

Security Attribute Evaluation Method

Shawn A. Butler

May 2003

CMU-CS-03-132

School of Computer Science

Carnegie Mellon University

Pittsburgh, PA 15213

Thesis Committee

Mary Shaw (chair)

Bill Scherlis

Jeannette Wing

Paul Fischbeck

*Submitted in partial fulfillment of the requirements for the Degree of
Doctor of Philosophy.*

Copyright © 2003 Shawn A. Butler

This research was sponsored in part by a grant from the National Science Foundation (NSF). Any opinions, findings and conclusions or recommendations expressed in this publication are those of the author and do not necessarily reflect those of the sponsor or other entity.

Keywords: Security, Cost-benefit, Multi-attribute, Risk Management, Security Architecture

Abstract

A security manager's selection of risk-mitigation controls for an information system's security architecture depends on the organization's risk-management process. Current security risk-management processes require security managers to thoroughly analyze their organization's threats, vulnerabilities, and assets before selecting cost-effective risk-mitigation controls. The most common risk-management method, Annualized Loss Expectancy (ALE), expects security managers to assess the probabilistic damage from different types of attacks, investing only in those risk-mitigation controls that cost less than the anticipated loss in asset value. The problem with current risk-mitigation-control cost-benefit analysis methods is that they attempt to give security managers the ability to make precise security investment recommendations or decisions based on imprecise information, such as estimated probabilities or expected economic loss in asset value.

This thesis proposes the Security Attribute Evaluation Method (SAEM) as an alternative to current risk-mitigation-control cost-benefit analysis methods. SAEM uses multi-attribute decision analysis techniques from the field of Decision Sciences to guide a security manager in his or her selection of risk-mitigation controls for the organization's information system security architecture. In contrast with current cost-benefit analysis methods, SAEM focuses on the *relative* benefit of risk-mitigation controls rather than the economic net value of the information system with and without the risk-mitigation control. In addition, SAEM integrates a new coverage-analysis model that allows security managers to evaluate how a risk-mitigation control contributes to the security architecture's defense-in-depth design, a fundamental security engineering design principle.

In this thesis, I present the results of using SAEM with the security managers of three different organizations—a large commercial company, a large government organization, and a small hospital. SAEM provided these security managers with insight into their risk priorities and, in two organizations, SAEM highlighted weaknesses in their security architectures. Overall, the security managers felt that SAEM's coverage-analysis model was very helpful in assessing how risk-mitigation controls support the organization's defense-in-depth security strategy.

ACKNOWLEDGMENTS

This has been quite a journey for me, but I never would have made it without the support of those that believed that I could finish. This is for Mary, who always believed that I could do a thesis even when I didn't. I also want to thank the Software Engineering Institute for their unqualified support and expert knowledge about security. I also want to thank my father, who offered many times to write this document for me, but knew that I had to do it myself. Finally, I want to thank Alice, who patiently and expeditiously read all the chapters. Thank You!

TABLE OF CONTENTS

| | | |
|------------|--|----|
| CHAPTER 1. | Introduction | 1 |
| 1.1 | Introduction..... | 1 |
| 1.2 | Background..... | 1 |
| 1.3 | Security Attribute Evaluation Method..... | 4 |
| 1.4 | Case Studies as Validation..... | 6 |
| 1.5 | Future Work | 9 |
| 1.6 | Thesis Roadmap..... | 9 |
| CHAPTER 2. | Background | 11 |
| 2.1 | Introduction..... | 11 |
| 2.2 | Risk Management: State of the Practice | 13 |
| 2.3 | Current Research..... | 19 |
| 2.4 | Security Attribute Evaluation Method (SAEM) | 21 |
| CHAPTER 3. | Multi-attribute Analysis Essentials..... | 23 |
| 3.1 | Introduction..... | 23 |
| 3.2 | Multi-attribute Analysis Applicability..... | 23 |
| 3.3 | Multi-attribute Analysis | 27 |
| 3.4 | Summary | 33 |
| CHAPTER 4. | Security Attribute Evaluation Method Process | 35 |
| 4.1 | Introduction..... | 35 |
| 4.2 | Security Architecture Development..... | 35 |
| 4.3 | The Security Attribute Evaluation Method..... | 36 |
| 4.4 | Risk Assessments | 38 |
| 4.5 | Benefit Analysis..... | 43 |
| 4.6 | Coverage Analysis | 46 |
| 4.7 | Security Technology Tradeoff Analysis..... | 49 |
| 4.8 | Summary | 50 |
| CHAPTER 5. | Commercial Case Study | 51 |
| 5.1 | Introduction..... | 51 |
| 5.2 | Case Study Description | 51 |
| 5.3 | The Risk Assessment..... | 52 |
| 5.4 | Final SAEM Risk Assessment | 58 |
| 5.5 | Benefit Analysis..... | 60 |
| 5.6 | Coverage Evaluation..... | 63 |
| 5.7 | Security Technology Tradeoff Analysis..... | 65 |
| 5.8 | Summary | 67 |
| CHAPTER 6. | Hospital Case Study..... | 81 |
| 6.1 | Introduction..... | 81 |
| 6.2 | Case Study Description | 81 |
| 6.3 | The Risk Assessment..... | 82 |
| 6.4 | Benefit Analysis..... | 88 |
| 6.5 | Coverage Analysis | 91 |
| 6.6 | Security Technology Tradeoff Analysis..... | 92 |
| 6.7 | Summary | 94 |

| | |
|--|-----|
| CHAPTER 7. Government Case Study..... | 101 |
| 7.1 Introduction..... | 101 |
| 7.2 Case Study Description..... | 101 |
| 7.3 The Risk Assessment..... | 102 |
| 7.4 Benefit Analysis..... | 113 |
| 7.5 Coverage Analysis..... | 115 |
| 7.6 Security Technology Tradeoff Analysis..... | 116 |
| 7.7 Summary..... | 118 |
| CHAPTER 8. Analysis of the Method..... | 136 |
| 8.1 Introduction..... | 136 |
| 8.2 Risk Assessment Analysis..... | 137 |
| 8.3 Benefit Analysis..... | 145 |
| 8.4 Coverage Analysis..... | 152 |
| 8.5 Security Technology Tradeoff Analysis..... | 153 |
| CHAPTER 9. Future Work and Observations..... | 157 |
| 9.1 Introduction..... | 157 |
| 9.2 Observations..... | 157 |
| 9.3 Future Work..... | 162 |
| APPENDIX A..... | 165 |
| REFERENCES..... | 173 |

CHAPTER 1. Introduction

1.1 Introduction

This thesis describes the Security Attribute Evaluation Method (SAEM), a quantitative cost-benefit analysis technique, developed to help guide security architecture design decisions. SAEM uses multi-attribute analysis techniques to prioritize an organization's risks and risk-mitigation controls and uses a new coverage model to evaluate a risk-mitigation control with respect to an organization's security architecture. The purpose of this thesis is to show the feasibility of using multi-attribute analysis techniques in guiding security architecture design decisions. In addition, this thesis provides a preliminary evaluation of the usefulness of the new coverage model. This model helps guide security design decisions because it gives security practitioners, i.e., engineers and managers, the ability to compare different security technologies with respect to the organization's security goals of protection, detection, and recovery.

Three real-world case studies validate the feasibility of using multi-attribute analysis techniques and evaluate the coverage model. The thesis shows that two case-study security managers added detection mechanisms to their organizations' security architectures based on SAEM. Furthermore, all three case-study security managers found the coverage model very useful in evaluating their information systems' security architectures. Although this thesis addresses the use of multi-attribute decision analysis techniques in guiding security architecture design decisions, the success from this research should encourage further researcher in using multi-attribute analysis techniques to guide other architectural attribute design decisions, such as dependability or maintainability.

1.2 Background

Security engineers and managers select risk-mitigation controls, in part, using a security risk-management process, but information system executives are often skeptical of the results of this process. They are skeptical because qualitative risk-management processes do not provide sufficient information to make finer grained decisions, and while quantitative processes offer the ability to make finer grained decisions, the quantitative results are based on highly subjective estimates of an organization's risk environment. Furthermore, quantitative processes require that security managers estimate the economic value of the organization's assets at risk. Converting non-tangible assets to dollars adds to the subjectivity and skepticism of quantitative results. SAEM addresses these problems because it provides a framework so that security managers and engineers can compare the relative value that a risk-mitigation control provides rather than trying to estimate the specific economic value.

1.2.1 Risk Management and Security

Risk management is the process of identifying risks, assessing risk-mitigation controls, and taking steps to reduce risk to an acceptable level (Stoneburner, Goguen et al. 2001). Information-system managers can use risk-management processes to address system architecture attributes, such as reliability and dependability. With respect to the security of

information systems, risk management is the process of: 1) identifying information system critical assets, threats, and system vulnerabilities, 2) estimating the expected impact from a successful attack on a critical asset, 3) assessing the expected effect that a risk-mitigation control has in reducing the impact of a successful attack, and 4) selecting the risk-mitigation controls that best meet the goals and objectives of the organization (King, Dalton et al. 2001).

Security engineers and managers are faced with limited budgets so they want to maximize the security architectures value; however, they have limited knowledge about their threat environment. They must make decisions with some degree of confidence that they are the correct ones and convince superiors these decisions are rational and will deliver the intangible “return on investment”, a standard against which other managers are evaluated. All other decision factors being equal, security managers want to select the most cost-effective risk-mitigation controls for their information system security architecture.

1.2.1.1 Qualitative and Quantitative Risk Management Methods

Currently, security managers can use qualitative or quantitative risk-management methods to help them select risk-mitigation controls for their organization’s security architecture. Currently, the National Institute of Standards and Technology recommends that government agencies use their qualitative risk-management process. In addition, the Software Engineering Institute developed OCTAVEsm, a qualitative risk-management process that helps security managers identify their threats and vulnerabilities. ALE, a quantitative risk-management process, is widely known among certified security practitioners in industry. Qualitative and quantitative methods are similar in the following ways:

The methods require the information system’s security manager (or a risk-analysis team) to identify the organization’s threats, vulnerabilities, critical information system assets, and the impact of attacks on the critical assets.

- The methods rely on the security manager’s best available information and expertise to estimate the organization’s level of risk for each asset from a threat exploiting a known vulnerability.
- The methods use the expected frequency of an attack and the expected outcomes of an attack to assess the level of asset-risk from a threat.
- The methods assess the value of a risk-mitigation control based on how well it reduces the level of threat-vulnerability risk.

Unfortunately, engineers and managers cannot consult canonical or industry-specific databases that might help determine their risks and the effectiveness of their risk-mitigation controls. Managers and engineers rely on empirical evidence and their experience when assessing their organizations’ threats and capabilities. Consistent with existing methods, SAEM captures their assumptions about the threat environment and the effectiveness of security technologies. This helps security managers and engineers to communicate the basis for their security technology selection decisions. When, or if, security managers gain better information, they can re-evaluate the impact of these changes to their underlying assumptions.

1.2.1.2 Qualitative Methods

Using qualitative risk-management methods, security managers assess the likelihood of an attack and its subsequent impact in general terms—high, medium, or low—rather than estimate specific probabilities to these events. However, in the National Institute of Standards and Technology’s risk management guide qualitative method, security managers assign probabilities (0 – 1) to the high-medium-low rankings and scale the impact of an attack, from 0 to 100, to determine a threat-vulnerability risk level. Despite this brief mapping of qualitative ratings to quantitative values, the result of the NIST risk-management method is a qualitative assessment of high, medium, or low risk for each asset. Furthermore, although NIST’s risk-management method recognizes the need for quantitative cost-benefit analyses, the value of the risk-mitigation control is rather simplistically based on a monetary assessment of an asset’s risk minus the cost of the risk-mitigation control. For example “the organization may not want to spend \$1,000 on a control to reduce a \$200 risk.”

1.2.1.3 Quantitative Methods

In contrast to qualitative risk-management methods, quantitative methods require the security manager to estimate: 1) the yearly probability of a vulnerability-threat event, i.e., an attack and 2) the expected economic damage incurred from the attack. In the most common risk-quantification method, Annualized Loss Expectancy, or ALE, these estimates are used to assess the asset-risk for a vulnerability-threat pair. In ALE, cost-benefit analysis of a risk-mitigation control is based on the difference between the ALE with and without the control, minus the cost of the control. ALE advocates recommend that a security manager only invest in risk-mitigation controls that provide a *positive net benefit* based on the cost-benefit analysis. In fact, recent research suggests that security managers should not invest more than 37% of the asset’s value in mitigating a risk (Gordon and Loeb 2002).

1.2.1.4 What is wrong with ALE?

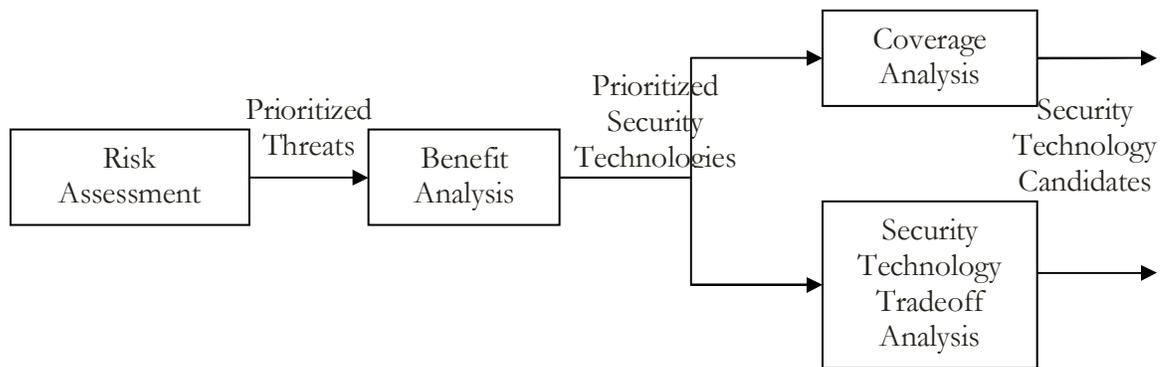
The fundamental problem with ALE’s security risk-management cost-benefit analysis is that it attempts to give security managers the ability to make *precise* security investment decisions based on a series of “best estimates” (e.g., estimated probability, expected damage, or high/medium/low ratings). The security manager’s uncertainty about the probability of attack and the uncertainty about the consequences of the attack appear to amplify the uncertainty about the risk overall. Therefore, decision maker’s lack confidence in arguments advocating “an investment of x dollars should result in a reasonable balance between security technology costs and benefits for the organization.” As one manager stated “it’s like making a decision on a house of cards”.

In addition, the ALE cost-benefit analysis requires a monetary assessment of all attack consequences, which contributes to the specious results of quantitative risk-management methods. Information system decision-makers may find it difficult to determine the economic loss of an attack that damages their public image, affects safety, or impacts worker productivity. These are the consequences from an attack that may really matter to an organization. The significance of these consequences may get diluted when risk analysts attempt to determine the dollar value of non-tangible assets.

1.3 Security Attribute Evaluation Method

SAEM provides a framework which helps security managers select risk-mitigation controls in a systematic manner and that establishes a pragmatic evaluation of risks given the real-world imprecision of threat estimations. SAEM helps security managers: 1) prioritize their organization's threats, 2) use the prioritized threats to assess the *relative* value of risk-mitigation controls, 3) evaluate how different risk-mitigation controls fit into their information system security architecture with respect to protection, detection, and recovery from threats using a new coverage model, and 4) select risk-mitigation controls using a multi-attribute analysis technique that compares alternative risk-mitigation controls based on non-technical purchasing objectives, such as user friendliness and maintainability, and implementation and maintenance costs. Figure 1-1 shows how the steps are related.

Figure 1 - 1 SAEM Steps



1.3.1 SAEM Steps

SAEM consists of four steps: 1) a risk assessment, 2) a security technology benefit analysis, 3) a coverage analysis, and 4) a security technology tradeoff analysis. These four steps help lead an organization's security manager in selecting the security components or risk-mitigation controls for the organization's security architecture. SAEM relies on the security manager's ability to identify organization's threats and most likely outcomes from an attack. Therefore, SAEM is similar to existing risk-management methods except that SAEM does not require security managers to identify system vulnerabilities. In fact, SAEM could use the threats identified during the organization's risk-analysis as input to the SAEM risk-assessment step.

The SAEM risk-assessment phase prioritizes the threats in terms of their risk to the organization. SAEM characterizes threats by their Relative Threat Index, which is a non-dimensional unit indication of their relative risk to the organization. This is a departure from existing methods in that the SAEM risk assessment develops the potential, but relative, cost that a threat could have against the organization's assets, without assessing the monetary value of the assets.

The benefit analysis step uses the security manager's attack frequency and outcome estimates to determine the overall benefit that a risk-mitigation control would have. Again, the risk-mitigation control is ranked on its relative effectiveness when compared to other risk-mitigation controls. In this way, SAEM differs from current approaches because the value of the risk-mitigation control is not defined in economic terms, but its relative value to other risk-mitigation controls.

In the third step, coverage analysis, the goal is to provide insight and design guidance to security managers when comparing or selecting risk-mitigation controls, so SAEM provides a coverage-analysis model. This model enables security managers to simultaneously see how risk-mitigation controls compare with respect to defense-in-depth (prevention, detection, and recovery) and breadth-of-defense coverage against multiple threats. No other risk-management method provides this perspective.

