial backgrounds with 31% married. Because of the diversity of the malicious insiders, profiling individuals based solely on personal attributes does not provide significant indicators about a person's threat potential.

In an additional study in May, 2005 the NTAC and CERT/CC viewed cases specifically dealing with computer system sabotage [21]. This study examined insider incidents across critical infrastructure sectors (banking and finance, information and telecommunications, transportation, postal and shipping, emergency services, continuity of government, public health, food, energy, water, chemical industry and hazardous materials, agriculture, and defense industrial base) in which the insider's primary goal was to sabotage some aspect of the organization. This research also found the majority of activities were planned in advance and 61% of the attacks involved simple attack methods using legitimate user commands, information exchange, or physical attack. In the 39% of cases involving one or more relatively sophisticated methods, a script program, autonomous agent, toolkit, or flooding was involved in the attack.

2.1.2 Trends. The annual FBI/CSI survey is currently in its tenth year and has provided valuable insight into threats associated with computer systems. The survey tracks computer security trends through responses from professionals in U.S. corporations, government agencies, financial institutions, medical institutions, and universities [16]. The 2005 report included 700 respondents that had a fairly equal split between the number of successful attacks originating from the inside and successful attacks from the outside as shown in Table 2.1.

The FBI/CSI survey also reports the number of detected insider attacks is on the decline. Figure 2.1 displays the percent of respondents that detected various types of attacks. This finding is in sharp contrast to what is commonly being reported in other literature [6]. According to Bingham, their database indicates incidents are on the rise [2]. The database maintained by Intrusic, called Insider Threat Watch, tracks

Table 2.1: Number of Successful Attacks Experienced by Organizations [16].

How Many incidents, by % of respondents	1-5	6-10	>10	Don't know
2005	43	19	9	28
2004	47	20	12	22
2003	38	20	16	26
2002	42	20	15	23
2001	33	24	11	31
2000	33	23	13	31
1999	34	22	14	29
How Many incidents from the outside by % of respondents	1-5	6-10	>10	Don't know
2005	47	10	8	35
2004	52	9	9	30
2003	46	10	13	31
2002	49	14	9	27
2001	41	14	7	39
2000	39	11	8	42
1999	43	8	9	39
How Many incidents from the inside by % of respondents	1-5	6-10	>10	Don't know
2005	46	7	3	44
2004	52	6	8	34
2003	45	11	12	33
2002	42	13	9	35
2001	40	12	7	41
2000	38	16	9	37
1999	37	16	12	35

malicious insider attacks and third party reports on each compromise. The disparity in views can perhaps be explained by insider incidents not accurately being reported. Reasons for this may include: insufficient evidence or damage to warrant prosecution; negative publicity; or more insiders are remaining undetected [21, 36]. Only 20% of the respondents for the FBI/CSI survey reported a compromise to law enforcement citing fear of negative publicity as the key reason for not reporting. Some researchers warn that survey data on computer crimes can be inaccurate due to the unreported or undetected acts, however, it can still be useful in characterizing a minimal level of threat and in drawing attention to security problems as a whole [36].




Figure 2.1: Trends in Detected Types of Attacks or Misuse [16].

Even with conflicting reports about some aspects of the insider threat, experts agree malicious insider attacks are not only more successful but also more costly than external attacks [28]. According to the SANS Institute, the most serious security breaches resulting in financial losses occur through unauthorized access by insiders [22]. 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. Additionally, the aforementioned study by NTAC and CERT/CC demonstrated the cost to the critical infrastructure in Table 2.2.

There is no doubt the insider threat is a serious concern and can cause significant damage since the individual already has access to the system he wants to compromise and can usually bypass the mechanisms in place that are designed to prevent an attack [33]. Additionally, the insider typically knows where the target is and can complete the objective without drawing attention to themselves. The fundamental

Table 2.2: Percentages of Organizations in the Critical Infrastructure Experiencing Financial Losses [21].

Percentage of Organizations	Financial Loss
42	\$1 – \$20,000
9	\$20,001 – \$50,000
11	\$50,001 – \$100,000
2	\$100,001 – \$200,000
7	\$200,001 – \$300,000
9	\$1,000,001 – \$5,000,000
2	Greater than \$10,000,000

problem in dealing with the insider threat is the malicious individual you are trying to prevent, is the same person you trust.

The different studies and trend analysis provide some important characteristics about the insider threat problem. The main characteristics are identified by the following list:

- A malicious insider may have little technical ability
- The malicious insider's action are typically planned in advance
- Malicious insiders do not share a common profile
- Gathering complete, dependable statistics about types/numbers of insider attacks is inherently difficult
- The malicious insider can cause significant damage
- The malicious insider is a trusted individual

2.2 Case Studies

The previous section demonstrates malicious insiders cause significant damage. This section expands on that by examining some historical malicious insider case studies. The nine specific cases that are summarized had many similarities with dozens of other insider cases that were reviewed from [12,21,29].

Aldrich Ames is perhaps one of the most widely known espionage cases. In 1985 Ames was assigned to a counterintelligence unit and had access to highly classified information [12]. He began selling secrets to the Russians around 1987, but wasn't actually arrested until 1994. During his time as a spy for the Soviets, he routinely removed bags of documents from CIA headquarters and deposited them at dead drops for his contacts. A search of his office after he was arrested revealed 144 classified intelligence reports not related to his current assignment.

Robert Hanssen was arrested in 2001 and charged with spying for Russia for more than 15 years [12,41]. The case is different from Ames because Hanssen had a high degree of technical expertise. Hanssen gathered classified information to sell by exploiting the FBI's computer system. He also used his access to a file system containing classified information about ongoing cases to see if he was being investigated. To exchange information Hanssen made extensive use of encrypted floppy disks, removable storage devices, and a handheld computer. Hanssen was responsible for providing over 6,000 pages of classified documents and the identities of three Russian agents working for the United States.

Anna Montes was a senior intelligence analyst at the Defense Intelligence Agency [12]. In 2001 she was arrested for transmitting sensitive and classified intelligence information to Cuba for over 16 years. Montes communicated with the Cubans through encrypted radio bursts and would meet with her contacts every three or four months to exchange encrypted disks of information. Montes left evidence behind by not removing all traces of the messages from her computer hard disk.

In 2000, Timothy Smith was a 37 year old civilian serving as an ordinary crewman on a US Navy ammunition and supply vessel [12]. He became upset by some mistreatment from his crewmates and decided to get revenge by stealing and selling classified materials to terrorist groups. When Smith was arrested 17 disks and five confidential documents were discovered in his possession, including one describing the transfer of ammunition and handling of torpedoes on US Navy vessels.

In 1978 Stanley Mark Rifkin, a 32 year old computer expert working as a consultant for the Security Pacific National Bank in Los Angeles, discovered the secret computer code the bank used to transfer funds to other banks telegraphically [13,33]. He used this information and the knowledge of the bank's computer system to transfer over ten million dollars to an alias account in New York. Using a phony passport and documentation, he then had the money transferred to an account in Zurich. Although Rifkin used no guns, bombs, or physical threats, his trusted position and understanding of the system led to one of the largest bank robberies in history.

The remaining case studies are taken from the Insider Threat Studies [21,29] in which the identities are not included in the literature. A city government employee became disgruntled when passed over for a promotion. Out of vengeance she deleted important files the day before the new person took office. It was never determined if all of the deleted files were recovered.

A system administrator who developed and managed the computer network for a manufacturing firm was angered by his diminishing role. In retaliation, he centralized the companies manufacturing processes software to a single server and planted a logic bomb. He then intimidated a coworker into giving him the backup tapes for the software and detonated the logic bomb, deleting the only remaining copy of the critical software causing an estimated \$10 million in damage and leading to the layoff of 80 employees.

An application developer who was laid off just prior to the Christmas holidays launched a systematic attack on his former employer's computer network using the username and password of one of his former coworkers to gain remote access. He modified several of the company's web pages by changing text and inserting pornographic images and sent the company's customers an email publicizing that the website had been hacked. A month and a half after the initial incident, he again remotely accessed the network and executed a script to change 4,000 pricing records and reset all network passwords.

In March 2002, ten billion files in the computer systems of an international financial services company were deleted by a logic bomb. The logic bomb had been planted by an employee that recently quit because of a dispute over the amount of his annual bonus. The incident affected over 1,300 of the company's servers and cost an estimated \$3 million to repair and reconstruct damaged files.

A common theme in the malicious insider cases is the attacker had both an opportunity and motive. The opportunity was provided by rights obtained through knowledge about the organization or granted permissions. The insider's motives varied from financial gain, ideology, ego gratification, under appreciation, disgruntleness, and/or revenge. The technical ability of the perpetrators also ranged greatly from highly advanced to simple user. Although their computer skills differed significantly, all of them left behind trails of suspicious activity when performing access, communication, or modification to the system.

2.3 Security Models

The case studies demonstrate the ability of users to leverage their rights in the system to cause damage. This problem can be mitigated by determining what rights a user can obtain and if the rights violate an organization's security policy. A method for examining this is a security model.

Security models provide a formal way to analyze the safety of a system. A system is considered safe if rights to system resources cannot be obtained by an unauthorized subject. Representing a security policy through the formal rules of a security model allow logical analysis to determine how violations occur and what mitigation techniques are required to maintain the policy. If the model shows a system is safe, then a system implementation using the model can result in a secure system [23].

2.3.1 Security Models for the Insider Threat. There has been relatively little work in developing a security model that includes the insider threat. Chinchani

et al. provide the only research found to date that formally attempts to model the insider threat [7]. Their methodologies propose a model that can be constructed by a security analyst and tailored to a specific organization's policy and structure. A physical entity is represented as a vertex in the graph, with each piece of information or capability that can be acquired by the vertex represented as a key. A directed edge between vertices represents a communication channel or access. A key challenge graph, similar to an access control list, defines actions that can take place in the system. Using the key challenge graph and procedures for traversing an edge, activity not allowed through the model can be detected.

Although this model addresses some concerns, it does not specifically define the safety of a system or provide a comprehensive representation of the threat. The scheme does not take into account some insiders use authorized rights to perform their attacks which would not be identified as malicious activity in this model. Additionally, this model does not address the safety question in a decidable manner. There is no process for tractable analysis and no maximal state exists to examine the transfer of rights.

2.3.2 Schematic Protection Model. There have been numerous formal security models developed to analyze the safety of a system. None of these, however, were designed with the insider threat in mind. One such model, the Schematic Protection Model (SPM), was developed in 1988 to answer the safety question for a generic but useful system [3]. The safety question determines whether or not a system can be formally proven secure. SPM is discussed in detail in subsequent sections because it is used as the basis for an insider threat security model.

2.3.3 Model Specifics. SPM defines the privileges possessed by subjects, called tickets, and determines how the tickets can flow amongst entities in the system [32]. If a system is in a safe initial state, then SPM can determine for a large class of systems whether derivable states are safe.

Each entity in the model has a protection type and a ticket flow based on this type. Entities are assigned a protection type upon creation that does not change. An entity is either a subject or object, where subjects are active entities and objects are passive with regard to ticket transfer. Entity types are partitioned into a subject set (TS) and an object set (TO), with the union of the two sets representing all types (T). The type of an entity X , is identified using the function $\tau(X)$.

The privileges that can be granted to a subject are identified as the rights in the system. The rights are specified as either an inert right (RI) or a control right (RC), and the set of all rights formed by their union (R). Inert rights cannot be used to change the protection state of the system. For example, read, write, and execute do not change the protection state of a system. Control rights, however, are rights that affect the protection state of the system and the distribution of privileges.

Tickets represent capabilities associated with a system and grant privileges to a subject. A ticket names a specific entity and a right associated with that entity. For example, the ticket Y/x authorizes the possessor to perform the operation associated with the right x on the entity Y . Multiple tickets for the same entity, like Y/u , Y/v , Y/w can be abbreviated to Y/uvw . The domain of a subject (Dom) specifies the set of tickets possessed by the subject.

The transfer of rights can occur if three conditions are met:

1. a copyable version of the ticket is in the domain of the subject transferring the right,
2. a link exists between subjects involved in the transfer, and
3. the filter associated with the link allows the tickets to pass over the link.

A copyable version of a ticket is specified by the copy flag (c). For instance, the ticket $Y/x:c$ is a copyable version of Y/x . Absence of the copy flag means the right cannot be copied to another entity.

A link predicate determines if there is a connection between two subjects. If the predicate evaluates to true, then a connection exists and can be used to copy tickets from one domain to the other. A connection exists between X and Y for any right $z \in RC$ for the conjunction or disjunction of the following:

1. $X/z \in \text{Dom}(X)$,
2. $X/z \in \text{Dom}(Y)$,
3. $Y/z \in \text{Dom}(X)$,
4. $Y/z \in \text{Dom}(Y)$,
5. true.

For example, the predicate:

$$\text{link}(X,Y)=Y/u \in \text{Dom}(X) \vee X/v \in \text{Dom}(Y)$$

evaluates to true if X has u rights over Y or Y has v rights over X. If the predicate is true, X and Y are connected. Rule 5 represents the universal link that does not depend on the entities rights. That is, a connection between X and Y exists regardless if the entities have tickets that refer to each other.