Although the benefit analysis phase determines which security technologies are the most effective in mitigating threats, security managers consider purchase cost, maintenance, skill level requirements, false positives, etc. before selecting a technology for inclusion in the security architecture. Security tradeoff analysis helps the security managers compare security technologies using multi-attribute analysis techniques to rank each security technology according to the organization's decision objectives. Rather than assigning a precise dollar amount, the SAEM method expresses relative value and makes clear the best investments for the organization's objectives.

1.3.2 Participants

Although the participants in each phase of SAEM vary among organizations, a multi-attribute analyst and an organization's lead security manager or specialist are the key participants. A multi-attribute analyst facilitates each phase of the process, eliciting from the lead security specialist and other participants their knowledge about the organization's risks, their expertise about the effectiveness of security technologies, and the key factors that the organizations' managers consider in selecting security technologies.

1.3.3 Summary of Contributions

This thesis offers two contributions to the state-of-the practice. First, this thesis shows the feasibility of using decision analysis techniques to guide security managers in selecting security technologies. Second, it shows that the coverage model is useful in helping security manager's select security technologies based on the organization's defense-in-depth and breadth-of-coverage objectives. In this thesis, the correlation of the final results with the case-study security manager's final results demonstrates the feasibility of using decision analysis techniques. The highly positive comments and scores on the case-study satisfaction surveys demonstrate the usefulness of the coverage model. Finally, the commercial and hospital case-study security managers selected security technologies based on the results of the coverage analysis, which provides some initial evidence that SAEM produces credible results on which security managers can make security-architecture design decisions.

As with current risk-mitigation methods, SAEM bases threat risk-mitigation control rankings on the estimated frequency and outcomes of attacks from threats. However, in contrast to existing risk-mitigation methods, SAEM helps the security manager find the most

effective risk-mitigation controls—specifically, technological risk-mitigation controls—without the need to assess financial risk of information system assets.

1.4 Case Studies as Validation

In this thesis, I used three case studies to validate the feasibility of SAEM in different organizations. I chose three different types of organizations—a large commercial company, a large civilian government organization, and a small hospital—so I could evaluate SAEM’s usefulness to security managers with different levels of experience and to organizations with different security requirements and resource-allocation constraints. These case studies helped determine the efficacy of SAEM by:

- establishing the level of effort and expertise required to use SAEM
- establishing whether security managers found the process useful and/or insightful
- evaluating the performance of the additive-model assumptions against actual threat and effectiveness estimates

Unfortunately, none of the organizations had completed a risk assessment prior to using SAEM, so I provided an initial set of threats as a starting point for discussion.¹

Before starting each risk assessment I asked the security managers to rank the organization’s threats. After completing the first iteration of the SAEM risk-assessment step, the security managers had the opportunity to adjust their estimates or their initial rankings. At the conclusion of the SAEM risk assessment, as part of the validation process, I compared final risk-assessment results with the security manager’s initial risk-assessment threat rankings to determine whether the risk-assessment process influenced the security manager’s final threat ranking.

In two case studies, the degree of correlation between the final SAEM rankings and the final security managers’ rankings were highly correlated, indicating that the SAEM risk assessment can approximate the security manager’s threat prioritizations. In addition, a comparison between the correlations of the first iteration rankings and the final threat rankings showed that the process significantly influenced the security manager’s final threat ranking. However, since SAEM uses the security manager’s estimation of threats and outcomes, one cannot determine whether the final results are *better* than before the SAEM process.

Although the results from these three case studies do not constitute a statistically significant sample size, the results show the feasibility of the method, especially for security managers with little or moderate experience. In addition, each case study security manager completed a satisfaction survey, which enabled them to assess the usefulness, ease, and general satisfaction with SAEM. The satisfaction survey was especially helpful in assessing the usefulness of the coverage-analysis phase. The satisfaction ratings were high among all of the case studies, particularly with the coverage model.

¹ Ideally, the risk assessment step within SAEM would use threats identified during a prior risk assessment.

1.4.1 Pre-validation Case Studies

Before I validated SAEM using three case studies, I developed and refined the method using three different case studies. During my initial development of SAEM, I tried to elicit expert threat information and effectiveness estimates that I could use to compare against the case-study manager estimates. In addition, expert testimony or evidence about the effectiveness of security technologies would have helped in cases where the case-study security manager was not familiar with a technology. Experts could not provide the threat and effectiveness estimate baselines because they felt that there were too many organization-specific factors that determine the actual effectiveness of a security technology. Despite the lack of threat and security technology effectiveness baselines, the pre-validation case studies were very useful because they established the elicitation protocol that I used during the validation case studies and showed that security managers were willing to make estimates about their perceived effectiveness of security technologies within their organizations.

1.4.2 Validation Case Studies

1.4.2.1 Commercial Case Study

I completed the first case study, a commercial manufacturer of retail products, in approximately two weeks, using a series of interviews to elicit threat and frequency estimates. As the analyst, I observed that the case-study participants were familiar with most of the threats and security technologies I presented, although they did not have much direct experience with many of them. In addition, although the company was large, the overall security budget was limited.

Overall, the security manager reported that the analysis was insightful and helpful in developing the organization's security strategies. The risk assessment highlighted the organization's conflict between an open and trusting work environment and the risk of being too lenient with security policies. The benefit analysis showed that some technologies that could help reduce risk from virus attacks had been overlooked because they would cause the organization to enforce stricter security policies. The coverage evaluation showed the security architecture was weak in detection mechanisms. Finally, the tradeoff analysis showed mixed results because the process assumes none of the security technologies that the security manager selected for comparison existed in the organization's security architecture.

1.4.2.2 Hospital Case Study

The second case study that I present is from a small hospital. The primary security responsibilities fall to the Technical Director, who handles the day-to-day operation of the information system. Of the three case studies presented, this organization's participants had the least experience and knowledge about risk-mitigation controls and threats. Since the hospital is small, security resources are very constrained, but Health Insurance Portability and Accountability Act (HIPAA) guidelines are putting pressure on the information-system executives to reduce security risks. In this case study, I used a combination of interviews and surveys to elicit the risk information and the security-technology effectiveness estimates over a two-month period, but the time required to identify and analyze the estimation data was about two person-weeks.

Overall, the Technical Director was very pleased with the results. Although the participants were not as familiar with the threats and security technologies as were other case study participants, the Director reported that:

“The time we spent made me aware of attacks/security holes that I was not aware of. Your assessment prepared me for the Red Siren assessment and HIPAA. Before you showed I was behind as far as network security goes, and now I have a plan created that will get the hospital [sic] where we need to be for HIPAA”.

A review of the SAEM results showed that the SAEM risk-assessment process moderately influenced the Director’s final prioritization of threats, and the satisfaction survey indicated that the Director gained significant insight about the value of some security technologies. In addition, he felt the coverage analysis would make it much easier to explain why a particular security technology should be purchased. Finally, although the security-tradeoff analysis ultimately showed a high positive correlation, the case study showed that it would have been more useful if the security-tradeoff analysis were used to evaluate security technologies in the context of organization’s security architecture, rather than independent of the security architecture.

1.4.2.3 Government Case Study

The third case study is from large government civilian organization. I chose this organization because the security staff was large, relative to the other case studies, and very experienced. Due to the highly sensitive nature of the information processed in the information system, the organization was fortunate to have a very large security budget. There were few security technologies that were not considered necessary and affordable. In addition, the organization operated a computer incident response team so they collected frequency data about many of the threats.

Despite the fact that SAEM did not appear to significantly influence the manager’s prioritization of threats, the participants found the results insightful. The manager’s prioritization of threats reflected his present concerns for the organization, and SAEM may have been more useful if the security manager could have used it to evaluate the information system’s residual risks, i.e., the remaining information system risks given the organization’s security architecture.

The benefit analysis resulted in high rankings for security technologies that mitigated threats that the risk assessment had determined were important to the organization. The organization had purchased all of the security technologies for those threats so the benefit analysis validated their previous selections. Finally, although the manager did not comment on the value of the coverage analysis, his supervisor felt that the coverage analysis was very important for communicating the organization’s needs.

1.4.3 Validation Limitations

1.4.3.1 Catastrophic Threats

Since SAEM uses the threats identified during the risk-assessment process, new threats are easily integrated into the method. Catastrophic events, such as terrorist attacks or major natural disasters, could be included in this method, just as they are in the other risk- management

methods, but their extremely low probabilities of occurrences would appropriately make them a low priority in SAEM's threat prioritizations. These types of events may be better handled as special exceptions in the risk assessment and need to receive special consideration based on an organization's policies.

1.4.3.2 Types of Risk Mitigation Controls

In this thesis SAEM is limited to evaluating only technological risk-mitigation controls, i.e., security technologies. Security procedures are an important element of any risk-mitigation plan, as is security training, but these procedures were not included in this thesis because operational security procedures are often unique to each organization so comparing case studies would have been problematic.

1.4.3.3 Improved Design Decisions?

Since SAEM captures the expertise and experience of the security manager the quality of the results, and the decisions made based on those results, are directly tied to the level of experience and expertise of the participants. In two case studies, the coverage analysis showed weaknesses in the security architecture and the security managers took action based on this analysis; however, this does not necessarily indicate that the method improves design decisions. One could argue that providing a structured framework and systematic way to evaluate risks and technologies results in *better* decisions, but the results of this thesis show that SAEM influenced the security managers' decisions, not that these decisions were better. Additional research is necessary to determine whether the method improves decisions.

1.5 Future Work

The greatest need for future research is to establish industry-specific risk information and risk-mitigation control effectiveness ratings. Currently, no such empirical or canonical risk databases exist, and there is only limited empirical data on some threats, such as viruses. Each case study security manager was keenly interested in how his or her estimates compared with those of the other participants. Since this thesis presents only three case studies, I could only make limited comparisons between the security managers' estimates. However, less experienced security managers could start with industry-specific threats, until empirical evidence showed them how to modify their estimates. In addition, most security managers would benefit from knowing the expected effectiveness rates of security technologies, even if the effectiveness rates depended on the organization's ability to maintain and correctly install risk-mitigation controls.

1.6 Thesis Roadmap

Chapter 2 of this thesis defines many of the security and multi-attribute analysis terms used throughout the thesis describes current risk-management methods and ongoing research in security cost-benefit analysis techniques. Chapter 3 provides the necessary background in multi-attribute analysis and argues that the additive model is appropriate for cost-benefit analysis of risk-mitigation controls. This chapter is useful for those readers unfamiliar with decision-analysis methods and terminology. Chapter 4 describes in detail the SAEM process that I used in eliciting threat and risk-mitigation control effectiveness information from the security managers of the three organizations used as case studies in the development of this

thesis. Chapters 5, 6, and 7 describe the case studies and the results from using SAEM. Chapter 8 analyzes the method and many of the additive-model assumptions described in Chapter 3, such as the linearity of the value functions. Chapter 8 also explores how frequency affects the threat prioritizations. Finally, Chapter 9 describes my observations from using SAEM with the security managers from each of the case-study organizations and describes future work that would significantly improve the benefit of SAEM for less-experienced security managers.

CHAPTER 2. Background

2.1 Introduction

The purpose of this chapter is: 1) to define the terms used throughout this thesis, 2) to describe the current practice of security risk-management, 3) to describe current research in security decision-analysis methods, and 4) to show how SAEM contributes to the current practice of security risk-management and differs from recent theoretical decision-analysis approaches. SAEM provides a quantitative analysis tool that allows security managers to gain insight into their assumptions about risk and security measures for the purpose of making decisions about risk-mitigation controls. In contrast to the most widely used quantitative method, Annualized Loss Expectancy, or ALE, that focuses on return on investment, SAEM allows the security manager to easily see the *relative* value of risk mitigation. Security managers can use the relative value of a risk-mitigation control to make security-architecture design decisions.

2.1.1 Terminology

Security practitioners and researchers tend to use common definitions, with only slight variation, for most security terms. The definition for security architecture² appears to be the only significant exception. Few authors of security architecture books have attempted to define security architecture, despite widespread use of the term throughout their books. In this thesis I will use these definitions for each of the following terms:

2.1.2 Security Terminology

- Asset – Any entity of value within an organization. Examples of organizational assets are employee productivity, reputation, public image, and revenue.
- Attack – An instance or realization of a threat. Attacks usually result in an information system security compromise.
- Outcome – The consequences or damages that result from the security compromise caused by a successful attack. For example, an outcome is the lost revenue and damage to an organization’s reputation that may result from a system security compromise.
- Risk – The possibility of asset damage due to a threat. In this thesis, risk is a function of the frequency and outcome of attacks.
- Risk Assessment – The process of determining an information system’s threats and vulnerabilities, and the value of information system assets. Risk assessment is a sub-process of risk management.

² Ramachandran attempts to define it but gives only vague definitions.

- Risk Management – The process of balancing the operational and economic costs of risk-mitigation controls with the expected reduction in risk those controls deliver.
- Risk-Mitigation Control– A security procedure or technology used to prevent, detect, or recover from an information system security compromise. Risk is reduced through use of risk-mitigation controls or mechanisms.
- Security Architecture – The security policies and risk-mitigation controls that are integrated into an organization’s information system architecture for purposes of reducing the risk from a threat.
- Security Compromise – A violation of the information system’s security policies or procedures that results in a loss of system availability, confidentiality or integrity.
- Security Policy – A statement of what is and what is not allowed with respect to an information system.
- Threat³ – A potential event that could lead to an information system compromise. Examples of threats are *Denial of Service* attacks, *Procedural Violations*, *IP Spoofing*, etc.
- Vulnerability – A flaw or defect in system security procedures, design, implementation, or internal controls that, if exploited, could result in a security compromise. Threats exploit vulnerabilities. Examples of vulnerabilities are: buffer overflow conditions, default passwords left during new software installation, disgruntled employees with access to sensitive information.

2.1.3 Decision Sciences Terminology

Although there are no universal definitions of the terms objective and attribute in decision sciences, the following informal definitions are used in this thesis:

- Attribute – A measure of the degree to which a given objective has been attained. 50 hours of lost productivity or \$5,000 in lost revenue are examples of attributes.
- Objective – Something that one’s efforts are intended to attain or accomplish. For example, minimizing risk and minimizing costs are objectives⁴. Other examples include minimizing damage to public reputation and minimizing lost revenue.
- Outcome – A vector of attributes. For example, the outcome from a successful attack may be 20 hours of lost productivity and \$1,000 in lost revenue.

³ In this thesis, threats and attacks are italicized.

⁴In this thesis, the protection against accidental or deliberate disclosure, modification, loss, or interruption of an information system’s critical assets are intermediate objectives.

2.2 Risk Management: State of the Practice

This section presents three risk-management methods that seek to provide security managers with systematic approaches for determining cost-effective, risk-mitigation controls. These qualitative and quantitative methods are not meant to represent all possible risk-management techniques, but rather those that are likely to be most familiar to a broad community of security professionals in the government and private sectors. These methods are presented because they 1) represent the most common state-of-the-practice security risk-assessment methods 2) show the essential elements of security risk-assessment methods, and 3) will facilitate the reader's understanding of SAEM's contribution to state-of-the-practice risk-assessment methods. A summary of the key points of these methods is presented at the end of this section so the reader can understand the SAEM's relative strengths and the weaknesses of these methods.

2.2.1 *Qualitative Risk Management: National Institute of Standard and Technology*

The National Institute of Standards and Technology is responsible for recommending standard federal security practices to government agencies. Government security managers are encouraged to follow the NIST risk-management method as defined in (Stoneburner, Goguen et al. 2001), which consists of three processes: 1) risk assessment, 2) risk mitigation, and 3) evaluation and assessment⁵. Although the risk-assessment process eventually results in a quantitative evaluation of risks, the management method is more qualitative than quantitative. Security managers use three levels of assessment—high, medium, and low—to establish the likelihood and impact of a threat-vulnerability realization, i.e. an attack. The qualitative nature of the method, combined with its limitations around understanding the benefits of specific mitigation technologies, handicap security managers as they attempt to make credible recommendations for security investments.

2.2.1.1 *Risk Assessment*

The purpose of the NIST risk assessment is to determine the extent of the potential threats and risks associated with an information technology system. NIST identified nine primary steps in the risk assessment process: 1) System Characterization, 2) Threat Identification, 3) Vulnerability Identification, 4) Control Analysis, 5) Likelihood Determination, 6) Impact Analysis, 7) Risk Determination, 8) Control Recommendations, and 9) Results Documentation.

The purpose of the System Characterization step is to define the scope of the risk assessment. The security manager identifies system boundaries, resources, personnel, and the sensitivity and criticality of the system assets. Next, the security manager identifies all of the organization's information-system threats and the potential sources of these threats. In the Vulnerability Identification step, the security manager identifies all the vulnerabilities that a threat might exploit and the potential assets targeted by the threat. The result of the vulnerability identification step is a set of vulnerability-threat pairs. Next, in the Control

⁵ The evaluation and assessment process requires security managers to continually update and the risk assessment and risk mitigation results, so only the risk assessment and risk mitigation processes will be discussed in this section.