For a copy to occur, the ticket must also be specified in the appropriate filter function. Each link predicate (denoted by subscript i) has a corresponding filter:

$$f_i : TS \times TS \rightarrow 2^{T \times R} \quad (2.1)$$

The function f_i maps $TS \times TS$ to the power set of $T \times R$ and simply specifies the range of copyable tickets that can be transferred between two subjects. This function is the final condition required for the transfer of rights to occur.

Thus, a right can be transferred from one entity to another provided three specific requirements are met: the right has a copy flag, a link exists between the two entities and the filter allows the transfer of rights. Formally, $Y/x:c$ can be copied from $\text{Dom}(A)$ to $\text{Dom}(B)$, if and only if all of the following are true for some i :

1. $Y/x:c \in \text{Dom}(A)$,
2. $\text{link}_i(A,B)$,
3. $\tau(Y)/x:c \in f_i(\tau(A),\tau(B))$.

The create operation introduces new subjects and objects into the system and is specified by can-create (cc). For a subject of type a to create entities of type b , then $cc(a,b)$ must hold. If can-create holds, the create-rule $cr(a,b)$ specifies the tickets that are created. The two different types of creates are: a subject creates an object or a subject creates a subject.

When a subject creates an object, the tickets for the object are placed in the domain of the subject. For example, let $RI=\{r:c,w:c,x:c\}$ and $cr(a,b)=\{b/r:c, b/w:c\}$. If subject A creates an object B, it adds tickets $B/r:c$ and $B/w:c$ to its domain under the rule specified by $cr(a,b)$.

When a subject creates another subject, the creator can be granted tickets over the new subject and the new subject can be granted creator tickets. The create-rule is specified as: $cr(a,b) = LEFT|RIGHT$. If subject A of type a creates Subject B of type b , the tickets specified by the left part of $cr(a,b) = LEFT|RIGHT$ are placed in the domain of A and the tickets specified by the right part are placed in the domain of B. For example, let $R=\{r:c,w:c,x:c\}$, $cr(a,b)=\{b/r:c,b/w:c \mid a/r:c\}$. If subject A creates subject B, A will add the tickets $B/r:c$ and $B/w:c$ to its domain and B will be created with $A/r:c$ in its domain. To avoid confusion, if two subjects of the same type are specified, such as $cr(a,a)$, the special symbol *self* is introduced to identify tickets associated with the creator. For example, the ticket *self/w* refers to the creator and a/w refers to the created subject.

The final two issues when creating an entity deals with attenuation of privileges and the rule of acyclic creates. The principle attenuation of privileges states no entity may have more rights than the entity that created it. This policy is enforced through the create-rules. The rule of acyclic creates limits the creation of subject types such that it is not possible for a subject to directly or indirectly create a new subject of

the same type as its creator. For the creation $cc(a, b)$, subject A can create a subject B, however subject B and none of its subsequent children can create a subject type a . This eliminates cycles that could otherwise be associated with cc .

2.3.4 Example. The following example from Bishop demonstrates an owner-based policy in which the owner of an object can authorized another subject access to the object [3]. Consider user Peter wants to give another user Paul execute permissions to a file he owns called *doom*. The SPM specification is: $\tau(\text{Peter}) = \tau(\text{Paul}) = \text{user}$, $\tau(\text{doom}) = \text{file}$, and $\text{doom}/x:c \in \text{Dom}(\text{Peter})$. All users are considered connected and any user can give rights away to any other user, so $\text{link}(\text{Peter}, \text{Paul}) = \text{true}$ and $\tau(\text{doom})/x \in f(\tau(\text{Peter}), \tau(\text{Paul}))$. Because the ticket has the copy flag, a link exists, and the filter includes the ticket, Peter can copy the ticket doom/x to Paul.

2.3.5 Extensions and Implementation. There have been several extensions to SPM since its initial development. One created a process for conditional tickets and a means to provide authentication [40]. Another extended SPM by incorporating the revocation of privileges [39]. These extensions demonstrate the ability to adapt the model to address different situations.

SPM provides a formal means to measure and analyze the safety of a system. It demonstrates the safety question is decidable if the schema is acyclic and attenuating. This attribute allows SPM to incorporate security policies into a model and demonstrate the effectiveness of the policies if implemented into a system. However, even with the expressiveness of the model and ability to characterize the security of a system, no publications were found that use SPM to model the insider threat.

2.4 Risk Analysis

The security of a system is different than the safety of a system. Security refers to the implementation of a protection system. It is possible for a system to be safe with respect to all rights but the implementation is not secure [3]. A security

model is used to perform a theoretical analysis of a system. Although this provides a fundamental baseline, the intricacies and complicated nature of systems make implementing a completely secure system unfeasible. Mechanisms designed to prevent or detect attacks are often full of vulnerabilities that can be exploited by a malicious individual [42]. The insider threat magnifies this problem because individuals are granted access rights to the system. For this reason, risk analysis is critical in maintaining the security of a system.

Risk analysis provides a means to formally identify individuals that pose a risk to the security of a system. There are many factors that may contribute to an individual's risk, ranging from behavioral characteristics to access on the system. These factors in the form of behavioral attributes or technical activities are indicators that can identify a potential insider threat.

2.4.1 Workshops. In August, 2000 an insider threat workshop sponsored by the Office of the Assistant Secretary of Defense for Command, Control, Communications and Intelligence (C3I) and the Defense Advanced Research Projects Agency (DARPA) met to discuss the malicious insider and develop common themes for mitigation. One of the significant findings is the requirement for a comprehensive process for identifying the level of risk posed by system users [1]. The group devised a conceptual framework, depicted in Figure 2.2, consisting of three major components: People, Tools, and Environment. This workshop became one of the first to address the need for formal risk analysis for the insider threat.

In 2004, the Advanced Research and Development Activity in Information Technology (ARDA) devised a six month insider threat challenge workshop. In this collaborative effort, experts in computer security met to create analysis methods to counter malicious insiders in the US intelligence community [26]. Figure 2.3 defines a taxonomy of cyber events derived from investigation of previous cases. Based on these indicators the primary detection strategies in Figure 2.4 were developed: profiling and data flow analysis (Stealthwatch); likely actions based on established patterns (Struc-




Figure 2.2: Framework for Insider Threat Risk Analysis Presented at the C3I/DARPA Workshop [1].




Figure 2.3: Taxonomy of Cyber Events Developed at the ARDA Workshop [26].

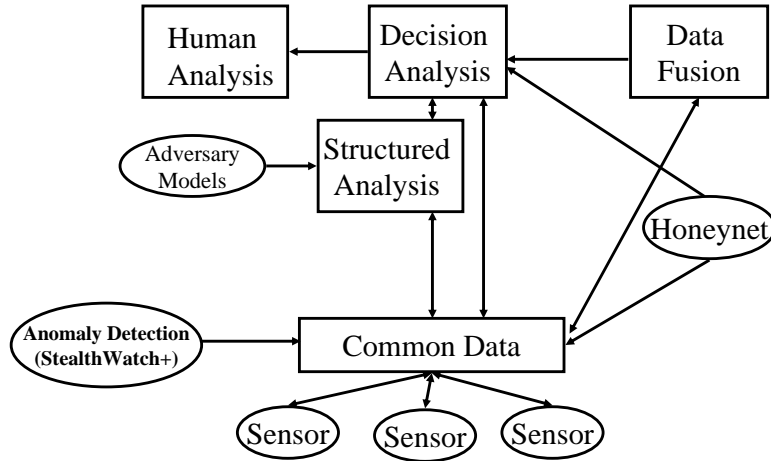


Figure 2.4: Detection Strategies Represented by ARDA Insider Threat Risk Analysis [26].

Structured Analysis); correlating inputs from the network application level (Data Fusion); and false production systems used to lure the malicious insider (Honeynets). These techniques associate behaviors and identify potential suspects through a Common Data Repository that sends indicators to the Decision Analysis for a human analyst to review. The workshop produced significant results using this framework to perform real-time detection of simulated malicious insiders on a live network test. The research and findings demonstrated a requirement for refined risk analysis methods, an observable taxonomy, and more sophisticated detection algorithms.

2.4.2 Frameworks for Identifying the Insider Threat. The aforementioned workshops emphasize the need for better identification of individuals that are potential insider threats. This section further reviews published work that attempts to formalize a method for identifying these threats.

Schultz devised the framework shown in Figure 2.5 for recognizing insiders through well-defined characteristics consisting of personality traits, verbal behavior, correlated usage patterns, preparatory behavior, meaningful errors, and deliberate markers [34]. His methodology examines these classifiers and determines through experience the type of attacks that are likely to occur. His belief is the indicators

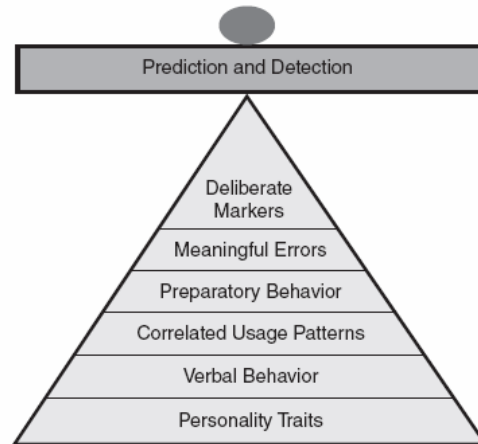


Figure 2.5: Insider Threat Framework Developed by Schultz [34].

can be quantified and expressed using multiple regression formulas and thus predict the likelihood of an attack. Although Schultz's methods are unproven, his framework emphasizes the importance of being able to analyze the threat in a measurable manner.

Wood presents a technique based on knowledge, tactics, and a predictable process [43]. He suggests a systematic method for simulating the behavior of the malicious insider by specifying the rationale and attributes of the perpetrator. The way an incident can be carried out is defined by distinct steps in whereby an individual is motivated, determines the target, plans the attack, and finally executes the attack. The insider is assumed to be able to obtain the privileges needed for an attack and has extensive knowledge of the system. The main problem with these assumptions is that it limits the scope of possible suspects and may lead to a malicious insider being overlooked. Additionally, a systematic method is important, but there is no formal way to analyze Wood's process. A risk analysis framework should be defined in a distinct manner that is capable of quantifying an individual's threat to the system.

Magklaras and Furnell have devised an insider threat prediction tool (ITPT) that estimates the level of threat based on certain profiles of user behavior [25]. They define a taxonomy that determines the threat level associated with an individual

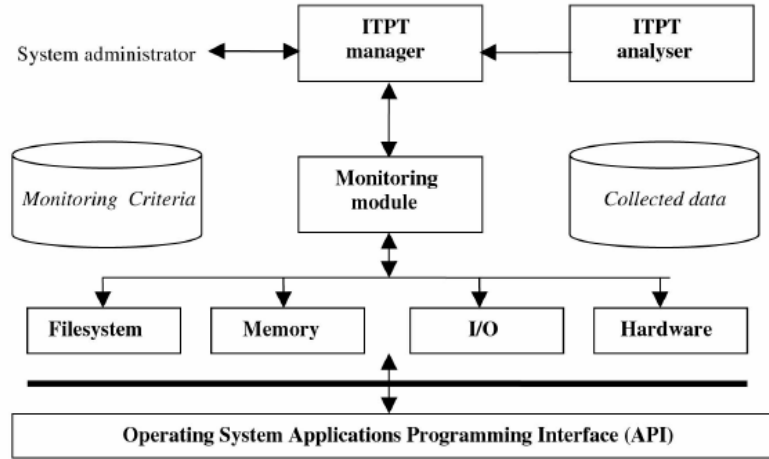


Figure 2.6: Insider Threat Prediction Tool Proposed by Magklaras [25].

by classifying users into three basic dimensions: system role, reason of misuse, and system consequences. A mathematical formula estimates threat levels based on the three dimensions and particular aspects of insider attributes and behavior. Once a threat level is identified, the ITPT architecture is defined and uses a variety of system modules to process and collect data. The collected data from the filesystem, memory modules, input/output, and hardware devices are delivered to the monitoring module for collaboration and detection techniques as shown in Figure 2.6. This methodology views the insider threat at system level and user profile. However, the framework lacks a practical implementation scheme and does not specifically address the indicators that produce potential threats. A more complete and formal method of classifying the risk for an insider threat is needed.

Reiher proposes using Anomaly Detection Systems (ADS) to identify malicious insiders [30]. ADS characterizes normal pattern usage and identifies behavior that does not conform to those patterns. Reiher analyzed file access by recording when a user performed an open, close, read, or write to a file. A time stamp measures the length of time a file was open and the time the operation occurred. Using the information gathered on 10 users over a 2-year time period, Reiher created a profile for what he labeled normal usage. This data was used to train the ADS so any pattern of

file access that did not match the profile was flagged as suspicious. Reiher tested the ADS by creating scenarios to see if what he believed to be malicious activities were identified. Although the research successfully classifies different file accesses, there needs to be real-world implementation or analysis against case studies to determine if file access is truly capable of determining malicious behavior. Additionally, while ADS algorithms are improving they are still susceptible to a high rate of false alarms [24]. Measuring an individual's usage patterns appear promising at producing indicators, but the use of this technique alone is likely not enough for system wide risk analysis.

The workshops and published works demonstrate risk analysis for the insider threat relies heavily on measuring behavior characteristics and access within the system. Unfortunately, the current methodologies have not demonstrated effectiveness through case studies or implementation. Identifying the malicious insider is an inherently difficult problem that expands across many areas of expertise such as social, behavioral, and technical disciplines. Perhaps the solution is in a multidisciplinary approach, where the factors from each area are leveraged together to produce an effective risk analysis framework.