Analysis step, the security manager evaluates the planned or existing risk-mitigation controls that will help him or her determine the likelihood of a vulnerability-threat attack, which fulfills the next step--Likelihood Determination--in the risk assessment process. The likelihood of an attack is rated according to the likelihood definitions in Table 2 - 1. (Stoneburner, Hayden et al. 2001)

Table 2 - 1 Likelihood Definitions

| Likelihood Level | Likelihood Definition |
|------------------|--|
| High | The threat-source is highly motivated and sufficiently capable, and controls to prevent the vulnerability from being exercised are ineffective. |
| Medium | The threat-source is motivated and capable, but controls are in place that may impede successful exercise of the vulnerability. |
| Low | The threat-source lacks motivation or capability, or controls are in place to prevent, or at least significantly impede, the vulnerability from being exercised. |

In step 6, the security manager assesses the impact of an attack using the impact definitions found in Table 2 - 2. These definitions are given by NIST in recognition of how difficult it is to quantitatively measure the non-tangible consequences of successful threat actions. However, the security manager uses the risk-level matrix shown in Table 2 - 3 to compute a risk level for each vulnerability-threat pair based on likelihood and impact. Therefore, every vulnerability-threat pair is assessed as having a high, medium, or low risk level.

Table 2 - 2 Impact Definitions

| Likelihood Level | Impact Definition |
|------------------|---|
| High | Exercise of the vulnerability (1) may result in the highly costly loss of major tangible assets or resources; (2) may significantly violate, harm or impede an organization's mission, reputation, or interest; or (3) may result in human death or serious injury. |
| Medium | Exercise of the vulnerability (1) may result in the costly loss of tangible assets or resources; (2) may violate, harm, or impede an organization's mission, reputation, or interest; or (3) may result in human injury. |
| Low | Exercise of the vulnerability (1) may result in the loss of some tangible assets or resources or (2) may noticeably affect an organization's mission, reputation, or interest. |

Finally, the security manager identifies risk-mitigation controls, documents the threats and vulnerabilities, and provides recommendations for risk-mitigation controls. Note that the security manager only identifies risk-mitigation controls in the risk assessment, but does not conduct a cost/benefit analysis of these controls until the next process: risk mitigation.

Table 2 - 3 Risk-Level Matrix

| | | Impact | | |
|-------------------|--------------|-----------------------------|--------------------------------|---------------------------------|
| | | Low (10) | Medium (50) | High (100) |
| Threat Likelihood | High (1.0) | Low $10 \times 1.0 = 10$ | Medium $50 \times 1.0 = 50$ | High $100 \times 1.0 = 100$ |
| | Medium (0.5) | Low $10 \times 0.5 = 5$ | Medium $50 \times 0.5 = 25$ | Medium $100 \times 0.5 = 50$ |
| | Low (0.1) | Low $10 \times 0.1 = 1$ | Low $50 \times 0.1 = 5$ | Low $100 \times 0.1 = 10$ |

2.2.1.2 Risk Mitigation (Cost/Benefit Analysis)

The first four steps of the risk-mitigation process require the security manager to 1) prioritize the risk-mitigation controls (from high to low), 2) evaluate and recommend control options based on feasibility and effectiveness, 3) conduct a cost-benefit analysis, and 4) select the controls based on the cost-benefit analysis. The remaining three steps of the risk-mitigation process require the security manager to develop a plan for implementation of the risk-mitigation controls.

The NIST-recommended, cost-benefit-analysis approach focuses primarily on the cost of the control and not the benefit of the control. The only reference to assessing the benefit of a risk-mitigation control in the NIST Risk Management Guide states:

“The organization will need to assess the benefits of the controls in terms of maintaining an acceptable mission posture for the organization. Just as there is a cost for implementing a needed control, there is a cost for not implementing it. By relating the result of not implementing the control to the mission, organizations can determine whether it is feasible to forgo its implementation.”

Although NIST endorses a qualitative or quantitative cost-benefit analysis, Jacobson (Jacobson 2002) argues that “Assessing a risk as ‘high,’ ‘unacceptable,’ or in other qualitative terms does not provide the information needed to support the decision to implement risk mitigation measures, which will always have quantitative implementation costs.” Since NIST did not provide a quantitative method for assessing the economic cost of a risk, it is unclear how security managers would conduct a quantitative cost-benefit analysis.

2.2.2 Risk Management: International Information Systems Security Certification Consortium (ISC)²

The most widely recommended and common (Anderson 2001) risk-management method is that put forth by the International Information Systems Security Certification Consortium,

(ISC)² (pronounced I S C squared), an international non-profit organization dedicated to maintaining a common body of information security knowledge⁶. (ISC)² recommends that security practitioners use a *quantitative* risk-management method based on Annualized Loss Expectancy, or ALE.

Although the (ISC)² (Knutz and Vines 2001; Harris 2002) defines risk management as the “process of identifying, assessing, and reducing risk to an acceptable level and implementing the right mechanisms to maintain that level of risk,” the emphasis in this (ISC)² risk-management process is on the risk-analysis sub-process. The (ISC)² gives little guidance on how to evaluate the benefit of risk-mitigation strategies or how to determine the residual risk that remains after implementing these strategies.

While the ALE method is effective in helping the security manager estimate the expected loss from a specific threat-vulnerability event, it requires the security manager to estimate the economic impact of all threat-vulnerability events before finding the most effective security technology. A good risk analyst could develop hundreds of threat-vulnerability scenarios requiring hundreds of financial loss estimates. Furthermore, method advocates recommend that security managers invest in risk mitigation controls that have the highest net benefit, i.e. greatest ALE reduction and lowest cost. Anderson (Anderson 2001) believes that few managers find the ALE results credible because the method produces estimates of risk, so the inputs are tweaked to produce acceptable results.

2.2.2.1 Risk Analysis

The first step in the (ISC)² risk-management method is to conduct a risk analysis, which has three steps: 1) estimate the potential losses to assets by determining the assets’ value, 2) analyze potential threats to these assets, and 3) compute the organization’s Annualized Loss Expectancy. Security analysis teams usually conduct the risk assessment. In this risk-analysis process, the security team must first determine a dollar value for the organization’s information and assets; then identify the organization’s threats and associated vulnerabilities that will lead to a loss if an attack occurs; and, finally, estimate the expected threat frequencies. For example, a security-team analyst may identify an attacker (the threat), who will take advantage of lax firewall settings (the vulnerability), to steal trade secrets (the asset).

Since (ISC)² does not provide a definitive set of threats or vulnerabilities, the complete identification of threats, vulnerabilities and assets is left to organizations’ security analysts, whose security expertise and experience vary greatly across organizations. In addition, one threat may take advantage of one of several vulnerabilities and potentially target many assets, each constituting a different threat-vulnerability-asset triplet. The set of triplets could get quite large depending on the imagination of the risk-analysis team.

The final step in the (ISC)² risk-analysis process is to define the ALE, which represents the product of the expected rate of a threat’s occurrence and the expected loss, expressed in monetary terms, resulting from a single occurrence. Before the security team can compute the ALE, they must first determine the Exposure Factor (EF), the Single Loss Expectancy (SLE),

⁶ The (ISC)² developed the Certified Information System Security Professional (CISSP) program, which has tested and certified over 40,000 security professionals. The CISSP certification test includes several questions on risk management and ALE.

and the Annualized Rate of Occurrence (ARO). The EF is the percentage loss that a realized threat event would have on an asset.

For example, a security manager may estimate that a denial-of-service attack may result in a 20% loss of revenue to the organization, so the Exposure Factor is .20. The next step is to compute the SLE, i.e., Asset Value (\$) * EF. For example if the lost revenue for the threat-vulnerability pair were estimated to be \$1,000, then the SLE is \$200. Next, the security team estimates the ARO, which is an estimation of the probability that a threat will occur during the next year. Finally, the team computes the ALE using the SLE and ARO:

$$ALE = SLE * ARO.$$

For example, if the ARO is .8, which indicates a high probability that the threat will occur, then the ALE is \$160 (.8 * \$200).

2.2.2.2 Cost/Benefit Analysis

Once the security team determines the ALE for each threat-vulnerability-asset combination, (ISC)² recommends that security managers conduct a cost/benefit analysis to determine the value of a risk-mitigation control. The manager determines the value of the control as follows:

$$\begin{aligned} &(\text{Value of control to the organization}) = (\text{ALE before implementing control}) \\ &\quad - (\text{ALE after implementing control}) - (\text{Annual cost of control}) \end{aligned}$$

Furthermore, the (ISC)² recommends that security managers should not invest in risk-mitigation controls that cost more than the ALE, which follows good business sense. Therefore, the security manager can prioritize the threat-vulnerability-asset triplets in order of ALE to determine which triplets pose the greatest risk to the organization. He or she can then determine the risk-mitigation controls that provide the greatest value by ordering the controls according to greatest reduction in ALE. However, the greatest disadvantage to the (ISC)² risk management process is that the risk analysis team must determine the dollar value of all assets, which may be difficult for intangible assets such as public reputation, customer perceptions, or for the value of a human life in a safety critical system.

2.2.3 Risk Analysis: Operationally Critical Threat Asset, and Vulnerability Evaluation (OCTAVEsm)

The Software Engineering Institute developed OCTAVEsm (SEI 2003) as a risk-analysis process that is intended to help organizations measure their information-system risk. There are three phases in the OCTAVEsm process: Threat Profile Development, Vulnerability Assessment, and Key Practice Evaluation. The goal of OCTAVEsm is to provide organizations with a systematic and thorough evaluation and measurement of their security practices, from which key stakeholders can make cost-effective decisions about risk-mitigation strategies. This section provides a brief overview of OCTAVEsm.

The first phase in the OCTAVEsm process is to develop threat profiles. In this step, organizational staff members identify important information assets, threats to those assets, and requirements for keeping them secure. As a minimum, a threat profile consists of an 1) asset, 2) actor (similar to a threat), and 3) outcome. In contrast to the definition provided in Section 2 of this chapter, the OCTAVEsm process specifically defines an outcome as “the immediate result of violating the security requirements of an asset (disclosure, modification, destruction,

loss, interruption).”(Alberts and Dorofee 2001) The result of the threat-profile development phase is the identification of an organization’s critical assets and the potential threats that, if realized, would result in disclosure, modification, destruction, loss/destruction, or interruption of the asset.

In the second phase of OCTAVEsm, the analysis team examines key operational components for vulnerabilities that can lead to a security compromise of the critical assets identified in the first phase. Therefore, each asset has an associated set of vulnerabilities that, if exploited, could result in one or more of the outcomes identified in the first phase.

In the third phase of OCTAVEsm, the analysis team conducts a risk analysis and develops a protection strategy. During the risk analysis, the team uses a qualitative scale (high, medium, and low) to evaluate the impact to critical assets of threats that successfully exploit vulnerabilities. Finally, the OCTAVEsm risk-analysis team creates protection strategies, i.e., risk-mitigation controls, to reduce the risks to critical assets.

Overall, the OCTAVEsm process provides a systematic and qualitative risk-management method that focuses on the risk-evaluation component of risk management, but does not suggest methods for cost-benefit analysis. In addition, information-protection decisions are based on risks to the confidentiality, integrity, and availability of critical information technology assets, rather than the risks to the organization’s more significant assets, such as public reputation or revenue.

2.2.4 Summary of Risk Management Techniques

In this section, I presented three different risk-management methods that security managers can use to help bring information security risks to an acceptable level within their organizations. Although two of these risk management processes are qualitative -- NIST’s Risk Management Process and OCTAVEsm -- and one quantitative—(ISC)²-- there are some key similarities and weaknesses among the processes. I highlight these similarities and weaknesses because, although SAEM relies on elements of these risk management methods, SAEM also addresses some of their weaknesses

2.2.4.1 Key Ideas Underlying Security Risk Management

First, the most significant similarity is that all three processes rely on the organization’s best available information concerning threats, vulnerabilities and assets. In each method, security managers analyze the organization’s threats, vulnerabilities, and assets as part of the risk-assessment process. Furthermore, all three processes require that the security manager or risk-analysis team determine which vulnerabilities a threat might exploit and consider the frequency and outcome of an attack in assessing the overall level of risk from a threat.

Second, each method’s risk-assessment process allows an organization to prioritize its risks. If the security manager uses the (ISC)² risk management process, he or she can prioritize threats based on ALE . If the security manager uses one of the qualitative risk-management methods, threats can be prioritized based on the high, medium, or low assessment. However, all of the risk-management methods prioritize threats based on the estimated frequency and consequences of the threat.

Another similarity among the three risk-management methods is that all methods evaluate security technologies based on a cost-benefit analysis. Each risk-management method advocates that risk-mitigation controls should be selected based on an assessed value of the technology. In the qualitative methods, the security manager determines which technologies reduce the risk to an acceptable level and then selects the most cost-effective technology. Using ALE, the security manager re-computes the ALE for a threat/vulnerability/asset triplet, assuming the security technology, and then selects from among the technologies with the greatest net values. Both qualitative and quantitative risk management methods recognize the need for selecting risk-mitigation controls based on their value to the organization—a relatively recent phenomenon in security engineering.

2.2.4.2 Problems with State-of-the-Practice

State-of-the-practice risk-management methods have two problems. Although security managers prefer quantitative over qualitative risk-management methods when conducting cost benefit analysis, information system executives are often skeptical of ALE results. In addition, ALE is not useful when security engineers need to evaluate the technology's value in the context of the security architecture design. This thesis attempts to improve the state-of-the-practice by addressing these two problems.

Ideally, security engineers and managers would like to use quantitative risk-management methods when selecting security technologies for the information system architecture because qualitative assessments may not be meaningful to decision makers trying to make finer grained decisions; however, ALE often lacks credibility with these decision makers(Anderson 2001). Decision makers are skeptical of making *precise* security-investment decisions based on subjective, albeit experienced, estimates of risk. In addition, ALE requires that the organization convert non-tangible assets to dollars before computing SLE. This adds to the subjectivity of the process and the skepticism of the results.

The second problem with state-of-the-practice is that the methods do not address how to value risk-mitigation controls, and, more specifically, how to value the risk-mitigation controls that simultaneously reduce the risk from several threats. For example, encryption is often effective against several threats so its value to an organization may be greater than a security technology that greatly reduces the risk of only a few threats, such as host-intrusion detection software. In addition, none of the risk-management methods help the security manager evaluate risk-mitigation controls with respect to security engineering goals, such as protection, detection, or recovery. Continuing with the example, if encryption provides stronger protection against already sufficiently protected threats, then detection technologies may be more useful in the security architecture. State-of-the-practice risk-management methods are not particularly useful in helping security engineers and managers with these types of engineering decisions.

2.3 Current Research

I developed SAEM, using multi-attribute analysis techniques, to support security managers in making cost-effective design decisions about security architectures. Therefore, it is necessary to review relevant research in techniques software engineering design decisions and security cost-benefit analysis.

2.3.1 Decision Methods

Previous research on design decisions in software engineering is sparse. In 1990, Thomas Lane (Lane 1990) developed a framework for classifying user interface design knowledge so that software designers could make good structural choices based on the user's functional requirements. Later, Kazman, et al. (Kazman, Asundi et al. 2000) use multi-attribute analysis techniques to estimate the costs and benefits of software architectural attributes, such as performance, security, and modifiability, so that software engineers can make tradeoffs among information system architectural design decisions.

Kazman was not the first to suggest using decision analysis techniques for making design decisions. Kontio (Kontio 1996) first proposed using a well-known decision analysis technique called Analytic Hierarchy Process (Saaty 1990) to help software engineers make systematic decisions about selecting commercial-off-the-shelf (COTS) products. However, only Finne (Finne 1998) has suggested that formal decision-making techniques should be applied to making decisions about information security. Unfortunately, his paper did not provide specific framework for using those techniques in selecting risk-mitigation controls.

Other security decision-analysis models have been qualitative in nature, such as the countermeasure matrix (Straub and Welke 1998), which compares two security risk-mitigation controls by highlighting their strengths and weaknesses relative to deterrence, prevention, detection, and remedies. However, the idea that security practitioners should categorize and evaluate a security technology by its capabilities is already established in NIST's (Stoneburner 2001) technical security service model. NIST categorizes security technologies according to three primary purposes—support, prevent, and recover (others (King, Dalton et al. 2001; Bishop 2003) have categorized security technologies slightly differently—prevent, detect, and recover). Regardless of which categorization is used, NIST encourages security managers to use these categories to help evaluate the benefit of a security technology. Unfortunately, security managers have little structured or quantitative guidance on how to actually evaluate and compare technologies based on these categories.

2.3.2 Security Cost-Benefit Analysis Research

The previous section explored current methods of security risk management, which emphasize the identification of threats, vulnerabilities and assets, but do not offer security managers any systematic cost-benefit analysis methods that could be used to select risk-mitigation controls. To date, research in security cost-benefit analysis methods has also failed in providing security managers with “real world” guidance in selection risk-mitigation technologies. Research in information security has focused primarily on technical defense mechanisms, such as encryption, intrusion detection, or access control or vulnerability taxonomies (Lanwehr, Bull et al. 1994; Corporation 2003); research related to the economic analysis of security investments is sparse (Gordon and Loeb 2002). Researchers offer theoretical cost-benefit models that are too generic to offer security managers any operational guidance in selecting risk-mitigation controls. This section discusses the current research in security cost-benefit analysis methods.

2.3.2.1 Economic Models

For many years, security managers followed the Orange Book (NCSC 1985) as the definitive security guide. The Orange Book did not include risk management in its guidance to security

managers, so practitioners were often criticized because they didn't understand economic tradeoffs. Only recently have researchers attempted to address the issue of how security managers should invest in risk-mitigation controls. Early cost-benefit models addressed economic tradeoffs for risk-mitigation techniques with respect to specific threats (Millen 1992; Meadows 2000), such as denial- of-service attacks. In 1998, Finne (Finne 1998) proposed a theoretical model as an approach to balancing risk-mitigation costs, economic losses, and levels of information security. This theoretical approach proposed that each organization must find its "specific point of optimization where the total costs are the lowest" given the level of information security required. Although Finne did not specifically address how one could measure the level of security, he thought that companies should be able to estimate the cost of not having risk-mitigation controls, which would help establish the benefit of these controls.