2.5 Summary

This chapter reviews the attributes and characteristics of a malicious insider by investigating some documented case studies and trend analysis. A security model for the insider threat is discussed along with defining the safety of a system. SPM is also reviewed to explore characterizing a system in a decidable and quantifiable manner. Risk analysis frameworks for the insider threat are also discussed. The remainder of this document explores these concepts to develop a systematic method to address the insider threat through a formal security model and risk analysis.

III. Methodolgy

This chapter outlines the methodology used to develop a formal security model and risk analysis for the insider threat. It provides the necessary information to formalize the insider threat problem in a systematic and definable manner.

3.1 *Problem Definition*

This research addresses two specific areas of concern for the insider threat: the safety of a system and identifying individuals that pose a threat. The process of analyzing how rights can transfer within the system determines the safety of a system and can be evaluated through security models. Risk analysis provides a formal means for identifying individuals that pose a risk to the system.

3.1.1 *Goals and Hypothesis.*

Security Model. There is currently no security model capable of determining the safety of a system against the insider threat in a quantifiable and deterministic manner. The first goal of this research is to develop a comprehensive security model that adequately determines the safety of a system against the insider threat.

Current mitigation techniques for the insider threat typically focus on the mechanisms for preventing or detecting an attack. Examples of these type of systems are Anomaly Detections Systems, event logs, file access monitoring, and Honeypots. A security model, however, uses a different strategy: expose weaknesses in the system before an attack by determining if a state that violates the security policy can be reached. For example, an organization's security policy states that only persons in department Q can access the network drive Z . When a security model of the system is analyzed, however, it is determined that an individual in department X can obtain access to the network drive Z . The process has discovered a vulnerability that violates the security policy and a mechanism is required to prevent the access.

It is expected that by using SPM as a foundation a formal means to measure and analyze the safety of a system against the insider threat can be produced. It has already been demonstrated SPM can be extended to model different conditions and the safety question is decidable for a system specified in SPM [29]. These attributes make SPM a formidable candidate for this research.

Risk Analysis. The second goal of this research is to present a risk analysis framework using a multidisciplinary approach capable of detecting potential malicious insiders.

Performing risk analysis for the malicious insider is an inherently difficult problem that transcends social, behavioral, and technical disciplines. Unfortunately, current methodologies to combat the insider threat have not proven effective primarily because techniques have focused on these areas in isolation. The technology community is searching for technical solutions while the law enforcement and counterintelligence communities 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.

This research proposes a multidisciplinary approach with a clearly defined methodology that attacks the problem in an organized and consistent manner. The hypothesis is that focusing on the collaboration of information to determine indicators and using statistical analysis to identify potential malicious insiders, effective risk analysis for the insider threat is possible.

3.2 Security Model for the Insider Threat

The model to evaluate the safety of a system against the insider threat is based on the principles of the original SPM and is referred to as the Schematic Protection Model for the Insider Threat (SPM-IT). SPM-IT analyzes security policies and implementation schemes to determine whether vulnerabilities exist. Using this information, mechanisms can be implemented or policies changed to mitigate the threat.

A security policy, at the simplest level, defines a set of safe states where distribution of privileges is consistent with the underlying objectives [32]. For example, the policy user X cannot read file Y defines a state for the system to maintain safety. Strict analysis of assigned privileges alone is not sufficient for determining the safety of a system. The distribution of rights must also be analyzed to ensure the possible states a system may reach are safe. The dynamic aspect of a security policy may allow the transfer of rights to result in an unauthorized state. In the previous example, if X cannot read Y initially but user Z transfers this privilege to X, then the system has reached an unsafe state. SPM-IT analyzes the dynamic allocation of rights and determines if transitions lead to an unsafe state in the context of the insider threat.

The safety of a system is an issue in systems that provide controlled sharing of information among multiple users [32]. The safety question is defined as determining if a given system with an initial state is safe with respect to a generic right [3]. The original SPM demonstrates the safety question is decidable and tractable provided the schema is acyclic and attenuating. The SPM-IT maintains this quality through the attributes of the SPM and demonstrates the safety question is decidable and tractable for the insider threat.

3.2.1 Approach. A taxonomy characterizing the insider threat through measurable and distinct actions is developed. The taxonomy is specified through the SPM attributes, extending the framework as necessary. The resulting SPM-IT is used for analysis to determine the safety of a system against the insider threat. Tractable analysis and implementation of mitigation techniques are demonstrated through one instance of the model.

3.2.2 Model Boundaries. To ensure the model adequately addresses the insider threat, it is necessary to clearly define the aspects that are encompassed by the model. An insider is any individual who has been granted any right in an information system. This description does not limit the insider to specific borders such as Firewalls, Routers, or a Local Area Network. The system itself could be a conglomeration

	File 1	File 2	Alice	Bob
Alice	own read write	read	own read write	
Bob	read	own read write		own read write

Figure 3.1: An ACM representation with two subjects and two objects.

of networks. What is important is that once the user has been granted any authorized right to the information system, they are now considered an insider and are included in the system protection state.

The protection state is the current state of all rights for all users and objects in the information system. The protection state encompasses all activities that are allowed according to organization policy or system access controls. The most precise model to describe a protection state is the Access Control Matrix (ACM) [3]. Figure 3.1 shows an ACM with subjects Alice and Bob and their rights to objects File 1 and File 2. Alice has own, read, and write rights over File 1 and is limited to read for File 2. SPM-IT captures the finite set of rights that can be represented through the protection state.

The malicious insider is any authorized user that uses rights to alter the system or the protection state of the system in an unauthorized way. For example, if an individual gains administrative rights and deletes files, they are a malicious insider. If an individual not associated with the system physically breaks into a building and places a packet sniffer somewhere on the network, they are not a malicious insider because the individual does not have rights on the system. They are more aptly categorized as a criminal. Physical access alone does not constitute an implied right

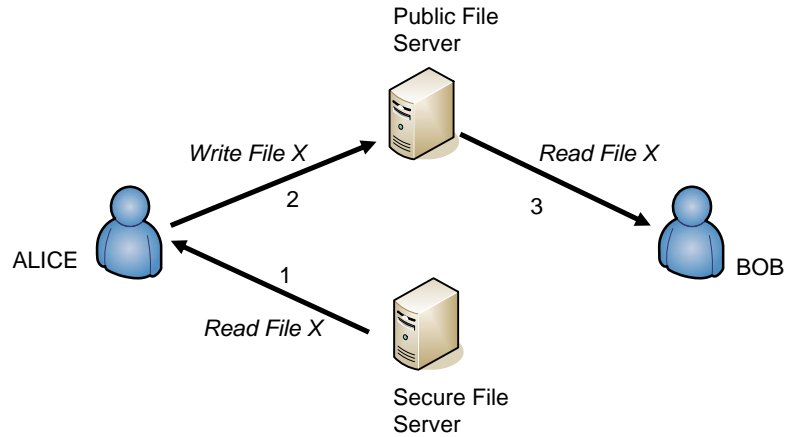


Figure 3.2: Example of Insider Threat Vulnerability.

and does not fall within the scope of this research and should be considered through different methodology.

3.2.3 Features. SPM-IT provides an ability to determine if a given system is vulnerable to attacks from malicious insiders. Specifically, it determines if the exchange of rights between entities can lead to a violation of the security policy. For example in Figure 3.2 Alice has read permissions over the secure object File X. Additionally, a communication channel exists between Alice and a public server so she can write the file to that server. From the public server, Bob can read the secure object, File X, even though he has no permissions to the protected server. To prevent this from occurring, the implementation scheme must be changed or a mechanism required that prevents Alice from writing a protected object to a public server.

3.2.4 Specifications. The model encompasses subjects and objects within a system along with their associated rights. Subjects are active entities that may perform operations (invoke rights) on another subject or object. Objects are considered passive and do not possess rights. Subjects and objects are assigned a protection type upon creation that is static.

Rights within the system are defined through the protection state and represent the operations that may be performed on subjects and objects. Tickets are the capabilities for the system and are specified through a specific entity and a right associated with that entity. A subject possessing the ticket X/r is said to have r rights over X . Ticket transfers may occur provided a link exists between the two entities, the filter allows the transfer, and the ticket has the copy flag set. Creation of a new entity is specified through can-create and create-rules.

3.2.5 Analysis. For analysis, a security policy is defined using SPM-IT. The policy is a simple one with the intention of expressing the effectiveness of SPM-IT in determining the safety of a system against the insider threat. Once defined through the SPM-IT, the system is analyzed to determine how tickets can transfer between subjects.

By design, analysis is performed using a worse-case scenario with respect to ticket transfer and possession. If a subject is capable of transferring a ticket, the ticket is transferred. Additionally, if a ticket is possessed by a subject, the subject will invoke the right.

The maximal state is the state which represents all possible ticket transfers. Analysis of the maximal state demonstrates where vulnerabilities exist in a system's policy and is used to determine the safety of a system. Mitigation techniques are then specified to demonstrate methods for eliminating or reducing the vulnerabilities.

3.3 Risk Analysis for the Insider Threat

The second goal of this research develops a risk analysis framework. For the insider threat, the concern is identifying individuals that display symptoms consistent with a malicious insider. By identifying these individuals, risk analysis prevents them from causing significant damage or more harm to the organization.

In the case studies reviewed in Chapter II, each individual had suspicious activities in the form of behavioral and technical attributes. Identifying these attributes in a comprehensive manner provides a means to perform risk analysis against the threat.

3.3.1 Approach. In this research, the Multidisciplinary Approach to Mitigating the Insider Threat (MAMIT) is the framework designed to perform risk analysis for the insider threat. Relevant indicators are identified by leveraging both behavioral and technical attributes. These indicators are quantified and combined to produce one threat level for each person in the system. The threat level for each individual is measured against the others in the organization to identify users that pose a threat. If an individual falls outside the acceptable statistic threat range, they are identified as a potential insider threat. The effectiveness of MAMIT is demonstrated using the well-known case study involving Robert Hanssen.

3.3.2 Specifications. An insider is any individual who has been granted access in an information system. Risk analysis identifies insiders that are planning to commit or have committed a malicious act against the organization using their rights. Indicators are the behavior and technical attributes produced by individuals within the system. Individuals that pose a risk exhibit a higher threat level than other users which if analyzed correctly can be identified through the indicators. Behavior indicators are the actions and the conduct of an individual such as financial activities or coworker interaction. Technical indicators are associated with usage of the system such as file accesses or document transfers. The consistent measurement and quantifying of indicators provides a formal method to identify potential threats using the MAMIT framework.

3.3.3 Attributes. The behavior attributes which produce relevant indicators are determined through analysis of historical case studies. The 150 case studies in the PERSEREC database [12] in conjunction with the Guidelines for Security Clearance [38] establish indicators that are most common and the best determinants consistent

with past malicious insiders. Technical attributes are determined through ability and opportunity of a user within the system.

The indicators necessary for measuring risk against the insider threat are identified within the context of the MAMIT framework. The indicators are assigned a quantitative value based on the perceived threat level of an individual for each specific indicator. The indicators are combined to produce a single identifier for risk analysis. The process is performed for each individual within the organization. Statistical analysis using the *t*-distribution and prediction interval determines an acceptable threat range for the organization. Anyone with a threat level not within the acceptable range is labelled a potential insider threat.

3.3.4 Analysis. The MAMIT framework identifies individuals that display a credible threat to the system. MAMIT is evaluated using the Robert Hanssen case study. Since the case study is well-documented and widely publicized, it is an effective candidate for analyzing the framework. The indicators are determined through documentation review and applied appropriately to the MAMIT framework. Statistical analysis within the framework specifies Hanssen's threat level in conjunction with an acceptable organization threat level. The study analyzes if Hanssen would likely have been identified using the MAMIT process.

3.4 Evaluation

Evaluation of both the security model and risk analysis is performed through analytical methods. SPM-IT is a theoretical model that examines the flow of rights within a system. Through the specified rules and enforcing the logical principles of the model, the correctness of the state transitions can be evaluated. Provided the flow of rights follow the principles and safety is assessed based on the definitions within SPM-IT, each instance of the model can be verified.

The MAMIT framework is a risk analysis process determined by quantifying identifiers and through statistical analysis. The principles and implementation of

MAMIT is validated through expert intuition. If an individual is a known malicious insider from a historical case, then the framework should identify him as an insider threat.

3.5 Summary

This chapter defines the methods used to develop a security model and risk analysis framework for the insider threat. The principles for analyzing a security policy and implementation scheme to determine the safety of a system are introduced. The development of SPM-IT by extending attributes of the original SPM is discussed. Additionally, a process for determining individuals within the organization that pose a threat is presented through a multidisciplinary risk analysis framework. The aspect of determining indicators through behavior and technical attributes is discussed. Finally, the analysis and evaluation techniques for the developed model and framework are presented.

IV. Developing a Security Model for the Insider Threat

This chapter presents a formal security model for the insider threat. Initially, a taxonomy is created that systematically defines the threat. The taxonomy is applied to the framework of SPM to develop the Schematic Protection Model for the Insider Threat (SPM-IT). An implementation of the model instantiates the safety analysis of a system and effectiveness of SPM-IT.

4.1 Taxonomy Development

To specify the insider threat through a security model the threat must be identified in a systematic, measurable manner. This requires a comprehensive taxonomy capable of classifying malicious insider activities. The taxonomy developed for this research decomposes an abstract threat into a solvable and analyzable process.

4.1.1 Approach. To develop a taxonomy for the insider threat, attack tree methodologies were examined. Researchers have proposed an attack tree is sufficient to address the outside threat and assess the security of a system against a compromise [20,27,37]. The attack tree structure places the goal of the attacker in the root node and different ways to obtain that goal depicted as leaf nodes.

Traditional attack trees, however, are not capable of capturing the insider threat effectively [7] since they do not provide a comprehensive model to reason about vulnerabilities [9]. One of the more significant problems is an insider may already have the rights needed to perform a malicious act. Additionally, the focus of the attack tree is obtaining a goal represented by the root node. Quantifying goals or motives of an attacker is difficult and still may not lead to an adequate representation of the threat.

Malicious insiders do not share a common profile, so there must be a different tangible way to produce a taxonomy if measurable results are to be obtained [29]. This research proposes a hierarchical tree capable of providing a malicious insider taxonomy using a systems engineering approach rather than the goal oriented objectives used

by attack trees. This representation focuses on the activities of the malicious insider and not their traits or attributes.

To represent a threat, actions that can lead to a violation of the protection state are identified. These activities are methodically investigated through functional decomposition, which addresses the problems associated with traditional attack trees. By decomposing actions with respect to the protection state, no user or threat is excluded. Additionally, an action either occurs or it doesn't so the methodology is measurable and analyzable. This systematic approach masks differences inherent with individuals and effectively classifies malicious behavior.

4.1.2 Methodology. The protection state defines through rights what actions are authorized within the system. The Access Control Matrix (ACM) is the classic representation of a protection state and defines operations through the finite set of actions [17]:

- Enter a right
- Delete a right
- Create a subject or object
- Destroy a subject or object

Consider the ACM in Figure 4.1. The right can be entered into the matrix to allow Alice write File 2. Removal or deletion of read File 1 from Bob is possible as well. File 2 can also be removed (destroyed). From a malicious standpoint, these actions can be performed by a subject to intentionally alter the protection state. Additionally, in the system represented by this example Alice can invoke read File 1 and potentially obtain unauthorized information. Alice may also grant write to Bob for File 1 because she is the owner.

	File 1	File 2	Alice	Bob
Alice	own read write	read	own read write	
Bob	read	own read write		own read write

Figure 4.1: An ACM representation with two subjects and two objects.

The insider threat taxonomy is specified through the actions that can result in malicious activity in the context of the protection state. These activities have been categorized into four distinct actions:

1. Invoke right to obtain unauthorized information (Snooping)
2. Enter right to gain unauthorized privileges (Elevation)
3. Delete or change a subject or objects rights, to include destroy a subject or object (Alteration)
4. Transfer a right to an unauthorized entity (Distribution)

Each activity is unauthorized if it violates organization policy or system access controls. These actions capture the possible malicious events that can produce a transition in the protection state. The malicious insider is therefore someone who violates the protection state of the system and is depicted in Figure 4.2 as the root node of the tree.

Alteration. Alteration occurs when a malicious insider changes a subject or object's rights in an unauthorized manner. A user deleting a file from the




Figure 4.2: The four actions represented in the first hierarchy of the tree.

system to deny access or launching a virus that corrupts entities on the system is an example of Alteration.

Distribution. Distribution is the transfer of rights to an unauthorized entity. This occurs when a user has appropriate system rights and a need to know to access a file but transfers it to an unauthorized entity. A user emailing a file to an unauthorized individual is an example of Distribution.

Snooping. Snooping occurs when a legitimate right is used to obtain unauthorized information on a user or object. This action is similar to Distribution except the user has appropriate system rights but lacks a need to know. The violation takes place when a user has permissions according to system access controls but the event violates organization policy. For example, an individual with administrative privileges is Snooping when he opens and reads another user's email to gain information. He has accessed something allowed according to the rights possessed but organization policy disallows it.

Elevation. Elevation takes place when a user obtains unauthorized rights in the system. A classic example of this is unauthorized acquisition of

administrative privileges. There are many different ways a malicious insider may accomplish this, from automated attacks to social engineering. Elevation addresses the notion of a malicious insider changing their rights and is an attempt to garner rights that are not already allowed by the system.

4.1.3 Example. The taxonomy created using functional decomposition ensures every activity of the malicious insider can be categorized in the context of the protection state. This establishes the underlying framework needed to identify a malicious insider in a deterministic fashion. That each activity can be captured by a specific action is an important and definitive concept.

It is best to explore this notion through a practical example. If Mallory gains administrator privileges by compromising a system and proceeds to delete Alice's email account, transitions in the protection state occur. Mallory is a malicious insider because her activities are intentional and deliberate. In this scenario there are two distinct actions that violate the protection state and subsequently there are two transitions of the protection state. The initial violation is Elevation by gaining administrator rights. The second violation is Alteration by deleting an email account and changing the system structure. Additionally, if Mallory accesses a secure document another violation occurs. Initially, when she gains administrator privileges through Elevation the protection state allows her access to the file. Although she now has these rights in the context of the protection state, Snooping has occurred because she still does not have an authorized reason (need to know) to view the file. Finally, if Mallory shares the document with Bob, Distribution has occurred because Bob has obtained rights to an object he shouldn't have. Thus, the problem is compartmentalized into distinct events.

4.1.4 Decomposition. Figure 4.3 demonstrates how to further decompose a malicious insider's activities. The threats are broken down step by step, beginning with the actions and continuing with the intermediate levels to the leaf node. This process is accomplished using a "how it can be performed" relationship between a

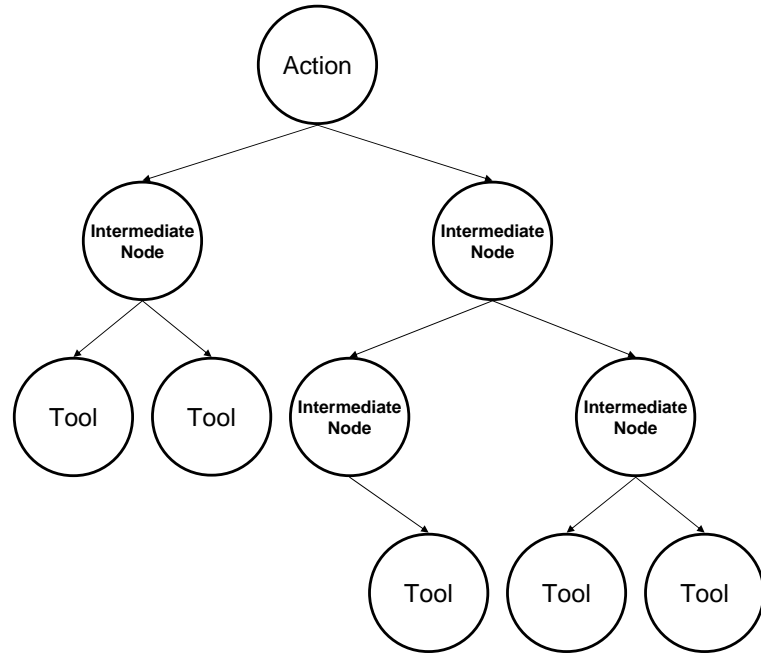


Figure 4.3: Decomposed tree representation.

parent and child node. The leaf node is the lowest level of abstraction and is the tool or technique the malicious insider uses to accomplish the activity. A path from the malicious insider (root node) to a tool (leaf node) forms a completely decomposed activity. The model is developed in an hierarchical acyclic fashion, meaning a malicious activity can only follow one specific path from the root node to a leaf node. This indicates that any activity is capable of being explicitly defined.

The following example uses this methodology for the Distribution action shown in Figure 4.4. The Distribution action can be performed through file sharing, via email, copying the file to storage media, online chat, or an electronic drop box. Sending email can be accomplished through a local or web based account. In addition, copying the file to storage can be performed by floppy disk, CD-ROM or USB drive.

Actions are limited to four distinct possibilities (Distribution, Snooping, Elevation, Alteration). The intermediate nodes, however, can use any number of children to describe its parent. This notion allows flexibility to tailor threats to the policies and

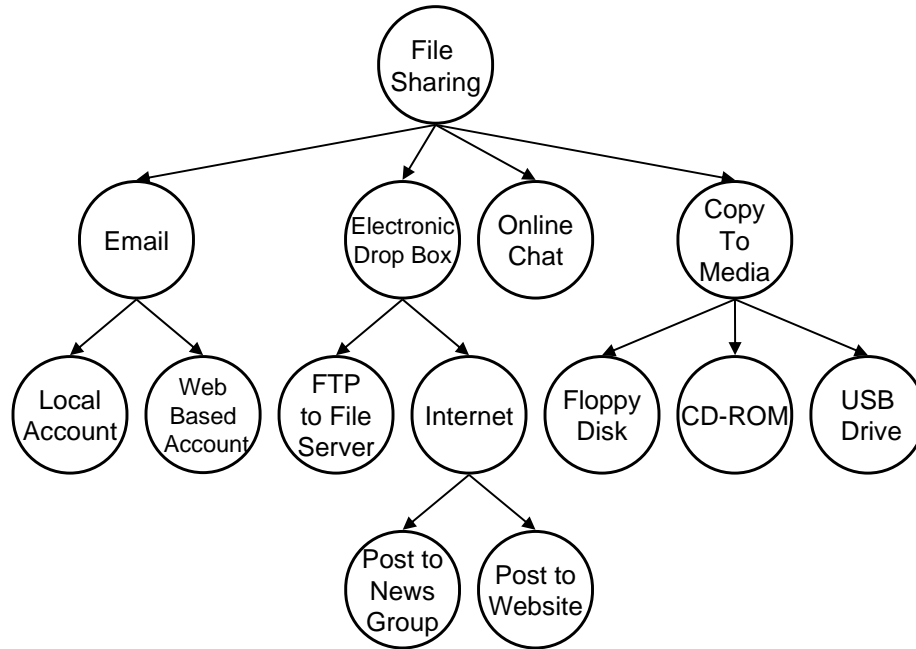


Figure 4.4: An example decomposing Distribution: file sharing.

specific environment of an individual organization, while still providing an analyzable and decidable taxonomy.

4.1.5 Taxonomy Attributes. The fundamental strength of this taxonomy is expressing the malicious insider through distinct actions capable of being decomposed and analyzed. It is a complete and well-defined representation of the insider threat because it defines actions and does not categorize individual attributes. Defining an attack using functional decomposition enables formal classification of the insider threat. This taxonomy is a systematic representation that can be applied to the development of analytical processes and security models.

4.2 SPM-IT

SPM-IT uses the original SPM, but extends it by incorporating the ability to model the insider threat using the above taxonomy and representing actions in the

context of the model. This formalization provides a framework to analyze the safety of a system and determine vulnerabilities of a given policy or scheme.

4.2.1 Rights. The original SPM specified two types of rights, control and inert. The purpose was to distinguish between rights that affect the protection state of a system (control rights) and those that do not (inert). For example, the read right is considered an inert right because reading a file does not change what entities have access to the document, and thus have no effect on the protection state [3]. Control rights, however, can alter the protection state. Because of this, SPM focuses on how control rights can be transferred between different entities. The insider threat, however, poses a different problem. Malicious activity can occur through both inert and control rights. For example, Snooping can occur by reading a file. By definition, for the insider threat this action is a violation of the protection state. For this reason, SPM-IT makes no distinction between the different types of rights and uses both for analysis.

In SPM-IT the two rights, read(r) and write(w), are used extensively to capture the actions associated with the insider threat. The right, r , indicates the ability to read an entity. The w right confers the ability to modify or delete an existing entity. These two rights are distinct and a subject is not required to have one to invoke the other.

Creating a duplicate of a subject or object is achieved through the combination of r and the can-create (cc) function. This functionality is referred to as the *save as* operation and is possible if a subject possesses r for a specific entity and cc for the subject includes the specified type of the entity.

Definition 1 (*save as operation*). $\text{Dom}(X)$ is the set of tickets possessed by Subject X. The creation of a replica of Subject or Object Y is possible iff:

1. $Y/r \in \text{Dom}(X)$,
2. $\tau(Y) \in cc(\tau(X))$

The resulting entity from the *save as* operation is indicated by $'$. For example, in the context of Definition 1, Y' is a copy of Y .

4.2.2 Actions. One of the benefits of the taxonomy is it encompasses the insider threat according to the types of actions that may occur to violate the protection state. These four distinct actions are now defined within the framework of SPM to characterize malicious activity. The representation of these four actions form the basis of SPM-IT and can be used to determine the safety of a policy or given scheme against the insider threat.

Alteration. This action encompasses modifying the information system structure in an unauthorized manner. Alteration occurs when a user maliciously changes a subject or object from one state to another. By definition, any subject that contains a ticket with w rights over another subject or object can perform an act of alteration.

Definition 2 (*Alteration*). Alteration can occur to Y by X iff:

1. $Y/w \in \text{Dom}(X)$

Snooping. This action obtains unauthorized information about a subject or object using legitimate access controls. The threat exists when a user has rights to perform the action, but should not because it violates organization policy. Snooping can occur whenever a subject has r rights over a subject or object.