Recently, Gordon and Loeb (Gordon and Loeb 2002) have proposed an economic model that suggests security managers should not invest more than 37% of the expected loss of an asset in risk-mitigation controls. This economic model is based on three probabilistic parameters: 1) the probability of a threat occurring, 2) the expected monetary loss to the organization caused by a security compromise, and 3) the probability that a threat would be successful once an attack is initiated. The authors concede that determining the probabilities for the model is not easy and that their model does not cover instances where a single investment in security protects multiple threats and vulnerabilities. Moreover, as in most economic models, their model would require security managers to estimate the financial loss of an attack, which can be difficult for non-tangible assets.

2.4 Security Attribute Evaluation Method (SAEM)

This thesis proposes SAEM as a quantitative cost-benefit analysis method that helps security managers compare risk-mitigation controls; therefore, SAEM is intended to be part of an organization's risk-management process. SAEM addresses the current problem with ALE because 1) SAEM de-emphasizes the actual computed value of the risk and emphasizes the *relative* values of the risks and risk-mitigation controls, and 2) security managers are not required to estimate the consequence of an attack in dollars in order to make meaningful comparisons of threats and risk-mitigation controls, and 3) SAEM integrates a new coverage-analysis model that allows security managers to evaluate how a risk-mitigation control contributes to the security architecture's defense-in-depth design, a fundamental security engineering design principle.

SAEM provides security managers with a framework to explore how their assumptions about the organization's security risk affect their selection of security technologies. Consistent with the previously described risk-assessment processes, SAEM prioritizes the threats based on the organization's best estimates of risk, i.e. frequency of a threat and the most-likely outcomes from a security compromise. The SAEM's benefit analysis step uses the information from the risk assessment to evaluate risk-mitigation controls and prioritize them based on their overall contribution to the security architecture.

One of the goals of SAEM is to provide insight to security managers when comparing risk-mitigation controls, so SAEM also provides a coverage-analysis model. This model enables security managers to compare risk-mitigation controls with respect to defense-in-depth (prevention, detection, and recovery) and breadth-of-defense coverage against multiple threats.

No other risk-management method provides this perspective. NIST's only guidance on the defense-in-depth security engineering principle is:

“Through user training and awareness, well-crafted policies and procedures, and redundancy of protection mechanisms, layered protections enable effective protection of information technology for the purpose of achieving mission objectives.”(Stoneburner, Hayden et al. 2001)

Since the security manager uses SAEM to help make comparisons among different risk-mitigation controls, he or she can focus on developing cost information on the controls that best meet the organization's objectives of minimizing risk or cost. Developing cost information can be time-consuming and expensive; a security manager's time is best spent developing cost information on risk-mitigation controls that are most promising to the organization's objectives. SAEM helps security managers identify technologies that will result in the greatest risk reduction for the organization.

CHAPTER 3. Multi-attribute Analysis Essentials

3.1 Introduction

Multi-attribute analysis provides a convenient framework for developing quantitative risk assessments, cost/benefit analyses, and security technology tradeoffs that result in a set of prioritized risk mitigation strategies. This framework, used in this thesis, relies on an *additive value model* to construct a prioritized list of risks and determine the benefit of different security technologies in an organization's security architecture. Decision makers (usually an organization's lead security specialists) can construct an additive value model when certain conditions of independence exist, such as *preferential independence* and *difference independence*. Since the Security Attribute Evaluation Method (SAEM) depends on the additive value model in three of its four phases (risk assessment, benefit analysis, and technology tradeoff analysis), it is important to understand the required steps for constructing such a model.

This chapter discusses why multi-attribute analysis techniques are appropriate for developing risk assessments, conducting benefit analyses, and comparing security technologies. Using the risk assessment as an example, this chapter also describes the additive model, the underlying assumptions of independence, and the four steps needed for constructing the model. Chapter 4 will describe how the benefit analysis phase adjusts the input parameters of the risk assessment's additive model and the security technology tradeoff analysis additive model.

3.2 Multi-attribute Analysis Applicability

Multi-attribute analysis techniques were developed to help decision makers evaluate alternatives. For the purposes of this thesis, the decision makers are the organization's lead security specialists and information technology (IT) managers, who need to evaluate different security technologies for integration into the information system's security architecture. Ideally, these individuals need a decision framework that allows them to systematically and consistently evaluate alternative security technologies based on their organization's risk and security environments. In addition, the decision framework should allow security specialists and IT managers to identify the security technologies that best meet the objectives of the organization, such as minimizing cost and maximizing adoptability of the security technology. The multi-attribute analysis techniques used in SAEM help an organization's decision makers find the security technologies that best meet their needs.

The multi-attribute techniques in SAEM have several advantages. These techniques allow security managers to identify their organizational risks, express their expectations about the consequences of successful attacks, and provide insights into how their assumptions about the effectiveness of their risk mitigation strategies affect their decisions about the security architecture. Security managers can conduct "what-if" analysis to see whether their decisions are sensitive to their assumptions, spending constrained resources to gain better information where it matters. The value of a multi-attribute analysis is not only in the numbers produced, but also in the insights that security managers gain during sensitivity analysis as well as in each refinement step of the evaluation.

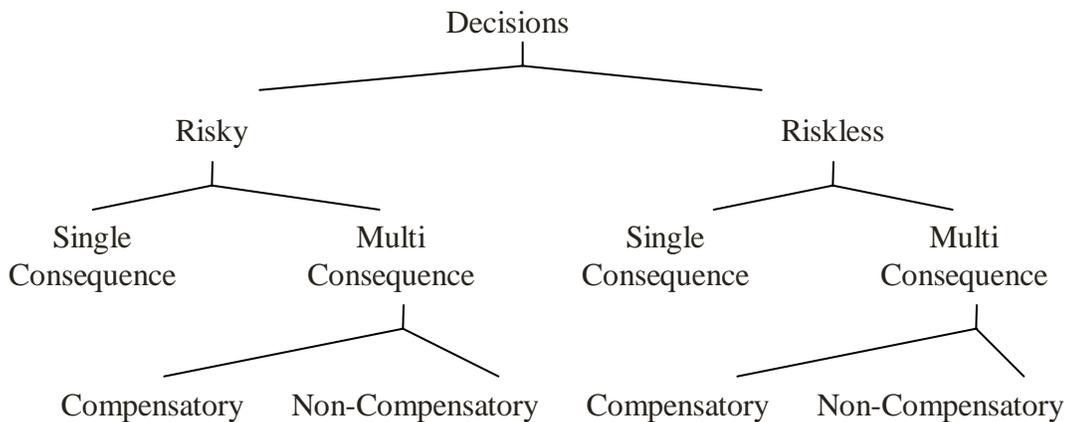
Another advantage to multi-attribute analysis techniques is that they provide security managers with a systematic and repeatable method for evaluating an organization's risks using the best available threat information. As security managers gain better threat and effectiveness information, the managers can easily update the models with new input data so they can measure the marginal effect of the new information on their decisions. In addition, as the organization discovers new threats and security technologies, the security manager can easily integrate these discoveries into to the analysis.

Finally, the multi-attribute analysis framework also provides the basis from which a security manager can systematically evaluate and compare alternative security technologies. Often an organization's decision to spend scarce financial resources to purchase security technologies depends on several factors, such as cost, complexity, maintenance, and organizational culture. The organization must make tradeoffs among these factors when comparing alternatives. Multi-attribute analysis techniques provide an efficient and consistent structure for evaluating different security technologies.

3.2.1 Types of Decisions

Decision analysts characterize decisions as either *riskless* or *risky*. Riskless decisions are ones in which the decision consequences are certain and risky decisions are ones in which the consequences are unknown. Decision analysts further characterize both types of decisions as *single-objective* decisions or *multi-objective* decisions. Single-objective decisions address a single consequence, whereas multi-objective decisions address multiple consequences. Finally, decision analysts further distinguish between *compensatory* decisions that allow tradeoffs among the consequences and those that do not, i.e. *non-compensatory*. Figure 3 - 1 shows the classification of decisions.

Figure 3 - 1 Decision Classifications



The multi-attribute analyst can structure the security technology benefit analysis as a *risky*, *multi-objective*, and *compensatory* decision problem. Before the analyst can evaluate the benefits of security technologies, the analyst first develops the risk assessment as the cost framework

against which the technologies can be evaluated. As part of the risk-assessment process, security managers identify their organization's threats⁷ and the potential consequences or outcomes from successful attacks. Since organizations identify several consequences, the benefit analysis is multi-objective. Furthermore, the benefit analysis is also compensatory because security managers can make tradeoffs among the consequences by specifically choosing security technologies that mitigate a given consequence.

The types of consequences (X_i) that can result from an attack constitute outcome *attributes* in the multi-attribute risk assessment and the actual attribute damage from an attack is the attribute's value (x_i). Therefore, each attack outcome can be described as a vector of attribute values $O_a(x_1, x_2, x_3, x_4)$. Similar attacks can have many different outcomes because each attack instance results in different consequences. For example, a denial of service attack could result in a few or many hours of lost productivity depending on the nature of the attack.

Compensatory decisions permit tradeoffs among attributes. The multi-attribute benefit analysis is characterized as compensatory because managers are often willing to trade one outcome for another. For example, some organizations are willing to trade some amount of revenue to avoid public embarrassment. Security managers in these organizations might prefer to invest significant resources in security technologies that reduce the risk of attacks that are likely to result in public embarrassment. Other organizations may be much more interested in investing in security technologies that reduce the risk of lost revenue.

Multi-attribute analysts use several models to help decision makers balance preferences by quantifying the subjective factors that influence their decisions. More specifically, multi-attribute analysis allows the security manager to express outcomes in non-economic terms, which is useful when describing outcomes that are difficult to quantify, such as the potential damage to corporate image from a successful attack. The *additive value* model provides a simple mechanism for combining the attributes values in order to determine a risk ranking for each threat. This model also enables security managers to see the overall risk mitigation that security technologies provide, and allows security managers to compare security technologies and balance organizational security objectives.

3.2.2 The Additive Value Model

The additive value model relies on the additive value function. The general form of an additive value function is:

$$V(x_1, x_2, \dots, x_n) = \sum_{i=1,n} w_i v_i(x_i)$$

where $v_i(x_i)$ is a single-attribute value function defined over levels of x_i , and w_i is a scaling constant that weights the value function for attribute value x_i . Constructing an additive value model involves four steps:

- Determine the single-attribute value functions V_1, V_2, \dots, V_n

⁷ Recall in this thesis, *threats* are defined as potential events, such as denial of service attacks, procedural violations, IP spoofing, which could lead to an information system compromise. An attack (a) is an instance of a threat that results in an information system compromise that has an outcome (O_a) of one or more consequences (X_i). For example, a single system compromise may ultimately result in lost revenue (X_1), public embarrassment (X_2), lost productivity (X_3), and damaged corporate image (X_4).

- Determine the weighting factors w_1, w_2, \dots, w_n
- Compute the value of each alternative and rank alternatives
- Conduct sensitivity analysis to see how sensitive the ranking is to model assumptions

The next two sections describe in more detail the additive multi-attribute value model; intuitively, the single-attribute value function ensures that the attribute values sum together using the weights that reflect the assessed preferences (e.g., security manager’s preference) and the additive value function is a linear combination of these single-attribute value functions.

3.2.3 Additivity Assumptions

The additive value model is valid if *transitivity*, *preferential independence*, *tradeoff independence*, and *difference independence* conditions exist among the attributes. Although it is not possible to prove that all of the requisite additivity assumptions hold in every case, there is strong evidence that even when there is not complete independence, the additive value model provides close approximations to “pure” additive value functions. (Yoon and Hwang 1995)

3.2.3.1 Transitivity

The *transitivity* condition holds if the decision maker can partially order the outcomes. More formally:

if $O_1, O_2,$ and O_3 are outcome vectors,

and $(O_1 \succeq O_2)$ and $(O_2 \succeq O_3),$

then $(O_1 \succeq O_3),$

where the symbol “ \succeq ” is read ‘is preferred or indifferent to.’

For example, three different security compromises result in varying degrees of lost revenue, public embarrassment, lost productivity, and increased oversight. If the security manager perceives the overall damage described by O_1 as less than O_2 , and the overall damage described by O_2 as less than that described by O_3 , then he or she would perceive the damage described by O_1 as less than O_3 . Transitivity is a normative assumption that any rational decision maker follows.

3.2.3.2 Preferential Independence

Assume that $\{X_1, \dots, X_i, \dots, X_j, \dots, X_n\}$ is a set of attributes and that x'_i and x''_i describe different levels of X_i . Then, if X_i is preferentially independent of $X_1, \dots, X_j, \dots, X_n$ then the security manager’s preference for different levels of X_i , holding all other attributes fixed does not depend on where we hold $X_1, \dots, X_j, \dots, X_n$ constant. Preferential independence is more formally expressed as:

if

$$(x'_i \succeq x''_i)$$

then

$$(x'_i, x_j) \succeq (x''_i, x_j), \text{ for all } x_j.$$

3.2.3.3 Difference Independence

Difference independence goes a step further than preferential independence. The additive model assumes that a decision maker can rank order the differences along a given dimension within an attribute. Difference independence requires that the ranking of the differences in values within the attribute do not change given fixed levels of outcomes in other attributes. For example, continuing with the example from preferential independence, managers prefer less lost productivity than more lost productivity. Also, a security manager's preference for lost productivity decreases the greater the loss. Therefore, the manager's preference for lost productivity is monotonically decreasing. Furthermore, the degree to which they prefer less damage to more should not change if other fixed amounts of other types of damage are inevitable. Difference independence is more formally expressed as:

if

$$\Delta(x'_i \succcurlyeq x''_i) < \Delta(x'_i \succcurlyeq x'''_i)$$

then

$$\Delta((x_i, x_j) \succcurlyeq (x''_i, x_j)) < \Delta((x'_i, x_j) \succcurlyeq (x'''_i, x_j)), \quad \text{for all } x_j.$$

3.2.3.4 Tradeoff Independence

Tradeoff Independence exists if the decision maker's preference ranking for two attributes does not depend on fixed values of other attributes. If $\{X_1, \dots, X_i, \dots, X_j, \dots, X_n\}$ represent decision attributes, then attributes $\{X_i, X_j\}$ are said to be independent of the remaining attributes if the decision maker's preference ranking of different levels of X_i and X_j (i.e., x_i and x_j), holding X_1, \dots, X_n constant, does not depend on where we hold X_1, \dots, X_n constant. For example, if X_i is lost productivity, X_j is lost revenue, and X_z is damaged public reputation, then lost productivity and revenue are preferentially independent of damaged public reputation:

if

$$(x'_i, x'_j, x'_z) \succcurlyeq (x''_i, x''_j, x'_z)$$

then

$$(x'_i, x'_j, x_z) \succcurlyeq (x''_i, x''_j, x_z), \quad \text{for all } x_z.$$

Although this is not a normative assumption, there does not appear to be any reason to believe that it does not hold for computer security problems. Intuitively, the attribute values represent the consequences of an attack and if one outcome is preferred over another, i.e., the security manager perceives that one outcome is more damaging than the other, then the addition of more damage of some other type to either outcome should not change the security manager's perception about the preference for the original two outcomes.

3.3 Multi-attribute Analysis

Once a multi-attribute analyst establishes that the additivity assumptions hold, the analyst can construct the additive model. This section describes in detail the steps in constructing a multi-attribute additive model using four outcome attributes: Lost Productivity, Lost Revenue, Damaged Public Reputation, and Additional Regulatory Penalties (the increased administrative burdens that an organization suffers from an external oversight agency because of a security

compromise). Assume that these four attributes represent a hypothetical organization's most significant concerns in the event of a security compromise.

3.3.1 Determine the Single Attribute Function ($v(x_i)$)

The first step in constructing a multi-attribute additive model is to assess a single-attribute *value function* for each attribute. The purpose of this function is to reflect preferences for outcomes over the relevant range for each attribute. The relevant range of each attribute depends on the attribute values provided by the security manager. For example, in one organization, the estimated hours of lost productivity could range from 1,000 hours to 0 hours and the estimated damage to public reputation could range from 1 to 7 on a well-defined Likert⁸ scale. In contrast, another organization's lost productivity could range from 500 hours to 0 hours. The MA analyst assesses a value function for each attribute.

The results of a single-attribute value function are standardized to a 0-1 scale to eliminate computational problems caused by the different units of measure. I used monotonically increasing functions to reflect the consequences and to normalize the attributes. The simplest form of the function is:

$$v_j(x_{ij}) = x_{ij}/x_j^*$$

where x_{ij} is the i^{th} attribute value of the j^{th} attribute and x_j^* is the maximum value for that attribute. This ensures that $0 \leq v_j(x_{ij}) \leq 1$, and as $v_j(x_{ij})$ approaches 1, the consequence is more severe. Although I initially used this linear function for simplicity, I developed other convex and concave functions for sensitivity analysis; these will be discussed later in the thesis. If the interview process reveals a non-linear relationship, then an analyst should use other forms of monotonically changing functions (e.g., convex or concave).

3.3.2 Assess Weighting Factors

The second step in constructing an additive multi-attribute model is to assess the attribute-weighting factors. These weights permit the decision maker to make trade-offs between the attributes. Although several weighting elicitation techniques have been developed (Stillwell and Seaver 1981; Von Winterfeldt and Edwards 1986; Borchering and Eppel 1991; Fischer 1995), I used the Swing-Weight Method in all case studies. It was easy to use and the security managers found it cognitively appealing.

In the *Swing-Weight Method*, the analyst asks the security manager to consider a hypothetical situation, in which the security manager discovers a new type of threat. This threat results in the worst level of damage for each attribute found in all assessed threats. The analyst gives the decision maker the option of improving the hypothetical outcome by changing one attribute to its best level. In succession, the security manager improves each attribute until all attributes are

⁸ A Likert Scale is a unidimensional scale whose intervals are assumed to be equidistant. Below is an example of a Likert scale from one of the case studies to help security managers estimate the amount of damage from an attack.

| | | | | | | |
|------|------|--------------------|----------|----------------------|--------|----------------|
| 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| None | Mild | Moderately Mild | Moderate | Moderately Severe | Severe | Most Severe |

ordered. Next, the security manager assigns a value of 100 to the security manager’s first chosen attribute and values the remaining attributes in relative importance to the first attribute. The actual weights are determined by dividing each of these values by the sum of all the values. The resulting weights sum to 1.0. Table 3 - 1 shows the results of using the swing-weight method for the risk assessment example.

Table 3 - 1 Assessed Weights

| Outcome Attribute | Rank | Rating | Weight |
|----------------------|------|--------|--------|
| Lost Productivity | 1 | 100 | .42 |
| Public Reputation | 2 | 80 | .33 |
| Regulatory Penalties | 3 | 40 | .17 |
| Lost Revenue | 4 | 20 | .08 |

3.3.3 Compute Value and Rank Alternatives

The third step in constructing the additive multi-attribute model is to compute the relative ranking of the alternatives. For the risk assessment and the benefit analysis, this means computing the *Threat Index*, which the analyst uses to rank the threats and security technologies. The threat index (TI) captures the relative importance of each type of attack. For SAEM, I assumed that the decision makers are risk neutral and that utility functions do not have to be assessed. Keeney and Raiffa (Keeney and Raiffa 1999) have developed techniques for integrating preferences of risk-averse decision makers been, but these techniques can be complex to implement. The threat index (TI) for each type of attack (*a*) is computed using the following equation:

$$TI_a = Freq_a * (\sum_{j=attributes} W_j * v_j(x_{aj}))^9$$

where w_j is the attribute weight and x_{aj} is the “most-likely” outcome attribute value of the attack, and $Freq_a$ is the estimated number of attacks per year. Finally, I compute the relative threat index using the following equation:

$$RTI_a = TI_a / TI^* * 100$$

where TI^* is the maximum threat index for the organization’s threats.

Table 3 - 2 shows the data and relative threat index of three threats (Procedural Violations, Theft, and Virus). The weights (w) are shown for each outcome attribute and the second column shows how often the security manager expected an attack to occur. The left-hand column under each of the outcome attributes shows the most likely consequence of an attack. In this example, the security manager expects a procedural violation to occur 4,380 times per year¹⁰ and result in approximately \$2 of lost revenue, have a *mild* impact (“mild” is 2 on the 1-7 Likert scale, where 1 means “no impact” and 7 means “most severe”) on the organization’s reputation, and lose 2 hours of productivity. The right-hand column under each of the outcome attributes shows the normalized values from the attribute value functions. The values in the TI column are dimensionless units, but the TI indicates the relative significance of each threat.

⁹ If x_{aj} is a Likert scale value then the value function normalizes $x_{aj} - 1$ to reduce the effect of the frequency term in the additive model equation on attacks that “most likely” do not have an impact

¹⁰ The security managers often provide threat estimates in frequency units of hours, days, weeks or months. In this thesis, I convert all estimates to a yearly rate.

Table 3 - 2 Outcome Attribute Values and Threat Frequencies

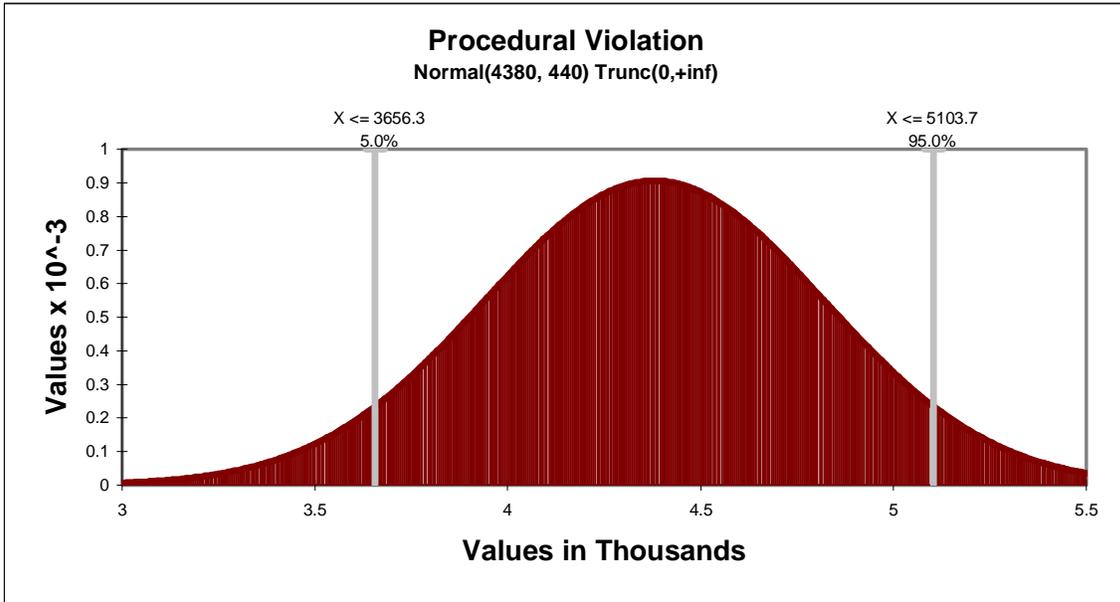
| | | Outcome Attributes | | | | | | | | | TI | RTI |
|---------|----------------------|--------------------|---------------------|-------|------------------------|-------|---------------------------|-------|----------------------------------|-----|-----|-----|
| | | Freq/yr | Lost Revenue (\$\$) | | Reputation (1-7 Scale) | | Lost Productivity (hours) | | Regulatory Penalties (1-7 Scale) | | | |
| | | | w=.08 | w=.33 | w=.42 | w=.17 | | | | | | |
| Threats | Procedural Violation | 4,380 | \$2 | .0008 | 2 | .25 | 2hrs | .0083 | 1 | 0 | 376 | 100 |
| | Theft | 24 | \$182 | .08 | 3 | .5 | 1hrs | .0042 | 2 | .67 | 7 | 2 |
| | Virus | 912 | \$0 | 0 | 1 | 0 | 3hrs | .0125 | 1 | 0 | 80 | 21 |

3.3.4 Conduct Sensitivity Analysis

The final step in constructing the additive model is to conduct sensitivity analysis. The purpose of sensitivity analysis is to determine how sensitive the analysis is to the security specialist's range of uncertainty about key variables. During the elicitation process, the analyst asks the security manager to provide an upper and lower bound to their estimates. The analyst uses these upper and lower bounds to create probability distributions, which the analyst uses as inputs when conducting sensitivity simulations. SAEM uses Automated Security Evaluation Support System (ASESS), a semi-automated decision support tool integrated with a commercial risk analysis package called @Risk developed by Palisade Corporation, to help construct probability distributions and run the simulations. Although Chapter 4 discusses the sensitivity analysis process in further detail, this section describes how ASESS constructs the simulations and the type of information produced at the end of a simulation.

SAEM requires that security managers make several estimates. For example, a security manager estimates three components in the risk assessment analysis: 1) attribute weights, 2) frequency of the attacks, and 3) the consequences or outcomes of an attack. For frequency and outcomes, the analyst asks the security manager to provide an upper and lower bound around his or her expected estimate. These bounds reflect the degree of uncertainty about how frequently and attack might occur, and what the actual consequences might be. From the upper and lower bounds and the expected estimate, ASESS constructs a probability distribution for each threat. ASESS constructs a normal distribution curve for attack frequencies using the expected estimate as the mean, with the upper and lower bounds determining the standard deviation. For example, Figure 3-2 describes the normal distribution curve if the expected frequency of procedural violations is 4,380 attacks per year, and the security manager's estimated lower and upper bound are 3,500 and 5,000 attacks per year respectively.

Figure 3 - 2 Normal Probability Distribution for Procedural Violations

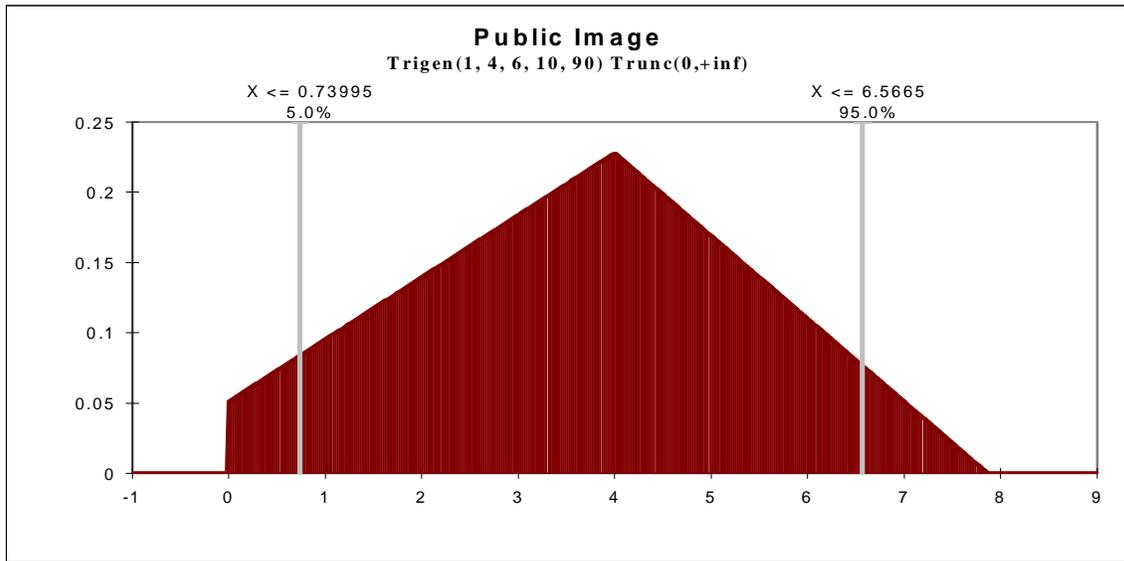


In this example, ASESS uses the lower bound to approximate two standard deviations from the mean (440 is one standard deviation). ASESS truncates normal distribution curves at zero. The minimum number of attacks can't be less than zero, but there can be an unlimited number of attacks in a year. The probability is low (less than 5%) that more attacks will occur than the security manager estimated as the upper bound.

ASESS constructs outcome attribute probability density functions slightly differently than it does the frequency density functions. ASESS constructs a *Triangular* probability density function using the security manager's "most likely" value as an estimate of the distribution's mode. ASESS uses the lower and upper bounds to establish the 10th and 90th percentile of the distribution. Figure 3 - 3 shows a *Triangular* probability density function for damaged public image with a most likely value of 4 and 6 and 2 as the upper and lower bounds respectively. Recall that the security manager rates the damage to an organization's public image on a scale from 1 to 7, but the upper bound is not truncated to ensure that extreme or "worst case" outcomes are possible in the sensitivity analysis. ASESS truncates the distribution at the lower bound of zero.

The third component that the security manager estimates is the weighted values of the attributes. Since there is uncertainty around the estimated attribute weights, ASESS also creates probability distributions for these values. ASESS creates Normal probability density functions using the estimated weights as the mean and 10% of the estimated weight as a standard deviation. Experiments using a slightly greater standard deviation did not change the risk assessment results.

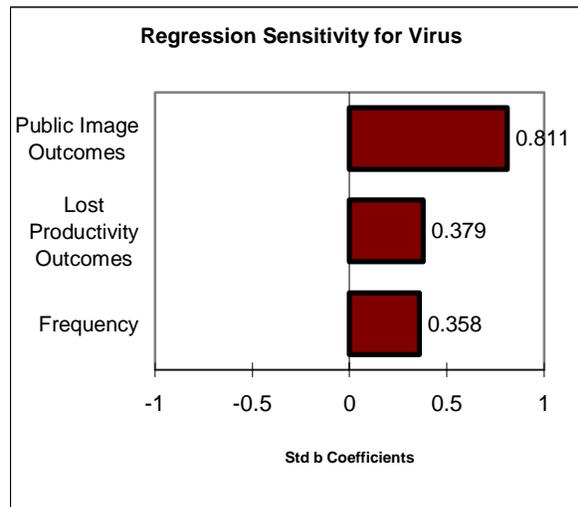
Figure 3 - 3 Triangular Distribution for Public Image



Once ASESS constructs the probability functions for each of the threat frequencies and outcome attributes, the analyst runs a simulation. The number of iterations can vary with the complexity of the problem, but 1,000 iterations are generally sufficient. Each @Risk iteration calculates the threat index for all threats by selecting values from the probability distributions for each threat. The @Risk tool computes the average threat index for each threat so that the multi-attribute analyst can rank the threats by their threat index. Simulations with a larger number of iterations, i.e., 3,000 iterations, did not change the results of the risk assessments.

Figure 3 - 4 Regression Tornado Graph for Virus

From the simulation results, the @Risk analysis tool produces linear regression *Tornado*¹¹ graphs. Tornado graphs give the multi-attribute analyst the Std b coefficient¹² for each variable in the additive value function. The Std b coefficient indicates which factors have the most influence on the threat's TI. For example, Figure 3 - 4 shows the Tornado graph for a *Virus* threat, which shows that a one standard deviation increase in Public Image Outcomes value increases the *Virus* Threat Index by .811 standard deviations, and a one standard deviation in the Lost Productivity Outcomes values increases the *Virus* Threat Index by .379



¹¹ Tornado graphs show either the correlation coefficients or the linear regression beta (Std b) coefficients in a bar format. The graph often looks like a tornado when there are both negative and positive coefficients.

¹² Std b coefficients are the variable coefficients that predict the TI.

standard deviations.

The TI for the organization's Virus threat shows that the Threat Index is most sensitive to the **Public Image Outcomes** attribute values. The *Virus* index is almost equally sensitive to the **Lost Productivity Outcomes** and the **Frequency** of attack. Linear regression coefficients are very useful in helping the analyst determine which inconsistencies are important between the security manager's initial ordering of risks and SAEM's computed rankings.

3.4 Summary

This chapter briefly described how multi-attribute analysis techniques provide a systematic and structured mechanism for developing a risk assessment. A multi-attribute analyst develops a risk assessment, or cost framework, by eliciting from the security manager the frequencies and outcomes from successful attacks. The analyst uses the results of the risk assessment to evaluate the benefits of security technologies so that the security manager can see how technologies affect the consequences (or attributes) of an attack. Chapter 4 will describe in detail how the risk assessment is used to evaluate and compare the benefits of various security technologies.

This risk assessment relies on an additive model, which allows the analyst to prioritize the organization's threats. The analyst is able to use an additive model because the consequences of the attacks are preferentially independent. Within the additive model, the additive value function computes a threat index for each threat which determines the relative significance of the threat to the organization. The additive value function uses the security manager's estimates about most-likely attack outcome values and expected frequencies to compute threat indexes. ASESS, a semi-automated decision analysis tool, models the security manager's uncertainty about his or her estimates so that an analyst can conduct sensitivity analysis, giving the security manager additional insight about the organization's threats and security technology benefits.

CHAPTER 4. Security Attribute Evaluation Method Process

4.1 Introduction

The Security Attribute Evaluation Method (SAEM) is a cost-benefit analysis process for analyzing security design decisions during the development or update of an organization's information system security architecture. SAEM consists of four steps: 1) a risk assessment, 2) a security technology benefit analysis, 3) a coverage analysis, and 4) a security-technology tradeoff analysis. These four steps help the lead security specialist select the security components and/or risk-mitigation strategies for the organization's security architecture. Although the participants in each step of SAEM vary among organizations, a multi-attribute analyst¹³ and the organization's lead security specialist are the key participants. A multi-attribute analyst facilitates each step of the process, eliciting from the lead security specialist and other participants their knowledge about the organization's risks, their expertise about the effectiveness of security technologies, and the key factors that managers use in selecting security technologies.

This chapter describes how SAEM supports the risk assessment and selection of security technologies during the security-architecture development process. In addition, this chapter describes the four steps of the SAEM process, the participants, and their role in each step of the SAEM process. The analyst uses an open protocol to elicit the organization's threats, outcomes, and security-selection objectives for the case studies in this thesis. Appendix A is an example of the protocol used by the analyst during the risk assessment. An open protocol allows the decision maker to select the outcome attributes.