Definition 3 (*Snooping*). Snooping can occur to Y by X iff:

1. $Y/r \in \text{Dom}(X)$

Distribution. This action transfers protected information and occurs when an unauthorized right is granted to a subject. Distribution occurs through the transfer of a ticket to an unauthorized subject or by creating a copy and distributing a ticket for the copy.

Definition 4a (*Distribution*). Let Subject U be an unauthorized subject and q be a generic right. Distribution of Y/q by X can occur when all of the following requirements are met:

1. $Y/q:c \in \text{Dom}(X)$
2. $\text{link}_i(X,U)=\text{true}$
3. $\tau(Y)/q \in f_i(\tau(X),\tau(U))$

For Distribution to occur, the subject must possess a ticket with a right for the target subject or object that contains the copy flag, a link must exist between the subject and unauthorized subject, and the filter must allow the ticket transfer from the subject to the unauthorized subject as shown in Figure 4.5.

A second way for Distribution to occur is for the subject to create a copy of the target entity using *save as* and distribute a ticket for the newly created copy. In this situation, the unauthorized subject gains a ticket for the replica and not the original.

Definition 4b (*Distribution of a copy*). Distribution of Y'/q by X can occur when all of the following requirements are met:

1. $Y/r \in \text{Dom}(X)$
2. $\tau(Y) \in cc(\tau(X))$
3. $\tau(Y)/q:c \in cr(\tau(X),\tau(Y))$
4. $\text{link}_i(X,U)=\text{true}$
5. $\tau(Y)/q \in f_i(\tau(X),\tau(U))$

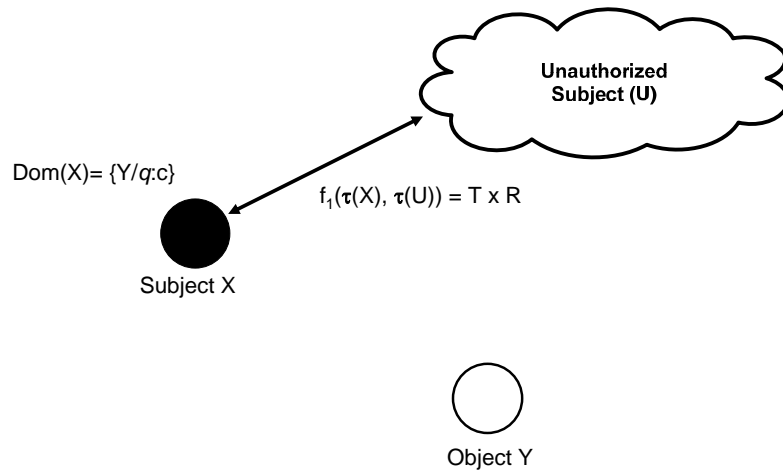


Figure 4.5: Distribution.

where the create-rule (cr) specifies the rights generated for the new entity after the cc function is invoked. The ticket Y/r in the domain of X and the cc function permits X to create a replica. Distribution can occur if a copy flag is included in cr , a link exists between the subject and unauthorized subject, and the filter allows the transfer as shown in Figure 4.6.

Distribution through association. The transfer of a ticket to an unauthorized subject can also occur through a combination of Alteration and Snooping even if no link exists between two subjects. That is, these two independent actions performed in conjunction can lead to an unauthorized ticket transfer. The threat as a whole is formally classified through Distribution. In this situation the unauthorized subject obtains a ticket for a replica and not the original entity.

Definition 4c (*Distribution through association*). Let Z be a Subject or Object. Distribution of Y'/q via X can occur when all of the following requirements are met:




Figure 4.6: Distribution of a copy.

1. $Y/r \wedge Z/w \in \text{Dom}(X)$
2. $\tau(Y) \in cc(\tau(X))$
3. $Z/r \in \text{Dom}(U)$

The ticket Y/r in the Domain of X coupled with the cc function allows X to perform a *save as* function. Additionally, X has the capability for Alteration to Z through the ticket Z/w . Together, these conditions allow X to effectively replace Z with a copy of Y . Once Z has been altered to match the original Y , U can invoke Z/r and obtain a replica of Y . This results in U performing Snooping by using granted rights to gain access to unauthorized information. The ability to gain a capability through Alteration followed by Snooping results in Distribution as shown in Figure 4.7.

Elevation. Elevation takes place when a user acquires unauthorized rights. This action implies some protection mechanism has been compromised to obtain the new ticket. Fundamentally, this threat is concerned with the security of a system and refers to the actual implementation of mitigation techniques [3]. Safety,

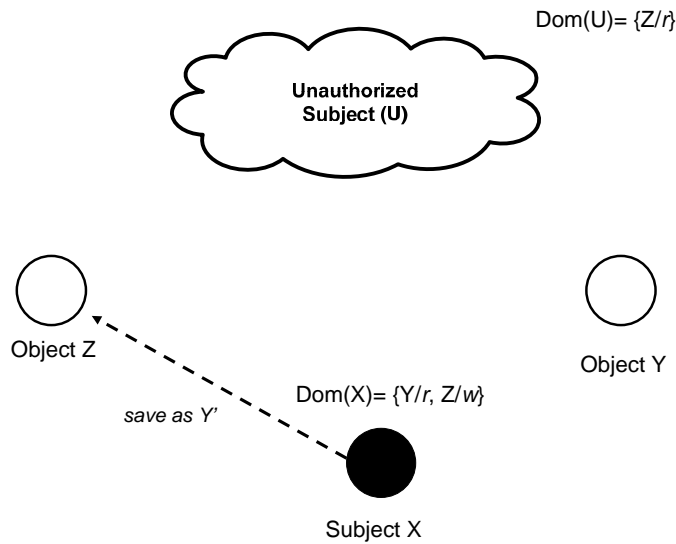


Figure 4.7: Distribution through association.

on the other hand, is an abstract model of the system and represents states the system can reach with respect to rights. It is possible for a system to be safe with respect to all rights but implementation of the system is not secure. For example, if the system is safe and rights cannot leak to unauthorized individuals but it can be exploited using a buffer overflow, the system is considered safe but the implementation of the system is not secure. Because SPM-IT analyzes the safety of a system, Elevation is not modeled and is a threat only if the system is not implemented in a secure manner.

4.2.3 Ticket Use and Transfer. There are subtle differences between how different vulnerabilities arise. Both Alteration and Snooping invoke a ticket. This means to prevent these threats, the focus should be on eliminating possession of the rights or ensuring they can only be invoked correctly. In contrast to those actions, Distribution occurs through the transfer of tickets. To mitigate this action, techniques on preventing the unauthorized flow of tickets between subjects is needed.

4.3 Analysis

This section applies SPM-IT and analyzes an implementation. The scheme considered expresses the power of the SPM-IT framework and demonstrates the process of determining the safety of a system using the model. Because SPM-IT preserves the generality and tractable analysis of the original SPM, the methodology can be extended to model more complex policies and systems.

4.3.1 Policy and Implementation. The scenario models a university setting where there are professors, students, and system administrators. Professors have access to a file server for professors as well as a common file server. Students have access to a student file server and also the common file server. The administrators have full control over the professors, students, and all servers. The rights are r which allows read and w which allows modify or delete for the associated subject or object. The organization policy states that users may only access files on their respective servers and the common server. Additionally, access should only occur if the information is required to perform their duty. Professors, students, and administrators have the capability to communicate through email. Through inheritance, a subject with a ticket for a server also assumes the ticket for all of the files on that server. For example, presenting the ticket for the shared file server automatically allows access to any file on that server. Formally, the policy is defined as follows:

1. $TS = \{\text{Prof, Stu, Admin}\}, TO = \{\text{server, file}\}$
2. $R = \{r:c, w:c\}$
3. $\text{link}_1(P, S) = \text{true}$
 $\text{link}_2(P, A) = \text{true}$
 $\text{link}_3(S, P) = \text{true}$
 $\text{link}_4(S, A) = \text{true}$
 $\text{link}_5(A, P) = \text{true}$
 $\text{link}_6(A, S) = \text{true}$

4. $f_1(\text{Prof}, \text{Stu}) = \text{T} \times \text{R}$
 $f_2(\text{Prof}, \text{Admin}) = \text{T} \times \text{R}$
 $f_3(\text{Stu}, \text{Prof}) = \text{T} \times \text{R}$
 $f_4(\text{Stu}, \text{Admin}) = \text{T} \times \text{R}$
 $f_5(\text{Admin}, \text{Prof}) = \text{T} \times \text{R}$
 $f_6(\text{Admin}, \text{Stu}) = \text{T} \times \text{R}$
5. $cc(\text{Prof}) = \{\text{file}\}$
 $cc(\text{Stu}) = \{\text{file}\}$
 $cc(\text{Admin}) = \{\text{Admin}, \text{Prof}, \text{Stu}, \text{server}, \text{file}\}$
6. $cr(\text{Prof}, \text{file}) = \{\text{file}/rw:c\}$
 $cr(\text{Stu}, \text{file}) = \{\text{file}/rw:c\}$
 $cr(\text{Admin}, \text{Admin}) = \{\text{self}/rw:c\} \mid \{\text{Admin}/rw:c\}$
 $cr(\text{Admin}, \text{Prof}) = \{\text{Prof}/rw:c\} \mid \emptyset$
 $cr(\text{Admin}, \text{Stu}) = \{\text{Stu}/rw:c\} \mid \emptyset$
 $cr(\text{Admin}, \text{server}) = \{\text{server}/rw:c\}$
 $cr(\text{Admin}, \text{file}) = \{\text{file}/rw:c\}$

The filters currently do not block any transfer through the links, which is a typical implementation scheme for most organizations. Additionally, the can-create function and create-rules specify an owner-based policy that allows each subject to create a file and retain control over the file.

This example for a policy is basic by design. Even so, it demonstrates the power and effectiveness of SPM-IT through a scenario similar to circumstances exploited by malicious insiders. For analysis, one instance of each subject and object are examined to determine the capabilities and type of threats in the implementation scheme. These entities are:

$$\tau(\text{P}) = \text{Prof}$$

$$\tau(\text{S}) = \text{Stu}$$




Figure 4.8: Initial State.

$\tau(A) = \text{Admin}$

$\tau(\text{server}_p) = \tau(\text{server}_s) = \tau(\text{server}_z) = \text{server}$

$\tau(\text{file}_{p1}) = \tau(\text{file}_{s1}) = \tau(\text{file}_{z1}) = \text{file}$

Figure 4.8 shows the initial state of the system. Notice the administrator possesses a copy flag for each ticket. This is indicative of a typical scheme that allows the administrator to have full control over the entities within the system. The copy flag permits the transfer of tickets to other entities, for example an administrator assigning a subject privileges to an object.

4.3.2 Alteration Threat. The scheme is now analyzed for the Alteration vulnerabilities. Subjects P and A are currently the only ones with the ability to perform Alteration on file_p1 (cf., Definition 2). Suppose file_p1 is a critical asset that cannot be modified. To prevent Alteration, a policy change or mechanism would have to be implemented to eliminate the ticket server_p/w for P and A. Additionally, if file_p1 has been altered, forensic analysis of the model can demonstrate what subjects

are capable of performing that action. The subjects that possess the capability for Alteration, as well as the other threats, will likely change through the transfer of tickets. The implications of performing threat analysis for the systems where tickets are transferred are discussed in detail in the maximal state section.

4.3.3 Snooping Threat. Using `file_p1` again, both P and A have the capability for Snooping (cf., Definition 3). By examining the model, an analyst can determine the subjects that have a valid reason for accessing the object. It is likely that P has a legitimate requirement for accessing the file and the vulnerability for the system exists through access via A. This example demonstrates the power of the system administrator that is prevalent in most systems. Snooping is generally a threat for administrators because they typically have access to the files but do not necessarily have a need to know for the information contained in the file. To eliminate or minimize the threat, a mandatory access control policy or mitigation technique must be implemented that reduces the span of control of the administrator.

4.3.4 Distribution Threat. Figure 4.9 demonstrates the threat to the system through Distribution (cf., Definition 4a). Subject A can transfer a copy of the ticket `server_p/rw` to S because A has the copy flag, a link exists, and the filter allows the operation. This transfer is synonymous with an administrator granting a student rights to the server.

A threat also exists for Distribution of a copy (cf., Definition 4b). Subject P uses *save as* to create a local file, `file_p1'`. P retains full rights to the new object through the owner-based policy specifications in the create rules. Using the existing connection between P and S, a ticket for `file_p1'` can be transferred. Figure 4.10 shows this transfer and how S can receive a ticket for a replica of the unauthorized object. In this situation, the violation can be generalized by a professor saving the file to his local machine and then using the email link to send the file to a student. At this point, the student now has the capability through *save as* function to create a local copy of the file if desired.




Figure 4.9: Distribution threat.




Figure 4.10: Distribution threat using a copy.

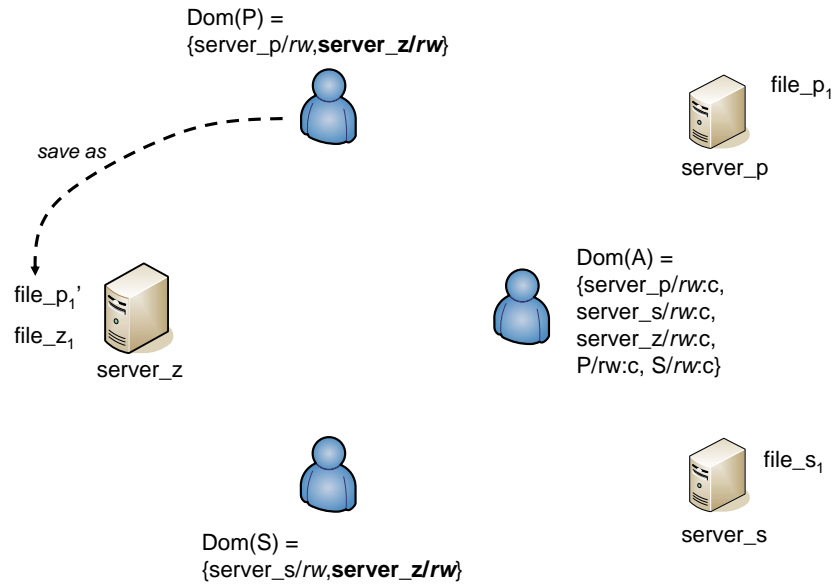


Figure 4.11: Distribution threat through association.

The final risk for Distribution occurs through association (cf., Definition 4c). Figure 4.11 shows the threat of Distribution for file_p1 by P. P can use *cc* and file_p1/*r* to *save as*. Instead of saving locally, however, P stores the new file to server_z by virtue of the ticket server_z/*w*. S can access the replica file, file_p1', by invoking server_z/*r* and S has now violated organization policy by obtaining access to a copy of an unauthorized file.

4.3.5 Maximal State. Examining each threat based on its initial state demonstrates how to analyze the system against various actions. The safety of the system, however, is determined by analyzing the events that can occur in a worse-case scenario. The maximal state is reached when all tickets that can be transferred have been transferred. The maximal state for an arbitrary system is not finite because more entities can be added through *can-create*. However, the safety question is decidable if the schema is acyclic and attenuating [32].

The first step in determining the maximal state is to identify how tickets can flow between subjects. In SPM-IT, this is accomplished via Distribution. Figure 4.12




Figure 4.12: Maximal State.

demonstrates a maximal state for the university setting being discussed. The results are intuitive and easily derived through Distribution (cf., Definition 4a), in which A can transfer every ticket to both P and S. This confirms the power of an administrator.

Once a maximal state has been determined, the system can be examined for other threats. S can perform Alteration on file_p1, file_s1, file_z1, and P. S can also perform Snooping on file_s1. The remaining threats for each subject can be obtained by applying the definitions for the actions. To mitigate these threats, a policy or implementation scheme that prevents the ticket transfers and reduces the risks associated with the other actions is needed.

4.3.6 Mitigating the Threats. This section demonstrates how mitigation policies can be modeled through SPM-IT to minimize the threat to a system.

Two-Person Integrity. Two-person integrity reduces the threat of Alteration and Snooping by prohibiting access to an object unless allowed by two authorized subjects. This makes it much more difficult for one subject to compromise

a file and can be implemented using a technique such as cryptographic secret sharing [35]. Cryptographic secret sharing encrypts a file and distributes unique keys to authorized subjects. To decrypt the file, a given number of subjects must present their key. If less than the required number of individuals present a key, it will not decrypt.

Thus, any attempt to invoke r or w requires the approval of another authorized subject. For the insider threat, this process imposes a check and balance system that forces a malicious insider to collude with another individual. For additional strength, the system could be set to randomly pick the second individual from the list of those authorized. This would force a malicious insider to potentially have to conspire with everyone that has access to the file.

Figure 4.13 is an implementation of the two-person integrity scheme. To specify the policy, additional subjects and rights have been introduced into the system:

1. $TS = \{\text{Prof, Stu, Admin, Verification Authority}\}$, $TO = \{\text{server, file}\}$
2. $R = \{r:c, w:c, x, q\}$
3. $cc(\text{Verification Authority}) = \emptyset$

$$\tau(V) = \text{Verification Authority}$$

The right x is execute and q is query. To obtain file_{z1}/r , P first must invoke V/x . The verification authority queries S for approval of the operation. If approved, V transfers file_{z1}/r to P. This reduces Snooping because another subject must approve the access. For w , if the operation is approved V will perform a write on behalf of the initiator. This prevents a subject with direct access from modifying the object. Alteration can still take place, but would require collusion.

Restricting File Copy. Although two-person integrity reduces the risk of Snooping and Alteration, a Distribution threat still exists. P can obtain




Figure 4.13: Two-Person Integrity Mitigation Scheme.

approval to access `file.z1` and then use `save as` to create a local copy. Based on the current policy, P can disseminate a ticket for the replica to an unauthorized entity. One way to prevent this is to alter `can-create` to eliminate the ability to create a file. In practice, this could be accomplished using display terminals that do not have the ability to save any information. For some the impact of this may be unacceptable. In a classified environment, however, the consequences of losing information may be so great that this is a viable solution.

Separation of Duty. It is evident that a system administrator poses a significant risk. To minimize the threat, a technique must be implemented to reduce the span of control. One such technique is separation of duty.

Figure 4.14 is an example of the system using separation of duty. In the new system, two additional administrators have been added. The scope of each administrator is now limited to a specific area. For example, A2 only possesses the ability to administer the part of the system associated with the professor. The rationale behind this process is to limit the amount of damage that can be caused by one entity to




Figure 4.14: Separation of Duty Mitigation Technique.

a smaller area. To enforce this policy filters must be added to restrict the ability of administrators to transfer tickets.

New Maximal State. Two-person integrity, restricting file copy, and separation of duty are mitigation strategies to reduce the threat to the system. The final policy with all of these changes is specified in Appendix A. Figure 4.15 shows this new policy. A1 has the capability to administer server_z and V1. The filters, however, only allow V1/x to be transferred to other subjects. Similarly, A2 can administer server_p, V2, and P. The transfer of tickets is also limited, only allowing V2/x to P. The transfer of tickets between P and S is allowed, however, this does not pose a threat because neither subject is capable of transferring an unauthorized right since they cannot gain access to an object with the copy flag and do not have can-create capability to perform a *save as* operation.

The policy can be examined through ticket transfer to determine the new maximal state. Based on the rules set by the new policy, the only change from the initial state is S has the capability to read server_z and server_s, P has the capability to read

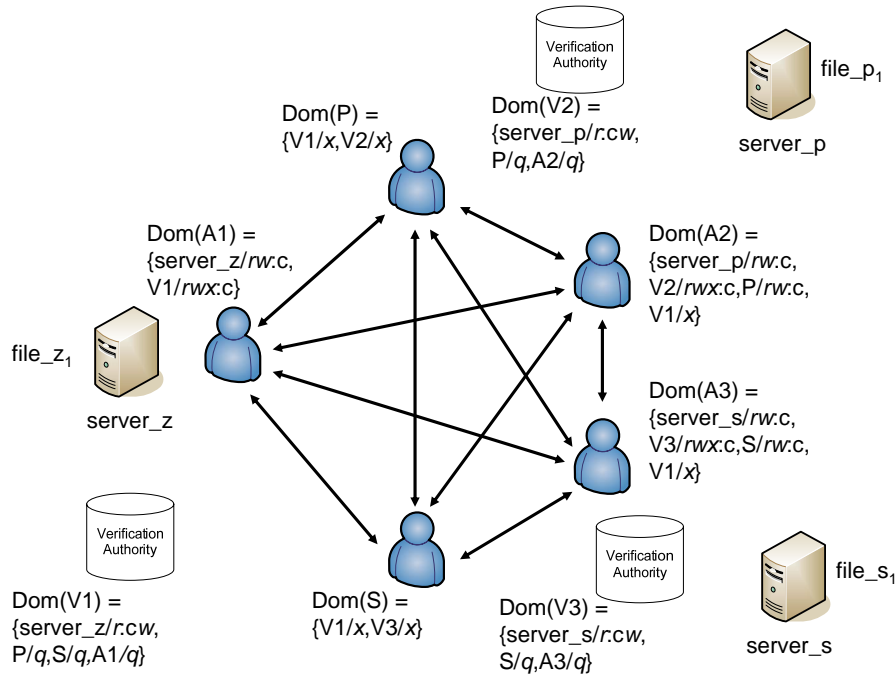


Figure 4.15: Implementation of Mitigation Techniques.

server_z and server_p, A2 has the capability to read server_z, and A3 has the capability to read server_z. Using the mitigation techniques, the implementation scheme closely mirrors the policy originally described. Professors only have access to files on the professor file server and common file server, and students only have access to files on the student file server and the common file server.

Further analysis of the scheme shows that some vulnerabilities are present in the system. The initial policy stated access should only occur if the information is required to perform their duty. A1 can still perform Snooping and Alteration on server_z and V1. The other administrators also have this ability in their respective areas. Even so, the threat has been compartmentalized and the risk to the system has been reduced.

4.4 Summary

This chapter introduces a formal methodology for modelling the insider threat. A taxonomy is developed that defines the threat through measurable and analyz-

able actions. The taxonomy is used to specify the threats to a system through the Schematic Protection Model for the Insider Threat. A process is demonstrated for analyzing a policy and implementation scheme using SPM-IT. Finally, the effectiveness of mitigation techniques are demonstrated through the attributes of the model.

V. Risk Analysis for Detecting Malicious Insiders

The previous chapter focuses on analyzing the safety of a system through a security model for the insider threat. This chapter examines the security of a system by developing a risk analysis framework to identify individuals that may pose a threat to the system. An attack cycle is first discussed to demonstrate the behavior and technical attributes that produce indicators. The risk analysis framework, a Multi-disciplinary Approach to Mitigating the Insider Threat (MAMIT), is presented and its effectiveness is demonstrated through the well-known case study involving Robert Hanssen.

5.1 Attack Cycle

To effectively develop mitigation techniques and prevention methods, it is important to first examine the attributes associated with the insider threat. The following four conditions are the characteristics identified as generally required before an individual betrays their organization and commits a malicious act [19]:

- 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

Two primary elements are required for a person to become a malicious insider: opportunity and motive. If either is missing, the individual does not pose a serious threat to the organization. For example, if an individual has high motivation to perform malicious behavior 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 to perform malicious behavior, 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,

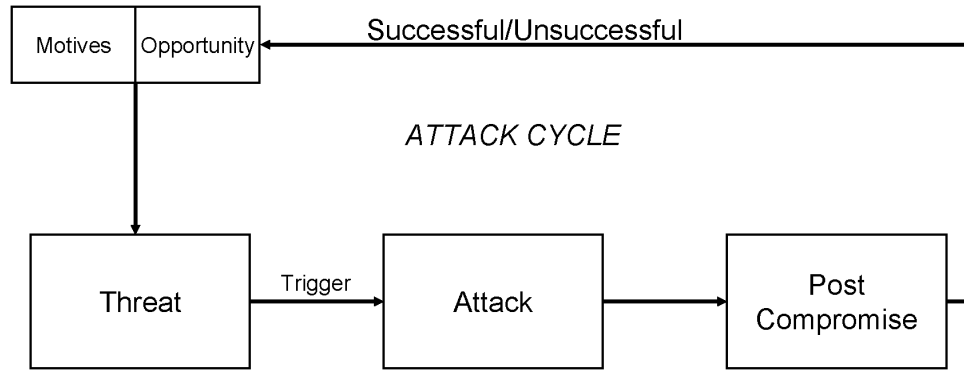


Figure 5.1: Attack cycle for the insider threat.

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 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.

Once the attack occurs, the individual evaluates how well the compromise went. This step forms the final phase of the attack cycle. Figure 5.1, the attack cycle for the insider threat, defines the anatomy of an attack. Each area is now discussed in more detail.

5.1.1 Opportunity. Opportunity for the insider can present itself through granted permissions, compromise of the system, or inadequate enforcement of organizational policies. The number 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 gains access through 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.

5.1.2 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 [18]. The findings identified motivating factors associated with the criminals, with the number one motivator being money. Also cited as motives were divided loyalties, a grudge against the employer, desire to please someone else, coercion, thrill seeking, and recognition.

Additionally, Dr. Mike Gelles of the Naval Criminal Investigative Service classified motivators for insiders into two personality disorders commonly found in spies: antisocial personality disorder and narcissism [15]. 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 excessive 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.

5.1.3 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 [19]. In those instances where a person performs a malicious act, their individual threat level has become so elevated that the inhibitions no longer prevent the attack from occurring. Upon reaching this

level, some event in their personal or professional life typically 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 [18]. 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 information theft.

5.1.4 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 is compromised and privileges are elevated. The effect of a 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.

5.1.5 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 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, sus-

picious travel plans, withdrawal from co-workers and changes in personal life. 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 Detection Systems, data mining event logs, and Honeypots. These techniques attempt to either determine when an attack occurs or make security controls so effective that a compromise is unfeasible.

Even with the advancement of technical countermeasures and law enforcement and counterintelligence strategies, current methodologies are unproven and have had limited success as trends indicate the insider threat is still a major concern. The current frameworks for risk analysis suggest independent methods that have failed to be effective because the insider threat is a problem that transcends functional areas. This research proposes a different mentality by leveraging the indicators from each area in a multidisciplinary approach and collaborates them in a cohesive manner that can be used to identify malicious insiders.