4.2 Security Architecture Development

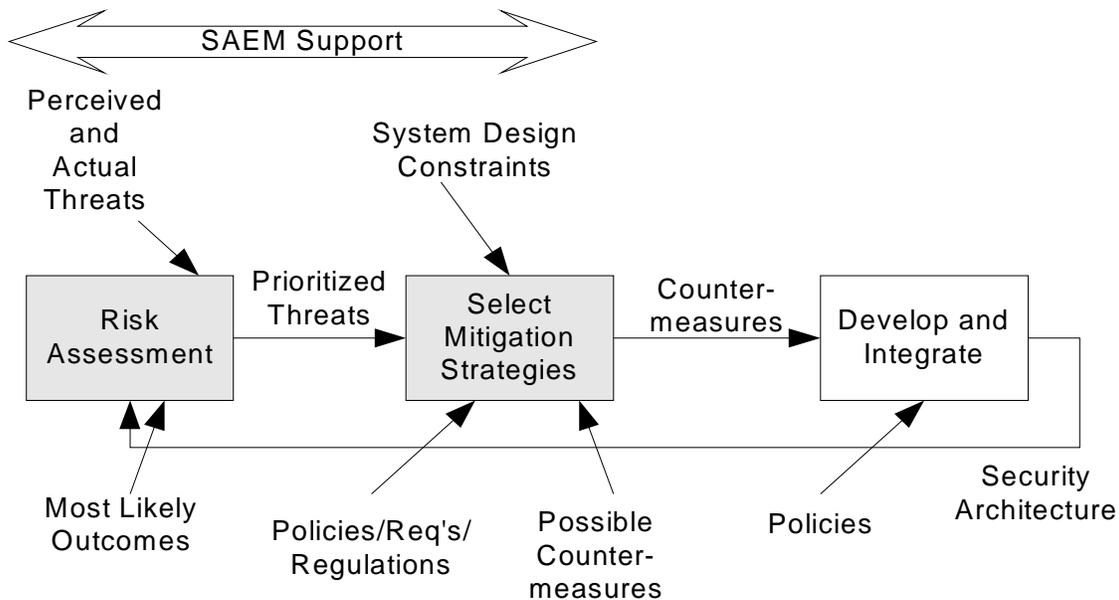
The security architecture of an information system consists of a collection of security technologies and procedures that satisfy the requirements of an organization's security policies (Ramachandran 2002). Best practice dictates that security architects develop a risk assessment, the first step of their risk management process, before selecting the security technologies (Fraser 1997) or risk-mitigation procedures for the security architecture. Security technologies and procedures implement an organization's security policies, but system designs and security regulations further constrain the security architect's selection of security technologies. Figure 4 - 1 shows the security architecture three-step development process, the information needed, and results at each step.

The security architect completes this process during system development and the organization's security manager revisits the process periodically to ensure that the security

¹³ Throughout this chapter any references to an analyst should be taken to mean a multi-attribute analyst

architecture is up to date with changes in the threat environment. The SAEM supports an organization’s security architect, engineer, or security manager in the development or periodic update of the information system’s security architecture. Specifically, SAEM is a flexible, systematic, and repeatable process that prioritizes threats and helps the security architect or manager select countermeasures, i.e., it assists with the first two steps of the security-architecture development process.

Figure 4 - 1 Security-Architecture Development Process

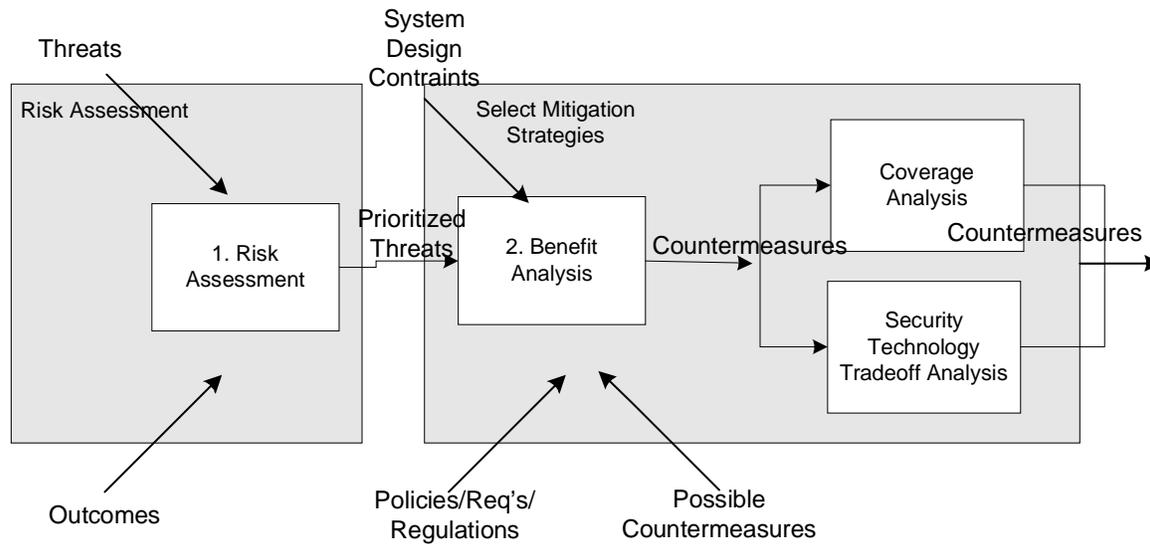


4.3 The Security Attribute Evaluation Method

The Security Attribute Evaluation Method supports the security architect and/or security manager in making security design decisions. More specifically, SAEM develops a *quantitative* risk assessment that security managers use to compare and select security technologies. In addition, one of the objectives of SAEM is to provide the security manager with insight into the selection of security technologies.

Each step of SAEM (Risk Assessment, Benefit Analysis, Coverage Analysis, and Security Tradeoff Analysis) consists of several steps. A multi-attribute analyst facilitates these steps by interviewing the security manager and senior information system managers, analyzing the results, and conducting sensitivity analyses. Since SAEM relies on the knowledge of an organization’s lead security specialist and security staff, the results of the method are dependent on their security expertise and experience with their organization’s threat and security environment. Figure 4 - 2 shows the SAEM process and the sequence of the four steps. SAEM does not replace all of the activities that are necessary for an organization to select security technologies, rather it supplements the Architecture development process. The background boxes depict the first two steps in the security architecture development process.

Figure 4 - 2 SAEM Steps within the First Two Steps of the Security Architecture Development Process



4.3.1 Automated Security Evaluation Support System (ASESS)

ASESS is a decision support tool prototype that I developed to help the analyst during each step of SAEM. It incorporates @Risk™, a commercially developed, risk-analysis tool from Palisades. The multi-attribute analyst uses ASESS to store an organization’s SAEM data, conduct simulations, rank the threat indexes and security technologies, and quickly show gaps via coverage analysis. The analyst and/or security manager can also use ASESS to conduct what-if analyses. Although ASESS is not necessary for SAEM analysis, it expedites the analysis.

4.3.2 Participants

The two key participants of SAEM are an organization’s information system lead security specialist or security manager and a multi-attribute analyst. Organizations rely on their security managers to make security design decisions and recommend risk-mitigation strategies for their information systems. Using an open protocol, the multi-attribute analyst prioritizes the information system risks, security countermeasures, and helps the security manager compare security technologies. The open protocol allows the security manager to address the organization’s specific risks and assess the technologies that he or she believes mitigate those risks.

Other participants in the SAEM process are information system managers, executives, and any other security personnel who may have a different or unique organizational perspective or specialized expertise. In the three case studies, the lead security managers relied on key members of their staffs for their security expertise, and senior information system managers reviewed and commented on the security manager’s input.

4.3.2.1 Lead Security Specialist Role

The lead security manager may be a dedicated security specialist, with information security expertise, responsible for the information system security architecture. In contrast, he or she may be an information system manager who has multiple responsibilities, one of which is security. Regardless of whether the security manager's job is full- or part-time, knowledge about information system risks and the effectiveness of countermeasures varies among security managers. For example, all the security managers who participated in the case studies had experienced viruses and denial-of-service attacks, so they were confident about their estimates of frequency of attack and consequences, but two security managers were less confident in estimating the frequency and potential consequences of compromising emanation attacks because they had never experienced that type of attack.

4.3.2.2 Multi-attribute Analyst Role

The role of the multi-attribute analyst is to facilitate the data collection during each step of SAEM, analyze the results, and help the security manager interpret the results. In addition to multi-attribute analysis techniques, the analyst must be familiar with information system security threats and security technologies. The analyst need not be a security expert, but such knowledge is helpful since the analyst helps clarify terms and definitions throughout the process.

4.3.2.3 Other Participants

Other participants in SAEM include information system managers or executives and members of the security staff. System managers and executives can help define the outcome attributes and weights. In two case studies, the information system managers used different outcome attributes and ranked the attributes differently than the security managers. Security personnel can assist the security manager with the attack frequency estimates, outcome values, and security technology effectiveness ratings.

4.4 Risk Assessments

Chapter 3 describes aspects of the risk assessment step, while this section details the sub-steps involved in conducting the risk assessment. The risk assessment is the first step of the SAEM and relies on the additive model to determine the relative ranking of an organization's threats. The risk assessment consists of six steps:

- 1) Determine the threats
- 2) Determine the attack outcome attributes
- 3) Elicit distribution of outcome attribute values
- 4) Weight the attributes
- 5) Compute and rank threats
- 6) Conduct sensitivity analysis and refine results.

During the risk assessment, the security manager determines which threats are potential risks to the organization and how often these threats result in an attack on the system. In addition, the security manager and/or information system manager determine the outcome attributes or which consequences are of most concern to the organization. Through a series of interviews and questionnaires, the analyst elicits the distribution of outcome attribute values. Next, the analyst asks the security manager to rank and assesses the attribute preferences to compute the threat indexes. Finally, the analyst reviews the results and follows-up with additional interviews to resolve discrepancies. The next six sections describe in more detail each step of SAEM, and a copy of a questionnaire used during the elicitation step of the risk assessment is attached as Appendix A.

4.4.1 Risk Assessment Step 1: Determine the Threats and Initial Risk Ordering

The first step in the risk assessment is for the security manager to identify which threats constitute a risk to the organization. The analyst presents the security manager with a set of cards, each of which contains an information-security threat and its definition. The security manager’s task is to identify which threats are potential risks to the organization. In addition, the security manager pre-sorts the cards into three piles: high-, medium- and low-risk threats. Finally, the security manager arranges the cards in complete or partial order from highest to lowest risk and eliminates any threats that are not a risk to the organization.

The analyst presents the pre-defined set of threats to help the security manager quickly identify the organization’s risks, but the security manager can further refine, redefine, add, or delete threats. For example, SAEM defines the Personal Computer Abuse threat as “The unauthorized use of information system assets for personal means (e.g., games, résumés, personal matters). One security manager found this definition too general; so, he further defined Personal Computer Abuse into five separate types. In contrast, another organization did not identify Personal Computer Abuse as an organization risk. Table 4 - 1 is a list of the initial set of threats and their definitions.

Table 4 - 1 Initial Threats and Their Definitions

| | THREAT | DEFINITION |
|---|-------------------------|---|
| 1 | Alteration | The modification, insertion, or deletion of data or lines of code, whether by an authorized user or not, that compromises the auditability, confidentiality, recoverability, availability, or integrity of the data or application. |
| 2 | Browsing | The unauthorized act of searching through electronic storage to locate or acquire information without necessarily knowing of the existence or the format of the information being sought. |
| 3 | Compromise | The unintentional release of information to someone not authorized access to the information. This includes information exempt from public disclosure, Privacy Act information, proprietary information, sensitive-unclassified information, and national security information. |
| 4 | Compromising Emanations | The unintentional release of data-related or intelligence-bearing signals that, if intercepted and analyzed, could disclose classified or sensitive-unclassified information being transmitted and/or processed. |
| 5 | Contamination | The intermixing of data of different sensitivity levels. |

| | THREAT | DEFINITION |
|----|--------------------------------|---|
| 6 | Cryptographic Compromise | Decryption of information or messages by an unintended recipient. |
| 7 | Data Entry Error | An error in the introduction of data that results in processing errors. |
| 8 | Denial of Service Attack (DoS) | A number of attacks, perpetrated from someone external to the system, designed to prevent access to system resources. |
| 9 | Distributed DoS Attacks | An attempt to use several legitimate, but unsuspecting, computers to coordinate a Denial of Service Attack against a different computer system. |
| 10 | Electronic Graffiti | Electronically defacing the public image by marking up home web pages. |
| 11 | Fraud/ Embezzlement | The deliberate, unauthorized manipulation of hardware, software, or data that could result in financial gain to the perpetrator. |
| 12 | IP Spoofing | Illegitimate attempts to enter the information system using an authorized or trusted IP address. |
| 13 | Logic Bomb | A resident computer program that triggers the perpetration of an unauthorized act when a particular state of the system is realized. |
| 14 | Message Stream Modification | The interception and modification of messages between two hosts. |
| 15 | Password Guessing | An automated or manual attempt to obtain user or system privileges by guessing the password. |
| 16 | Password Nabbing | Unauthorized capture of passwords. |
| 17 | Personal Computer Abuse | The unauthorized use of an information system asset for personal means (e.g., games, resumes, personal matters). |
| 18 | Procedural Violation | The unintentional violation of an established procedure or regulation. |
| 19 | Signal Interception | The unauthorized interception of communications between two hosts. |
| 20 | System Scanning | An attempt to detect vulnerabilities and weaknesses in a system. |
| 21 | Theft | The unauthorized taking of information for personal gain. |
| 22 | Trap Door | A hidden software or hardware mechanism that can be triggered to permit system protection mechanisms to be circumvented (Back Door). |
| 23 | Trojan Horse | A computer program with an apparently or actually useful function that contains additional (hidden) functions that surreptitiously exploit the legitimate authorizations of the invoking process to the detriment of security or integrity. |
| 24 | Vandalism | The malicious and generally motiveless destruction or damage to data or software. |
| 25 | Virus | A self-propagating software program composed of a mission component, a trigger component, and a self-propagating component. |
| 26 | WEB Page Spoofing | Deliberately misdirecting WEB page links to illegitimate sites. |

The initial list of security threats and their definitions was compiled from the risk assessments of several government and commercial organizations, trade journals, and automated security-news mailings. The security manager should not consider this initial list to be the definitive and complete list of all potential information-system risks to the organization. Security managers can, and did in some cases, identify threats not originally included in the threats presented. The important result of this step is that the security manager identifies the organization's risks regardless of any initial list. The starter list is helpful, but not essential, in identifying a wider range of threats than the security manager is likely to come up with unprompted.

4.4.2 Risk Assessment Step 2: Determine the Attack Outcome Attributes

The second step in the risk assessment is to determine the attack outcome attributes or consequences. Although most organizations' security managers quickly identify lost productivity and lost revenue as the two most obvious consequences of an attack, the analyst's job is to help the organization determine other significant outcome attributes of an attack. In addition to minimizing lost productivity and revenue, organizations can also be concerned about other consequences, such as damage to corporate image, inability to conduct key operations, and damage to customer relationships. Although the analyst interviews the security manager to determine the outcome attributes, it is also useful to interview senior- and executive-level managers since their concerns sometimes differ from those of the security manager. In fact, information system managers provided additional outcome attributes or changed the security manager's outcome attributes in all three case studies.

4.4.3 Risk Assessment Step 3: Elicit Expected Values

Once the analyst determines the organization's important attack outcome attributes, the next step in the risk assessment is to elicit the expected frequency and the most-likely consequences of an attack for each threat, i.e., the attribute outcome values. The analyst can elicit these values either through interviews with the security manager and security staff, or through a questionnaire that the security manager can fill out at his/her convenience. Appendix A gives an example of the questionnaire used during the industrial case study.

For each threat, the security manager provides an expected frequency. Expected threat frequencies are the security manager's best estimates of how often an attack of each threat would occur *if the system did not have any security countermeasures in place*. For example, most security managers estimate that their organizations would experience several hundred virus attacks each week if anti-virus software and other security countermeasures were not in place.

For each threat, the security manager also provides the most likely attack outcome attribute values. The most-likely attribute values are the consequences that the security manager would expect to see if an attack was successful, but not necessarily the average or mean consequence of the attack. For example, most security managers said they would most likely experience a few hours of lost productivity if a virus infiltrated the information system. In contrast, some viruses, such as the "I love you" virus, caused two of the case-study organizations thousands of hours of lost productivity. This extreme but rare event significantly changed the *average* cost of a virus attack for these organizations.

Recall that one of the key benefits of using multi-attribute analysis in the risk assessment is that the security manager can state the outcome values in non-economic terms. For example, in all case studies, the security managers describe lost productivity as the number of hours lost, but security managers also rated other types of outcomes, such as damage to public image, inability to conduct tax administration, etc., using a 1-7 Likert¹⁴ scale, with 1 meaning no or negligible damage and 7 meaning most severe damage. The assessed value function described in Chapter 3 normalizes and standardizes the values of different attributes so the analyst can compute a unified threat index.

4.4.4 Risk Assessment Step 4: Weight the Attributes

The next step in the risk assessment is to weight the outcome attributes. In this thesis, the analyst uses the swing-weight method described in Chapter 3 to assess the security manager's preference for each outcome attribute. The security manager must determine the order in which he or she would mitigate the consequences of an attack. After selecting the order, the security manager values each consequence relative to the others with the most important consequence (as determined by the ordering) receiving 100 points.

4.4.5 Risk Assessment Step 5: Compute and Rank Threats

In step 5, the analyst enters the organization's risk information into ASESS, including the attack outcome attributes, attribute values, and the assessed attribute weights into ASESS. ASESS creates probability density functions for each additive value function variable, i.e., threat frequency, threat outcome values, and attribute weights, to serve as inputs to a simulation model. From the simulation results, ASESS determines the relative ranking of each threat using the additive value function and threat index averages described in Chapter 3.