5.2 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 5.2 illustrates the proposed cohesive process for countering the malicious insider.

5.2.1 Likelihood Matrix. The MAMIT process requires a methodology to combine the different indicators to form one specific threat level that identifies an

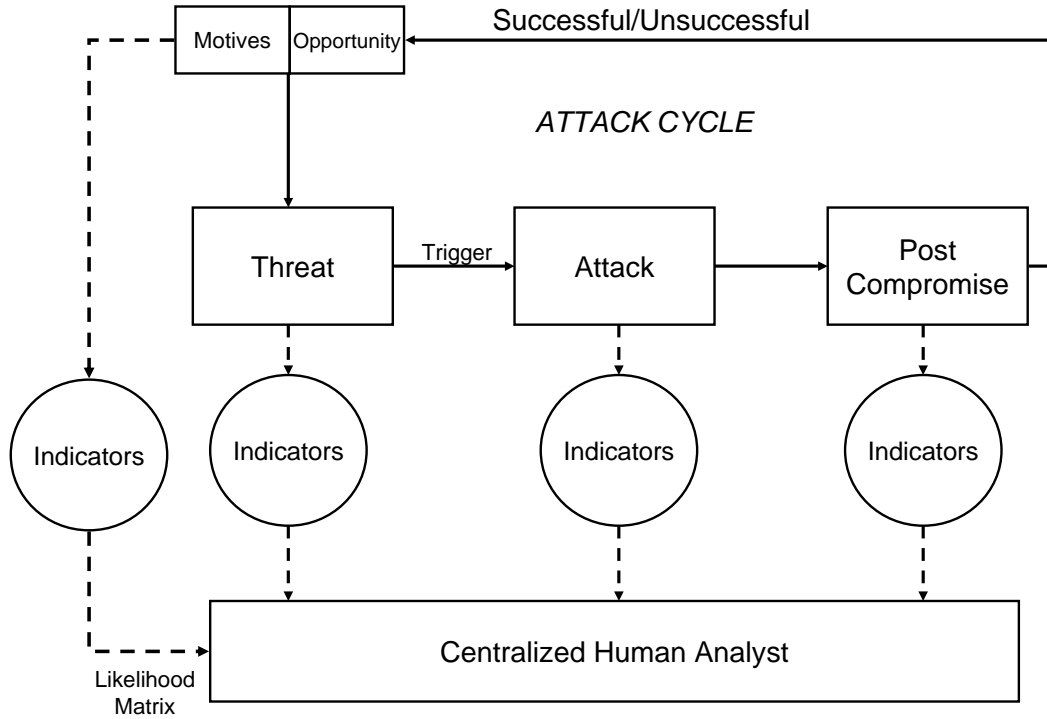


Figure 5.2: Multidisciplinary framework for mitigating the insider threat.

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 develops a method for quantifying the behavior and technical indicators.

Behavior Indicators. A study of past American spies determined that 80% exhibited one or more conditions of security concern defined in the Guidelines for Security Clearance [18]. 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 [38]. The adjudicative process assesses the risk associated with entrusting an individual with sensitive information. The 150 case studies in the PERSEREC database [12] were examined along with the thirteen specific guidelines outlined in the directive. The relevant areas that appear to indicate behavior trends for the malicious insider in these case studies are:

- 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. As such, these areas are used as the behavior indicators for the MAMIT framework. 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.

Technical Indicators. Technical indicators focus on specifying opportunities within the system. Quantifying these attributes is 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 measures 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, comprise ten factors used to determine an individual's threat level.

Putting it together. The first step in building the Likelihood Matrix is an initial analysis of an individual by assigning a numerical value for each of the ten factors. The values are 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. A prediction interval for the mean threat level of the organization is determined using the t -distribution. If an individual's threat level falls above the interval, they are identified as a potential insider threat. In statistical analysis the t -distribution is used to estimate the mean with an unknown population variance [11]. The individual threat levels are random variables because the characteristics of individuals will inevitably introduce some variance in the assigned values. When the random variables are averaged together the organization's mean threat level is produced. The prediction interval is used to determine if any single individual's threat level lies outside the norm for the organization's calculated threat level. The underlying goal of the statistical analysis is to characterize the normal threat level for an organization and use a method to identify the individuals that pose an unacceptable threat level. The following formula is used for determining the interval for the mean threat level:

$$\bar{X} \pm t_{\frac{\alpha}{2}} s \sqrt{1 + \frac{1}{n}} \quad (5.1)$$

where the prediction level is $100(1 - \alpha)\%$,

\bar{X} is the overall mean threat level of the organization,

t is the t variate at $\frac{\alpha}{2}$,

s is the standard deviation, and

n is the number of individuals within the organization.

The assigned prediction level determines the threshold for identifying threats. For example, a 90% prediction interval states that there is a 5% chance that an individual threat level will fall above the calculated range. Incidentally, there is also a 5% chance it will fall below the interval, but this is not a concern for identifying an insider threat. When determining the prediction level to choose, a larger percentage results in a greater interval. 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.

Updates. 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, 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. Involvement of supervisors is critical because they are typically in the best position to notice changes when they first occur [8].

5.2.2 Centralized Human Analyst. Ultimately, the purpose of MAMIT is to identify potential insider threats by using the different indicators. The effective implementation of this scheme requires 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 maintains two-way communication to each area involved in the process. The CHA should consist 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.

5.3 Case Study

The case involving Robert Hanssen is 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 [10, 12, 31].

5.3.1 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 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 U.S. 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 provided.

5.3.2 Indicators and Likelihood Matrix. For analysis of 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 strongly rejects 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 not being promoted late in his career. He also displayed strong financial indicators by living more extravagantly than his salary allowed. 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

Table 5.1: 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

information within the organization and was a recognized expert in computer systems and technology. Table 5.1 summarizes the values assigned for the indicators based on the relevant information from the case studies.

5.3.3 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 assumed. 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 (cf., (5.1)), 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 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.

5.4 MAMIT Implementation Scheme

A significant concern that arises from the MAMIT framework is the impact on organizational alignment. The implementation scheme of the MAMIT process calls into question how security fits into typical organizational structures. 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 [14]. 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.

The MAMIT framework demonstrates a need to shift 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.

5.5 Summary

This chapter introduces a formal framework for risk analysis of the insider threat. The MAMIT approach focuses on the combining of indicators using a multi-disciplinary approach producing a single identifier for risk analysis. The effectiveness of MAMIT is illustrated through the case study of Robert Hanssen which demonstrates the process would likely have identified him as an insider threat. Finally, the impact on organization alignment is discussed.

VI. Conclusions

This chapter presents the research conclusions, significance, and recommended areas for future research.

6.1 Problem Summary

Addressing the insider threat using systematic and formulated methodologies is an inherently difficult problem. The threat is typically viewed in an abstract manner without a method to represent the threat or means for formally identifying risks to the system. A security model provides a process to formally analyze the safety of a security policy and implementation scheme. Additionally, risk analysis formally identifies threats to the system by determining and measuring potentially dangerous indicators. Current techniques for risk analysis on the insider threat have not demonstrated an effectiveness through case studies or implementation. Formalizing the insider threat through security modelling and risk analysis provides a method to ensure sound policy implementation and identify individuals that pose a threat.

6.2 Conclusions of Research

The goals of this research were to develop a comprehensive security model that adequately determines the safety of a system against the insider threat and to present a risk analysis framework using a multidisciplinary approach capable of identifying potential malicious insiders.

6.2.1 Security Model. A formal security model, SPM-IT, analyzes the safety of a system against the insider threat. Initially a comprehensive taxonomy is developed using functional decomposition that characterizes the threat through definable and measurable actions. The actions are specified using SPM to create a comprehensive and formal security model for the insider threat. SPM-IT maintains acyclic and attenuating rules, thereby preserving the generality and tractable analysis of the original SPM. Through the maximal state, the safety of a security policy and implementation scheme for the insider threat can be determined. Mitigation techniques

can be implemented or changes in the policy can reduce the identified vulnerabilities within the system. The power and expressiveness of SPM-IT is demonstrated through a model of a simple policy and implementation scheme assuming a university setting.

6.2.2 Risk Analysis. The research introduces a formal framework for risk analysis of the insider threat. The MAMIT approach effectively identifies possible malicious insiders based on their threat level. The strategy focuses on the combining of information using a multidisciplinary approach producing a single identifier for risk analysis. Because the indicators are provided to a central analyst, individuals with an elevated threat level can be identified earlier and techniques enacted to mitigate the threat. Statistical analysis of individual threat levels against their organization identify the individuals that exhibit indicators consistent with malicious insiders. Since the framework analyzes threat levels against co-workers within the same organization, the methodology effectively adapts to different organizations and can be applied with the prediction interval set to the desired threshold. The effectiveness of MAMIT is illustrated through the case study of Robert Hanssen, demonstrating the process would likely have identified him as an insider threat.

6.3 Significance of Research

This research proposes the first known formal security model capable of analyzing the safety of a system against the insider threat. Included in the process is the development of a well-defined taxonomy that expresses the malicious insider through distinct actions capable of being decomposed and analyzed. Additionally, the research demonstrates there is no distinction between inert and control rights for the insider threat. This notion is a significant finding that demonstrates both types of rights can alter the protection state of a system and thus, both must be considered when modelling the insider threat.

The research also introduces a framework for risk analysis that is demonstrated through a well-documented case study. The MAMIT framework proves effective at

identifying a malicious insider through practical implementation. The multidisciplinary process also questions how security fits into typical organizational structures. The MAMIT process demonstrates a need to shift away from the standard organization structure to one having a central authority specifically for security matters.

The premises for much of this research was validated by independent scholars through the publishing of this work in the proceedings of two international security conferences. The process of defining a comprehensive and analyzable taxonomy for the insider threat using functional decomposition was presented in St. Petersburg, Russia at the 2005 Mathematical Methods, Models, and Architecture for Computer Network Security workshop [4]. Additionally, the MAMIT framework is to be presented in Maryland at the 2006 International Conference on Information Warfare and Security [5].

6.4 Recommendations for Future Research

The generality of SPM-IT means it can model more complex policies and systems. To further demonstrate the expressiveness and effectiveness of SPM-IT, the implementation of more complex policies and systems should be performed. Additionally, SPM-IT analysis can be automated using a simulation engine such as AFIT's Cyber Operations Emulator (CORE). The OPNET tool embedded in CORE allows specification of granular objects such as users, objects, and rights. Security policies and implementation schemes of significant magnitude can be specified, allowing the simulator to generate the ticket transfers for analysis. This method creates an automated process for identifying vulnerabilities and determining the safety of a system for a large class of systems. Another area for research is to determine whether the taxonomy developed in this research can be specified using the frameworks of other established security models such as the take-grant model. It would be interesting to compare the expressiveness of SPM-IT against these other models for the insider threat.

The MAMIT framework proved effective at identifying Robert Hanssen as a potential malicious insider. The case study, however, only demonstrates the effectiveness against an espionage threat to the United States. More research is required to determine if MAMIT is effective for other insider threat types. Additionally, the next logical step for analysis is implementing the MAMIT framework in a real-world environment. This would allow the determination of false-positive rates and the effectiveness and practicality of implementing the framework. Finally, the legal basis for collecting certain information for the MAMIT process must be examined.

Appendix A. Security Policy with Mitigation Strategies

This is the complete security policy for the SPM-IT university example with the mitigation strategies included.