4.4.6 Risk Assessment Step 6: Conduct Sensitivity Analysis and Refine

Once ASESS computes the threat indexes, the analyst shows the security manager the results so that the analyst can investigate counter-intuitive results. The analyst and security manager can explore discrepancies between the security manager's initial ordering and SAEM results. The security manager can adjust the assessed values or the initial ordering if the results uncover errors or provide the manager with a clearer idea of the process. The security manager should avoid tweaking the input values to obtain preconceived results. The security manager may need to identify additional attributes if the initial set does not capture the organization's concerns.

One type of discrepancy that should be explored is when the manager's estimated values of threat indicate that the threat should be ranked lower or higher because similar threats have different estimates. For example, if the security manager estimates that denial of service attacks occur less frequently and most likely result in less damage than scanning attacks, then SAEM will rank denial of service attacks lower than scanning attacks. However, if the manager ranked denial of service attacks higher than scanning attacks, then the analyst would point out the inconsistency between the security manager's rank and the SAEM rank.

¹⁴ A Likert Scale is a one-dimensional, equal appearing scale.

In addition to exploring ranking discrepancies, the analyst could also conduct “what-if” analysis for the security manager. The security manager may want to revise some estimates to determine whether these new estimates affect the final threat prioritization. For example, there may have been disagreement among the organization’s security staff about how often an attack is likely to occur. The analyst could use ASSESS to determine how different estimates would affect the final risk assessment results.

4.5 Benefit Analysis

The benefit analysis step determines which security technologies provide the greatest risk mitigation. Benefit analysis uses the results of the risk assessment and consists of three sub-steps:

- 1) Threat/security technology mitigation identification
- 2) Effectiveness elicitation
- 3) Effectiveness computation

The result of the benefit analysis step is a prioritized list of security technologies, which ranks the security technology according to effectiveness given an organization’s risk assessment. Intuitively, this step provides a way to identify those security technologies that have the greatest effect in mitigating the organization’s risks. SAEM usually rates security technologies that mitigate the highest risk threats as more effective than those technologies that mitigate low risk threats, but SAEM could rate security technologies as highly effective if they reduce the risk from several moderately ranked threats.

4.5.1 Benefit Analysis Step 1: Mitigation Identification

In the first step of the benefit analysis step, for each threat, the security manager identifies all security technologies that he or she believes mitigate the consequences of an attack. The analyst facilitates the identification step by presenting the security manager with a set of cards that list security technologies. Table 4-2 shows a list of initial security technologies. The security manager should identify additional security technologies if they are not included in the list. The result of this step is a list of security technologies for each information security threat. When independent, objective research assesses the actual effectiveness of technologies, this step could include that information, and then the security manager would calibrate the effectiveness to reflect the organizational environment. Table 4-3 shows a hypothetical example of three threats and their corresponding risk-mitigating security technologies.

4.5.2 Benefit Analysis Step 2: Effectiveness Elicitation

The next step in the benefit analysis step is effectiveness elicitation. In this step, the analyst asks the security manager to estimate the effectiveness of each security technology given a threat. For example, if the security manager identified anti-virus products as a security technology that mitigates the risk of viruses, then the security manager might estimate that anti-virus software would stop 90% of incoming viruses given his beliefs about the organization’s ability to correctly configure and maintain the anti-virus product and the strength of the product. As with the threat attribute or consequence values, the security manager identifies an upper and lower value to establish a most likely range of effectiveness.

Figure 4-3 shows an example of the percentage effectiveness estimates for the same three threats and their risk-mitigating security technologies. Each technology shows the range of effectiveness as indicated by the line length, and the tick mark along the line indicates the security manager's most likely estimate of effectiveness

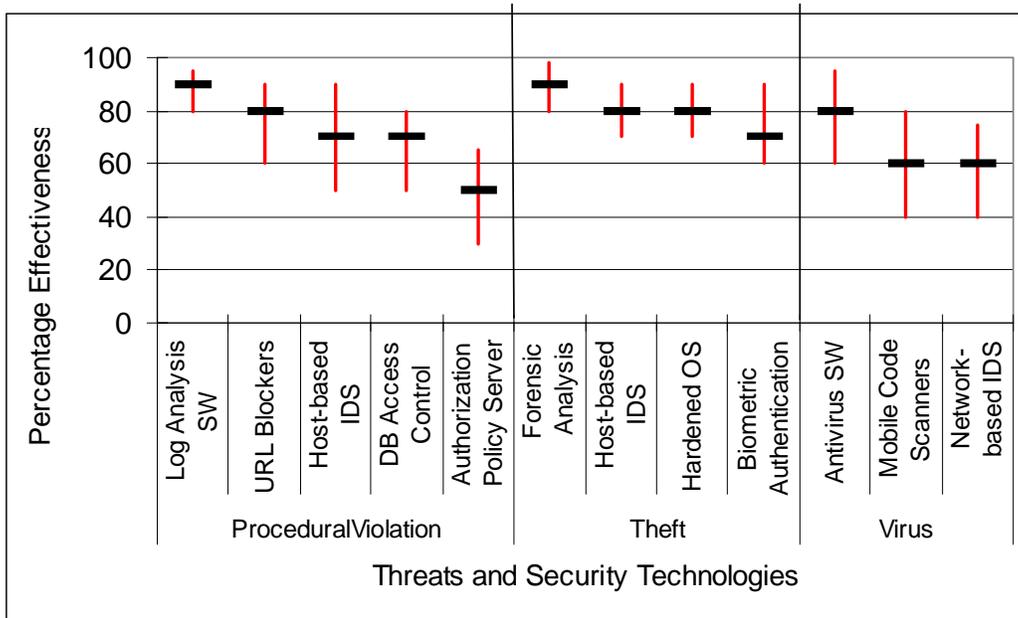
Table 4 - 2 Initial Security Technologies
(in alphabetical order)

| | | | |
|-----------------------------------|-----------------------|---------------------------|-------------------------|
| Anti-SPAM Filters | Electronic Signature | Modem Access Controls | Single Sign-On Apps |
| Antivirus Products | E-Mail Filters | Network-based IDS | Smart Card Products |
| Auditing Tools | Forensic Software | Network Monitoring Tools | Sniffer Detection |
| Authentication Tokens | Hardened OS | One-time Passwords | Software Lockout |
| Authorization Policy Servers | Hardware Lockdown | Packet Filter Firewalls | URL Blockers |
| Automatic Logout SW | Host-based IDS | Penetration Testing Tools | Virtual Private Network |
| Biometrics Authentication | Key Stroke Replicator | PKI/Cert Products | VLAN's |
| Centralized Security Management | Line Encryption | Proxy Firewalls | Web Access Control |
| Cryptographic Cards/Hardware | Load Balancers | Secure E-mail | |
| Database Security Access Controls | Log Analysis Software | Secure OS | |
| DB Encrypted Data Storage | Mobile Code Scanners | Secure User ID/Password | |

Table 4 - 3 Threats and Their Risk-Mitigation Technologies

| Threats | Technologies |
|-----------------------|---|
| Procedural Violations | Host-based IDS Database Security Access Control Authorization Policy Servers URL Blockers Log Analysis Software |
| Theft | Host-based IDS Forensic Software Hardened OS Biometrics |
| Virus | Antivirus Products Mobile Code Scanners Hardened OS Network-based IDS |

Figure 4 - 3 Effectiveness Estimates



4.5.3 Benefit Analysis Step 3: Effectiveness Computation

The final step in the benefit analysis step is to re-compute the threat index for all the threats, but with security-adjusted frequency and outcome values. When a security technology mitigates a threat, ASESS uses the security manager’s estimated effectiveness percentages to reduce the threat’s frequency of attack or the consequence estimates. Since a security technology can mitigate several threats, ASESS calculates a new total threat index for each security technology. Therefore, every security technology that mitigates as least one threat will have a total corresponding threat index less than the original risk assessment’s total threat index.

The analyst uses these new threat indexes to rank each security technology. The result of the benefit analysis step is a ranking of all security technologies according to their overall effectiveness against all threats. Intuitively, security technologies that are effective against the highest rated threats will rank higher than equally effective security technologies that are effective against lower ranked threats. For example, if a security manager estimates that a Host-based IDS is most likely 50% effective against Procedural Violations and Forensic Analysis is 50% effective against Theft threats, but Procedural Violations have a higher Threat Index than Theft, then Host-based IDS will rank higher than Forensic Analysis.

ASESS uses the security manager’s estimated effectiveness percentages to reduce the frequency of a threat and reduce the outcome of an attack. If the security technology stops an attack, then ASESS reduces the frequency by the security manager’s estimated effectiveness. If the security technology is effective in detecting the attack and/or helping the security manager recover from the damage or preventing further damage, then ASESS reduces the outcome values.

As with the outcome attributes in the risk assessment, ASESS creates a Triangular probability density distribution for each security technology/threat combination. For example, using Table 4-4, ASESS creates two probability density distributions for Host-based intrusion detection mechanisms (IDS), i.e., one for Host-based IDS/Procedural violations and one for Host-based IDS/Theft.

Using ASESS, the analyst runs another simulation, and ASESS computes the average change in total threat index that corresponds to each security technology. For example, the total threat index for the three threats (Procedural Violations, Theft, and Viruses) presented in Chapter 3 was 463. Table 4-5 shows the impact on the risk assessment's threat indexes of four different security technologies. In this case, host-based IDS provides the greatest reduction in total threat index. Individual threat index changes are highlighted in bold. In this simple example, the security manager's estimated effectiveness reduced the frequency as a result of each security technology, but in the case-study analyses, the outcomes were reduced when addressed by security technologies that detect intrusions or help the security manager recover from an attack.

Table 4 - 4 Changes in Threat Index Caused by Security Technology

| | | Security Technologies | | | | |
|---------|----------------------|-------------------------|-------------|------------|----------------|-------------------|
| | | Initial Risk Assessment | Hardened OS | AntiVirus | Host Based IDS | DB Access Control |
| Threats | Procedural Violation | 376 | NC | NC | -113 | -113 |
| | Theft | 6.8 | -5.4 | NC | -5.4 | NC |
| | Virus | 80 | NC | -64 | NC | NC |
| | Total | 463.8 | -5.4 | -64 | -118.4 | -113 |

4.6 Coverage Analysis

Coverage analysis evaluates the security manager's overall mitigation of information-system threats. Coverage analysis is dependent on the security technologies identified during the first step of benefit analysis step. Coverage analysis is based on the military's defense-in-depth strategy (Stoneburner, Hayden et al. 2001), which encourages military commanders to have multiple lines of defense to ensure that if the enemy penetrates one line of defense, the military unit is not completely vulnerable at the point of penetration. Since security managers cannot depend on any single security measure to ensure complete risk mitigation, they should rely on a combination of different types of countermeasures to reduce their vulnerability to attack.

The analyst uses coverage analysis to show the security manager how the current security architecture *protects* against potential attacks, *detects* unauthorized intrusions, and helps the staff *recover* from successful attacks. In addition, coverage analysis can show general gaps in the security architecture and allows the security manager to see how new security technologies fit into the security architecture. Therefore, coverage analysis depicts the role existing security technologies play in mitigating attacks, i.e., protection, detection, or recovery roles.

The analyst categorizes every security technology according to its primary function in the security architecture. Security technologies that prevent an attack from succeeding are *protection* mechanisms. Examples of protection mechanisms are firewalls, antivirus software, scanning packages, and authentication mechanisms. Although some of these security technologies detect an attack first, their primary function is to stop or prevent unauthorized access. *Detection* technologies alert security personal that an attack is in progress. Host-based intrusion detection mechanisms and network management applications are examples of detection security technologies. Finally, if an attack has occurred, *recovery* mechanisms help the security manager determine the amount of damage, restore system integrity, and possibly discover the perpetrator. Table 4-5 is a list of the security technologies and their classifications as a protection, detection, or recovery mechanism. Although the analyst classifies each technology, the security manager may change the classifications as needed.

Table 4 - 5 Security Technology Role Classification

| | Protection | | Detection | Recovery |
|-----------------------------------|---------------------------|-------------------------|--------------------------|----------------------------|
| Anti-SPAM Filters | E-Mail Filters | Secure E-mail | Centralized Security | Auditing Software |
| Antivirus Products | Hardened OS or Secure OS | Secure OS | Host-Based IDS | Back-up and Recovery Tools |
| Auditing Tools | Hardware Lockdown | Secure User ID/Password | Network Monitoring Tools | Forensic Software |
| Authentication Tokens | Line Encryption | Single Sign-On Apps | Network-Based IDS | Key Stroke Replicator |
| Authorization Policy Servers | Mobile Code Scanners | Smart Card Products | Packet Sniffers | Load Balancers |
| Automatic Logout SW | Modem Access Controls | Software Lockout | | Log Analysis Software |
| Biometrics Authentication | One-time Passwords | SPAM Filters | | |
| Cryptographic Cards/Hardware | Packet Filter Firewalls | URL Blockers | | |
| Database Security Access Controls | Penetration Testing Tools | Virtual Private Network | | |
| DB Encrypted Data Storage | PKI/Cert Products | VLAN's | | |
| Electronic Signature | Proxy Firewalls | Web Access Control | | |

In the first step of the coverage analysis step of SAEM, the security manager identifies the technologies that exist in the organization's current security architecture. Next, the analyst selects several threats, usually the top threats, and constructs a coverage model depicting the security technologies that mitigate selected threats (that the security manager had identified these in the first step of the benefit analysis) and that exist in the organization's security architecture.

Continuing with the hypothetical example from the benefit analysis Section, Figure 4-4 shows an example of a coverage model with the three threats and the security technologies placed appropriately in each ring of the model. For this example, assume that the security architecture included all of the mitigation technologies identified for each of the threats in

Figure 4 - 4, except for host-based intrusion detection. In addition, darker shading within a ring indicates more coverage than lighter shading; so white means there are no security technologies in the organization that fulfill that role. From this example, it is clear that the organization does not have many detection mechanisms to help the security manager determine the presence of an ongoing threat. Since the security manager identified host-based intrusion detection systems as capable of mitigating procedural violations, adding this technology to the model could show how it fills the detection gap, as shown in Figure 4 - 5. Security managers can compare alternatives by populating the model with the alternatives.

Figure 4 - 4 Coverage Analysis Model

Generally, security managers want to have at least one security technology in every ring for each threat, but that might not be either economically feasible or the best mitigation strategy. Another strategy might be to increase the defense against a highly probable or damaging threat, leaving gaps against other, less likely threats. Regardless of the security manager's risk-mitigation strategy, the coverage model provides an overview of the depth and breadth of defense that security technologies provide.

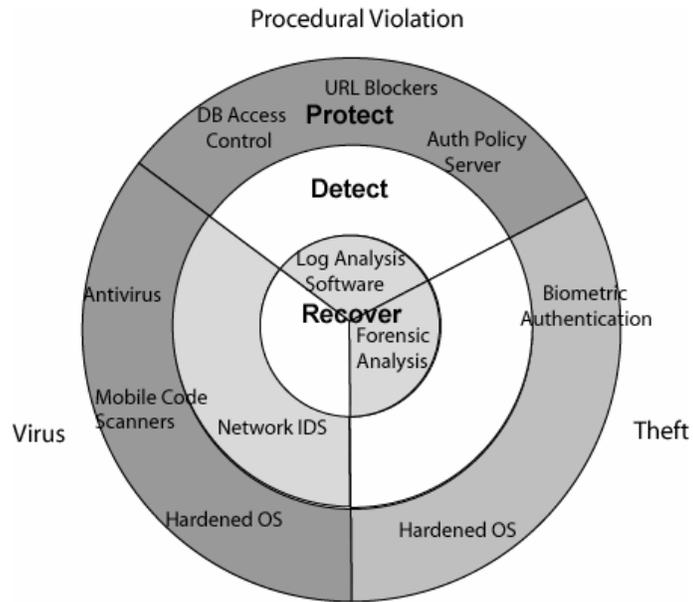
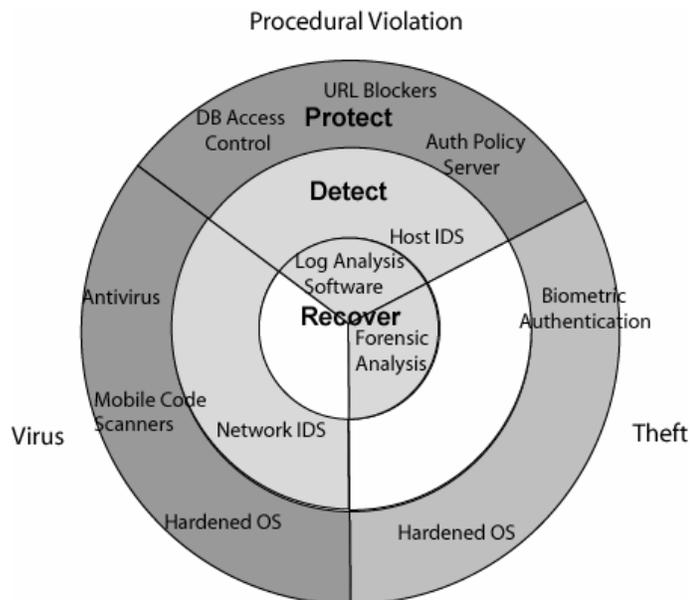


Figure 4 - 5 Coverage Analysis Model with Host IDS



4.7 Security Technology Tradeoff Analysis

Although the benefit analysis step determines which security technologies are the most effective in mitigating threats, security managers must consider purchase cost, maintenance, skill level requirements, false positives, etc. before selecting a technology for inclusion in the security architecture. In fact, the security manager often considers purchase cost or maintenance to be more important than the effectiveness of the technology. The security-tradeoff analysis helps the security managers compare security technologies using multi-attribute analysis techniques to rank each security technology according to the organization's decision objectives.