1. $TS = \{\text{Prof, Stu, Admin, Verification Authority}\}$, $TO = \{\text{server, file}\}$
2. $R = \{r:c, w:c, x:c, q\}$
3. $\text{link}_1(P, S) = \text{true}$
 $\text{link}_2(P, A1) = \text{true}$
 $\text{link}_3(P, A2) = \text{true}$
 $\text{link}_4(P, A3) = \text{true}$
 $\text{link}_5(S, P) = \text{true}$
 $\text{link}_6(S, A1) = \text{true}$
 $\text{link}_7(S, A2) = \text{true}$
 $\text{link}_8(S, A3) = \text{true}$
 $\text{link}_9(A1, P) = \text{true}$
 $\text{link}_{10}(A1, S) = \text{true}$
 $\text{link}_{11}(A1, A2) = \text{true}$
 $\text{link}_{12}(A1, A3) = \text{true}$
 $\text{link}_{13}(A2, P) = \text{true}$
 $\text{link}_{14}(A2, S) = \text{true}$
 $\text{link}_{15}(A2, A1) = \text{true}$
 $\text{link}_{16}(A2, A3) = \text{true}$
 $\text{link}_{17}(A3, P) = \text{true}$
 $\text{link}_{18}(A3, S) = \text{true}$
 $\text{link}_{19}(A3, A1) = \text{true}$
 $\text{link}_{20}(A3, A2) = \text{true}$
 $\text{link}_{21}(V1, P) = \text{true}$
 $\text{link}_{22}(V1, S) = \text{true}$
 $\text{link}_{23}(V1, A2) = \text{true}$

$$\text{link}_{24}(V1, A3) = \text{true}$$

$$\text{link}_{25}(V2, P) = \text{true}$$

$$\text{link}_{26}(V3, S) = \text{true}$$

$$4. f_1(\text{Prof}, \text{Stu}) = T \times R$$

$$f_2(\text{Prof}, \text{Admin}) = \emptyset$$

$$f_3(\text{Prof}, \text{Admin}) = T \times R$$

$$f_4(\text{Prof}, \text{Admin}) = \emptyset$$

$$f_5(\text{Stu}, \text{Prof}) = T \times R$$

$$f_6(\text{Stu}, \text{Admin}) = \emptyset$$

$$f_7(\text{Stu}, \text{Admin}) = \emptyset$$

$$f_8(\text{Stu}, \text{Admin}) = T \times R$$

$$f_9(\text{Admin}, \text{Prof}) = \{\text{Verification Authority}/x\}$$

$$f_{10}(\text{Admin}, \text{Stu}) = \{\text{Verification Authority}/x\}$$

$$f_{11}(\text{Admin}, \text{Admin}) = \{\text{Verification Authority}/x\}$$

$$f_{12}(\text{Admin}, \text{Admin}) = \{\text{Verification Authority}/x\}$$

$$f_{13}(\text{Admin}, \text{Prof}) = \{\text{Verification Authority}/x\}$$

$$f_{14}(\text{Admin}, \text{Stu}) = \emptyset$$

$$f_{15}(\text{Admin}, \text{Admin}) = \emptyset$$

$$f_{16}(\text{Admin}, \text{Admin}) = \emptyset$$

$$f_{17}(\text{Admin}, \text{Prof}) = \emptyset$$

$$f_{18}(\text{Admin}, \text{Stu}) = \{\text{Verification Authority}/x\}$$

$$f_{19}(\text{Admin}, \text{Admin}) = \emptyset$$

$$f_{20}(\text{Admin}, \text{Admin}) = \emptyset$$

$$f_{21}(\text{Verification Authority}, \text{Prof}) = \{\text{server}/r\}$$

$$f_{22}(\text{Verification Authority}, \text{Stu}) = \{\text{server}/r\}$$

$$f_{23}(\text{Verification Authority}, \text{Admin}) = \{\text{server}/r\}$$

$$f_{24}(\text{Verification Authority}, \text{Admin}) = \{\text{server}/r\}$$

$$f_{25}(\text{Verification Authority}, \text{Prof}) = \{\text{server}/r\}$$

$$f_{26}(\text{Verification Authority}, \text{Stu}) = \{\text{server}/r\}$$

5. $cc(\text{Prof}) = \emptyset$

$$cc(\text{Stu}) = \emptyset$$

$$cc(\text{Verification Authority}) = \emptyset$$

$$cc(\text{Admin}) = \{\text{Admin}, \text{Prof}, \text{Stu}, \text{Verification Authority}\}$$

6. $cr(\text{Admin}, \text{Admin}) = \{\text{self}/rw:c\} \mid \{\text{Admin}/rw:c\}$

$$cr(\text{Admin}, \text{Prof}) = \{\text{Prof}/rw:c\} \mid \emptyset$$

$$cr(\text{Admin}, \text{Stu}) = \{\text{Stu}/rw:c\} \mid \emptyset$$

$$cr(\text{Admin}, \text{Verification Authority}) = \{\text{Verification Authority}/rwx:c\}$$

$$\tau(P) = \text{Prof}$$

$$\tau(S) = \text{Stu}$$

$$\tau(A1) = \tau(A2) = \tau(A3) = \text{Admin}$$

$$\tau(V1) = \tau(V2) = \tau(V3) = \text{Verification Authority}$$

$$\tau(\text{server}_p) = \tau(\text{server}_s) = \tau(\text{server}_z) = \text{server}$$

$$\tau(\text{file}_p1) = \tau(\text{file}_s1) = \tau(\text{file}_z1) = \text{file}$$

Bibliography

1. Anderson, R., T. Bozek, T. Longstaff, W. Meitzler, M. Skroch, and K. Van Wyk. "Research on Mitigating the Insider Threat to Information Systems - #2 Proceedings of the Insider Workshop". CF-163-DARPA, Arlington VA 2000.
2. Bingham, J. "Insider Threat Going Nowhere – or Maybe Everywhere". <http://www.intrusic.com/ThreatReport.htm>, 2004.
3. Bishop, Matt. *Computer Security: Art and Science*. Addison-Wesley, Boston MA, 2003.
4. Butts, J., R. Mills, and R. Baldwin. "Developing an Insider Threat Model Using Functional Decomposition". *Proceedings of the 2005 Mathematical Methods, Models, and Architecture for Computer Network Security workshop*, 412–417. St. Petersburg Russia, September 2005.
5. Butts, J., R. Mills, and G. Peterson. "A Multidiscipline Approach to Mitigating the Insider Threat". Accepted for presentation and publication in *The International Conference on Information Warfare and Security*, to be presented March 2006.
6. Caruso, V. *Outsourcing Informatoin Technology and the Insider Threat*. Master's thesis, Graduate School of Engineering, Air Force Institute of Technology (AETC), Wright-Patterson AFB OH, March 2003. AFIT/GIR/ENG/03-01.
7. Chinchani, R., A. Iyer, H. Ngo, and S. Upadhyaya. "Towards a Theory of Insider Threat Assessment". *Proceedings of the 2005 International Conference on Dependable Systems and Networks*, 108–117. Yokohama Japan, June 2005.
8. CSO Magazine, U.S. Secret Service and CERT Coordination Center. "2004 E-Crime Watch Survey". CSO Magazine, May 2005.
9. Daley, K., R. Larson, and J. Dawkins. "Automated generation and analysis of attack graphs". *Proceedings of the 2002 IEEE International Conference on Parallel Processing Workshops*, 49–64. Vancouver Canada, June 2002.
10. Davey, M. "Secret Passage". Chicago Tribune, 21 April 2002.
11. Devore, J. *Probability and Statistics for Engineers and Scientists*. Brooks/Cole, Belmont CA, 2004.
12. DSS. "ESPIONAGE CASES 1975-2004 (Defense Personnel Security Research Center)". <http://www.dss.mil/training/espionage>, December 2004.
13. Epstein, E. "The Diamond Invention". <http://edwardjayepstein.com/diamond/chap20.htm>, November 2005.
14. Fitzgerald, M. "All Over the Map". CSO Magazine, June 2003.

15. Gelles, M. "Exploring the Mind of a Spy". Defense Personnel Security Research Center, <http://www.dss.mil/search-dir/training/csg/security/Treason/Mind.htm>, 2001.
16. Gordon, L., M. Loeb, W. Lucyshyn, and R. Richardson. "2005 CSI/FBI Computer Crime and Security Survey". Computer Security Institute, 2005.
17. Harrison, M., W. Ruzzo, and J. Ullman. "Protection in operating systems". *Communications of the ACM*, 19(8):461–471, 1976.
18. Herbig, K. and M. Wiskoff. "Espionage Against the United States by American Citizens 1947-2001". Defense Personnel Security Research Center: Technical Report 02-5, 2002.
19. Heuer, R. "The Insider Espionage Threat". Defense Personnel Security Research Center, <http://www.dss.mil/search-dir/training/csg/security/Treason/Insider.htm>, 2001.
20. Jha, S., O. Sheyner, and J. Wing. "A Structural Framework for Modeling Multi-stage Network Attacks". *Proceedings of the 15th IEEE Computer Security Foundations Workshop*, 5–10. Cape Breton Nova Scotia, August 2002.
21. Keeney, M., E. Kowalski, D. Cappelli, A. Moore, T. Shimeall, and S. Rogers. "Insider Threat Study: Computer System Sabotage in Critical Infrastructure Sectors", May 2005. U.S. Secret Service and CERT Coordination Center/SEI.
22. Kratt, H. "The Inside Story:A Disgruntled Employee Gets His Revenge". SANS Institute, GIAC Certified Incident Handler, December 2004.
23. Landwehr, C. "Formal Models for Computer Security". *Computing Surveys*, 13(3):247–278, September 1981.
24. Lazarevic, A., L. Ertoz, V. Kumar, A. Ozgur, and J. Srivastava. "A Comparative Study of Anomaly Detection Schemes in Network Intrusion Detection". *Proceedings of the 2003 SIAM International Conference on Data Mining*, 108–120. San Francisco CA, May 2003.
25. Magklaras, G. and S. Furnell. "Insider Threat Prediction Tool: Evaluating the probability of IT misuse". *Computers & Security*, 21(1):62–73, 2002.
26. Maybury, M., P. Chase, B. Cheikes, D. Brackney, S. Matzner, T. Hetherington, B. Wood, C. Sibley, J. Marin, T. Longstaff, L. Spitzner, J. Haile, J. Copeland, and S. Lewandowski. "Analysis and Detection of Malicious Insiders", 2005. Submitted to 2005 International Conference on Intelligence Analysis, McLean VA.
27. Phillips, C. and L. Swiler. "A graph-based system for network-vulnerability analysis". *Proceedings of the 1998 workshop on New security paradigms*, 71–79. Charlottesville VA, September 1998.
28. Rammel, A. *Assessing the Usefulness of Visualization Tools to Investigate Hidden Patterns Within Insider Attack Cases*. Master's thesis, Graduate School of

Engineering, Air Force Institute of Technology (AETC), Wright-Patterson AFB OH, March 2005. AFIT/GIR/ENV/05M-14.

29. Randazzo, M., M. Keeney, E. Kowalski, D. Cappelli, and A. Moore. "Insider Threat Study: Illicit Cyber Activity in the Banking and Finance Sector", August 2004. U.S. Secret Service and CERT Coordination Center/SEI.
30. Reiher, P. "File Profiling for Insider Threats". Air Force Research Laboratory Technical Report: AFRL-SN-WP-TR-2002-1102, February 2002.
31. Rodriguez, P. "Diary of a Spy". Insight Magazine, 16 July 2001.
32. Sandhu, R. "The Schematic Protection Model: Its Definition and Analysis for Acyclic Attenuating Schemes". *Journal of the Association for Computing Machinery*, 35(2):404–432, April 1988.
33. Schneier, Bruce. *Secrets and Lies*. Wiley Publishing, Danvers MA, 2000.
34. Schultz, E. "A Framework for Understanding and Predicting Insider Attacks". *Computer & Security*, 21(6):525–531, October 2002.
35. Shamir, A. "How to Share a Secret". *Communications of the ACM*, 22(11):612–613, November 1979.
36. Shaw, E., K. Ruby, and J. Post. "The Insider Threat to Information Systems". *Security Awareness Bulletin*, 27–46, June, 2002.
37. Sheyner, O., J. Haines, S. Jha, R. Lippman, and J. Wing. "Automated generation and analysis of attack graphs". *Proceedings of the 2002 IEEE Symposium on Security and Privacy*, 273–284. May 2002.
38. U.S. Policy. "Adjudicative Guidelines for Determining Eligibility for Access to Classified Information". Defense Personnel Security Research Center, <http://www.fas.org/sgp/spb/class.htm>, Approved by the President 24 March 1997.
39. Varadharajan, V. "Extending the Schematic Protection Model II: Revocation". *ACM SIGOPS Operating Systems Review*, 31(1):64–77, January 1997.
40. Varadharajan, V. and C. Calvelli. "Extending the Schematic Protection Model I: Conditional Tickets and Authentication". *Proceedings of 1994 IEEE Computer Society Symposium on Research in Security and Privacy*, 213–229. Oakland CA, May 1994.
41. Verton, D. "Spy Case Demos Insider Threat". *Computer World*, February 2001.
42. Viega, J. and G. McGraw. *Building Secure Software*. Addison-Wesley, Boston MA, 2002.
43. Wood, B. "An Insider Threat Model for Adversary Simulation". SRI International Cyber Defense Research Center, Albuquerque NM, July 2000.

REPORT DOCUMENTATION PAGE

*Form Approved
OMB No. 074-0188*

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of the collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) 23-03-2006		2. REPORT TYPE Master's Thesis		3. DATES COVERED (From - To) Feb 2005-Mar 2006	
4. TITLE AND SUBTITLE Formal Mitigation Strategies for the Insider Threat: A Security Model and Risk Analysis Framework				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Butts, Jonathan W., First Lieutenant, USAF				5d. PROJECT NUMBER If funded, enter ENR #	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAMES(S) AND ADDRESS(S) Air Force Institute of Technology Graduate School of Engineering and Management (AFIT/EN) 2950 Hobson Way WPAFB OH 45433-7765 937-255-3636				8. PERFORMING ORGANIZATION REPORT NUMBER AFIT/GE/ENG/06-57	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT This research extends the Schematic Protection Model to produce the first comprehensive security model capable of analyzing the safety of a system against the insider threat. The model is used to determine vulnerabilities in security policies and system implementation. Through analysis, mitigation strategies that effectively reduce the threat are identified. Furthermore, an action-based taxonomy that expresses the insider threat through measurable and definable actions is presented. A risk analysis framework is also developed that identifies individuals within an organization that display characteristics indicative of a malicious insider. The framework combines behavior and technical attributes to produce a single threat level for each individual within the organization. Statistical analyses using the t-distribution and prediction interval on the threat levels reveal those individuals that are a potential threat to the organization. The effectiveness of the framework is illustrated using the case study of Robert Hanssen, demonstrating the process would likely have identified him as an insider threat.					
15. SUBJECT TERMS Insider threat, security model, risk analysis, information security					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 92	19a. NAME OF RESPONSIBLE PERSON Robert F. Mills, PhD (ENG)
REPORT U	ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (Include area code) (937) 255-6565 ext 4527; email: robert.mills@afit.edu

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39-18