In the security technology tradeoff analysis step of SAEM, the security manager identifies and ranks four or five technologies and decides which factors are important when selecting security technologies. The analyst constructs an additive model, which ranks each of the selected technologies according to these decision factors. In the tradeoff step, the decision factors are the additive value function's attributes and the analyst weights the attributes (w) using the swing-weight method described in Chapter 3. The next section describes a simple example of how the analyst compares three different security technologies using a multi-attribute, additive-value function.

4.7.1 Tradeoff Analysis Example

Assume that the security manager wants to compare three technologies: host intrusion detection mechanisms, biometric authentication, and log analysis software. Also, assume that the security manager has identified purchase cost, maintenance, and threat effectiveness as the three most important factors (i.e., attributes) to be considered before selecting a security technology, and furthermore, that these factors meet the additive model assumptions of independence. Table 4 - 6 shows how the security manager might rank each technology against each of the decisions factors.

Table 4 - 6 Tradeoff Analysis

| | | Security Technologies | | | | | |
|------------|---------------------------------|-----------------------|--------|--------------------------|--------|-----------------------|--------|
| | | Host-based IDS | | Biometric Authentication | | Log Analysis Software | |
| | | Rank | Weight | Rank | Weight | Rank | Weight |
| Attributes | Purchase Cost $w=.43$ | 33 | .18 | 100 | .77 | 50 | .29 |
| | Maintenance $w=.21$ | 50 | .27 | 20 | .15 | 100 | .59 |
| | Threat Effectiveness $w=.36$ | 100 | .55 | 10 | .08 | 20 | .13 |
| | Technology Total | 183 | 1 | 130 | 1 | 170 | 1 |

In Table 4-6, the analyst computes the security technology rating using the following equation:

$$Technology\ Index_i = \sum_{j=attributes} w_j * v_j(s_{ij})$$

Where $Technology\ Index_i$ is the rank of the i^{th} security technology, w_j is the selection attribute weight and s_{ij} is the selection attribute value for the j^{th} security technology. The value function is assumed to be linear in this example, and is similar to the one used in the risk assessment step, i.e. s_{ij}/s_j^* where s_j^* is the maximum value for the j^{th} attribute. In this example, host-based intrusion detection ranked the highest (.42). As with the risk assessment and benefit analysis steps, the analyst can conduct sensitivity analysis to determine how sensitive the results are to the security manager's estimate.

4.8 Summary

The four steps of SAEM help the security manager determine the organization's threat priorities and the security technologies that best reduce the organization's risks. In addition, the process gives the security manager insight about how his or her threat assumptions affect the organization's selection of security technologies. The security manager can explore the assumptions during the sensitivity analysis. The benefit analysis helps the security manager determine which security technologies are candidates for weaknesses in the security architecture, but the coverage analysis helps the security manager determine those weaknesses. Together, the benefit and coverage analyses provide a visual model for evaluating the organization's security architecture and comparing security technology alternatives. Finally, the security tradeoff analysis step helps the security manager see how selection objectives other than effectiveness affect the selection of a security technology.

CHAPTER 5. Commercial Case Study

5.1 Introduction

This chapter describes the participants, activities, and results from the commercial case study. Serving as the multi-attribute analyst, I elicited all of the data for this case study through a series of interviews conducted during a one-week period. The Global Security Architect (Architect) and two technical advisors provided the input for each SAEM phase, but the Architect made the final decisions during each refinement step. The risk assessment highlighted the organization's conflict between an open and trusting work environment and the risk of being too lenient with security policies. The benefit analysis showed that some technologies that could help reduce risk from virus attacks had been overlooked because they would cause the organization to enforce stricter security policies. The coverage evaluation showed the security architecture was weak in detection mechanisms. Finally, the tradeoff analysis showed mixed results because the process assumes none of the security technologies that the Architect selects for comparison exist in the security architecture. Overall, the Architect reported (Satisfaction Survey at end of Chapter) that the analysis was insightful and helpful in developing the organization's security strategies.

5.2 Case Study Description

The first case study is from a large commercial organization (\$4 Billion/year in Revenue) that has four information system (IS) departments throughout the world. In addition to developing the organization's global security architecture, corporate security personnel must ensure that the security architecture is consistent across IS departments. Since some of the IS departments are located in foreign countries, U.S. software and cryptographic export restrictions constrain the selection of security technologies. Although the company is large, the overall IS security budget is limited (less than \$1 million).

The corporate culture is very relaxed and informal. Policies governing use of corporate computing resources for personal use, such as personal email, games, and web browsing, are not very restrictive or tightly enforced. At the time of the interviews, the corporate security office was focused on reducing the risk from virus attacks, because the company experienced significant operational disruptions from the "I love you" virus. However, the Architect thought that internal threats posed the greatest portion of security risks for the organization. Since the organization did not track security incidents, operate an incident response center, or have an existing risk assessment, the security staff made all estimates based on subjective assessment.

5.2.1 Case Study Participants

The Architect was the organization's primary participant in this case study. The Architect was responsible for the corporation's overall security architecture, developing security requirements, recommending changes to the security architecture, such as the purchase of new

security technologies; managing security policy changes; and preparing the organization’s security budget. Two individuals assisted the Architect by providing technical advice and expertise about threats and countermeasures. These individuals also participated in the case study, but were not available for all phases. When disagreements about the inputs occurred among the participants, the Architect resolved the disputes and provided the final input value to the model. Collectively, the participants were familiar with most of the threats and security technologies presented, but didn’t have much direct experience with them. Finally, I served as the multi-attribute analyst for this case study.

5.3 The Risk Assessment

This section outlines the details, sequence of events, and data collected during the risk assessment phase of the case study. The Architect and two technical advisors provided the threat frequency and outcome data during one interview session, which lasted approximately four hours. The next day the analyst presented the SAEM results, and the Architect made revisions to her initial risk priorities and the risk assessment input data.

5.3.1 Initial Iteration

Table 5 - 1 shows the Architect’s initial ranking of threats. The Architect added *Internal Vandalism* (ranked 22nd in Table 5 - 1) to the analyst’s initial threat list. The organization had recently experienced a few internal security compromises and wanted to differentiate between vandalism committed by employees and vandalism from outside the organization.

5.3.2 Outcome Attributes

The Architect identified three outcomes that were important to the organization: Damaged Public Image, Damaged Customer Relationships, and Lost Revenue. The company has a widely recognizable public logo and positive reputation, so they were sensitive to any security compromise that could damage their public image. Similarly, they had well-established and strong relationships with other corporations and spokespersons that were

Table 5 - 1 Initial Threat Rankings

| Rank | Threat |
|------|---------------------------------------|
| 1 | Virus |
| 2 | System Scanning |
| 3 | Distributed Denial of Service Attacks |
| 4 | Trojan Horse |
| 5 | Denial of Service Attacks |
| 6 | WEB Page Spoofing |
| 7 | Compromise |
| 8 | Contamination |
| 9 | Alteration |
| 10 | Theft |
| 11 | Compromising Emanations |
| 12 | Trap Door |
| 13 | Logic Bomb |
| 14 | Password Guessing |
| 15 | Procedural Violation |
| 16 | Vandalism |
| 17 | Electronic Graffiti |
| 18 | Data Entry Error |
| 19 | Browsing |
| 20 | Personal Computer Abuse |
| 21 | Fraud/Embezzlement |
| 22 | Internal Vandalism |
| 23 | Message Stream Modification |
| 24 | Signal Interception |
| 25 | Password Nabbing |
| 26 | IP Spoofing |
| 27 | Cryptographic Compromise |

affiliated with the corporation and thus, could be affected by a security compromise. Finally, the company wanted to mitigate security compromises that could cause lost revenue.

The participants used the 7-point scale described in Chapter 4 to assess the impact of an attack on the outcome attributes Damage to Public Image and Customer Relationships. Since many security compromises result in lost employee productivity, rather than a direct revenue cost, the participants chose to convert lost productivity into lost revenue. The group used a \$100/hour cost to compute the lost revenue value. For example, if an attack resulted in 2 hours of lost productivity, then the lost revenue was \$200. Table 5 - 2 represents the results of the security manager's rating of the outcome attributes using the swing-weight method as described in Chapter 3.

Table 5 - 2 Outcome Attribute Weights

| Attribute | Rating | Weight |
|--------------------|--------|--------|
| Public Image | 100 | 0.38 |
| Customer Relations | 85 | 0.33 |
| Lost Revenue | 75 | 0.29 |

5.3.3 Outcome and Frequency Estimates

Table 5 - 3 shows all of the Architect's estimates for attack outcome values and frequencies. The table is organized according to average frequency with the most frequently occurring attacks listed first. Notice that several internal threats, such as *Contamination*, *Personal Computer Abuse*, and *Procedural Violations*, are listed first. In addition, the Architect believed these internal attacks would most likely result in no damage to the company's Public Image or Customer Relationships. Table 5 - 3 also shows that *Theft* is the most significant threat in terms of Lost Revenue when compared to other threats.

The process of estimating the outcome values and frequencies can be tedious for the participants and they can quickly forget how they arrived at with many of their estimates. In developing the estimates, the security staff participants either relied on specific security compromises that the organization had previously experienced, or they developed scenarios that helped them arrive at their estimates. The process of capturing scenarios and recalled experiences helped the participants maintain consistency through iterations of the risk assessment, since the analyst recorded the justification for their estimates. Participants could refer back to these rationales for their estimates and re-evaluate the scenarios rather than make significant changes because they had forgotten the reason for the original estimate.

TABLE 5 - 3 INITIAL ESTIMATED OUTCOME AND FREQUENCY VALUES

| Threats | Frequency/Year | | | Lost Revenue (\$\$) | | | Damaged Public Image | | | Customer Relationships | | |
|--------------------------|----------------|---------|---------|---------------------|-------------|-----------|----------------------|-------------|------|------------------------|-------------|------|
| | Low | Average | High | Low | Most Likely | High | Low | Most Likely | High | Low | Most Likely | High |
| Contamination | 10,000 | 365,000 | 6.6m | 0 | 0 | 7,500 | 1 | 1 | 1 | 1 | 2 | 3 |
| Personal Computer Abuse | 182,500 | 365,000 | 730,000 | 1 | 2 | 5 | 1 | 1 | 2 | 1 | 1 | 2 |
| Procedural Violation | 36,000 | 72,000 | 120,000 | 0 | 0 | 25,000 | 1 | 1 | 2 | 1 | 1 | 2 |
| Virus | 2,190 | 4,380 | 7,300 | 2,000 | 4,000 | 250,000 | 1 | 1 | 2 | 1 | 1 | 2 |
| Password Guessing | 730 | 1,825 | 7,300 | 0 | 0 | 8,000 | 1 | 1 | 2 | 1 | 1 | 2 |
| Alteration | 156 | 730 | 2,920 | 0 | 2,000 | 100,000 | 1 | 2 | 5 | 1 | 3 | 7 |
| Compromise | 1 | 52 | 365 | 0 | 0 | 30,000 | 1 | 2 | 5 | 1 | 1 | 4 |
| Signal Interception | 40 | 50 | 200 | 0 | 0 | 12,000 | 1 | 1 | 2 | 1 | 1 | 2 |
| Scanning | 6 | 12 | 48 | 0 | 600 | 2400 | 1 | 1 | 2 | 1 | 1 | 2 |
| Web Page Spoofing | 2 | 10 | 20 | 5,000 | 10,000 | 25,000 | 2 | 4 | 6 | 1 | 3 | 5 |
| Data Entry Error | 2 | 6 | 12 | 0 | 0 | 10,000 | 1 | 3 | 5 | 1 | 3 | 5 |
| Internal Vandalism | 2 | 6 | 12 | 100 | 1,000 | 4,000 | 1 | 2 | 4 | 1 | 2 | 4 |
| Fraud/Embezzlement | 2 | 5 | 50 | 25 | 50 | 10,000 | 1 | 2 | 5 | 1 | 2 | 6 |
| Password Nabbing | 1 | 4 | 6 | 0 | 0 | 8,000 | 1 | 1 | 2 | 1 | 1 | 2 |
| Compromising Emanations | 1 | 5 | 10 | 0 | 0 | 35,000 | 1 | 1 | 2 | 1 | 1 | 2 |
| Cryptographic Compromise | 1 | 5 | 10 | 0 | 0 | 35,000 | 1 | 1 | 2 | 1 | 1 | 2 |
| Trojan Horse | 2 | 4 | 10 | 0 | 0 | 8,000 | 1 | 1 | 2 | 1 | 1 | 2 |
| Denial of Service | 0 | 3 | 10 | 0 | 0 | 35,000 | 1 | 2 | 5 | 1 | 2 | 4 |
| Theft | 1 | 2 | 5 | 25,000 | 500,000 | 2,000,000 | 2 | 4 | 5 | 2 | 3 | 4 |
| Vandalism | 0.5 | 1 | 3 | 0 | 1,000 | 25,000 | 2 | 3 | 5 | 2 | 3 | 5 |
| IP Spoofing | 0.5 | 1 | 3 | 0 | 0 | 500 | 1 | 1 | 2 | 1 | 1 | 2 |
| Browsing | 1 | 2 | 4 | 0 | 200 | 1,500 | 1 | 1 | 3 | 1 | 1 | 4 |
| Electronic Graffiti | 0.5 | 1 | 3 | 0 | 1,000 | 2,000 | 3 | 4 | 6 | 2 | 3 | 5 |
| DDoS | 0 | 1 | 2 | 0 | 0 | 70,000 | 1 | 3 | 6 | 1 | 3 | 5 |
| Logic Bomb | 0.33 | 0.5 | 1 | 100 | 1,000 | 4,000 | 1 | 1 | 2 | 1 | 1 | 2 |
| Trap Door | 0.2 | 0.5 | 2 | 0 | 0 | 8,000 | 1 | 1 | 2 | 1 | 1 | 2 |
| Message Stream Mod | 0.2 | 0.33 | 0.5 | 25 | 50 | 1,000 | 1 | 1 | 3 | 1 | 1 | 4 |

Table 5 - 4 SAEM Threat Rankings

5.3.4 Initial Results

Table 5 - 4 compares the results from SAEM with the Architect’s initial estimated values. Differences in rankings occur because SAEM ranked high-frequency, but low-impact, threats higher than the Architect ranked them. This lead to, SAEM ranking internal threats highest (i.e., *Personal Computer Abuse*, *Procedural Violations*, and *Contamination*). The correlation between the SAEM Rank and the Architect’s initial rank is 0.2, a weak relationship between the two ranks.

5.3.5 Refinement

The Architect reviewed the results of SAEM and made some revisions. Since SAEM ranked high-frequency threats as the most significant based on the relative threat indexes, the Architect looked more closely at her internal threat frequency estimates and her initial ranking of threats. The Architect revised her estimates for 18 of the 27 threats, leaving nine unchanged. In some cases, such as *Contamination* and *Personal Computer Abuse*, the frequency changes were large. In contrast, the Architect only slightly changed the frequency of *Theft*, but did change her estimates of Lost Revenue outcome. Table 5 - 5 shows the Architect’s revised estimates, most of which are frequency revisions.

| Threat | Relative Threat Index | SAEM Rank (S) | Initial Rank (I) | S - I |
|-------------------------------|-----------------------|---------------|------------------|-------|
| Personal Computer Abuse | 100 | 1 | 20 | 19 |
| Procedural Violation | 23 | 2 | 15 | 13 |
| Virus | 2.0 | 3 | 1 | 2 |
| Contamination | 1.9 | 4 | 8 | 4 |
| Alteration | 0.9 | 5 | 9 | 4 |
| Denial of Service Attack | 0.9 | 6 | 5 | 1 |
| Password Guessing | 0.5 | 7 | 14 | 7 |
| Password Nabbing | 0.4 | 8 | 25 | 17 |
| Browsing | 0.4 | 9 | 19 | 10 |
| Compromise | 0.1 | 10 | 7 | 3 |
| WEB Page Spoofing | 0.02 | 11 | 6 | 5 |
| Signal Interception | 0.01 | 12 | 24 | 12 |
| Theft | 0.01 | 13 | 10 | 3 |
| Data Entry Error | 0.008 | 14 | 18 | 4 |
| System Scanning | 0.007 | 15 | 2 | 13 |
| Fraud/Embezzlement | 0.007 | 16 | 21 | 5 |
| Internal Vandalism | 0.006 | 17 | 22 | 5 |
| Electronic Graffiti | 0.002 | 18 | 17 | 1 |
| Distributed Denial of Service | 0.002 | 19 | 3 | 16 |
| Compromising Emanations | 0.002 | 20 | 11 | 9 |
| Cryptographic Compromise | 0.002 | 21 | 27 | 6 |
| Vandalism | 0.002 | 22 | 16 | 6 |
| Trojan Horse | 0.001 | 23 | 4 | 19 |
| IP Spoofing | <0.001 | 24 | 26 | 2 |
| Message Stream Modification | <0.001 | 25 | 23 | 2 |
| Trap Door | <0.001 | 26 | 12 | 14 |
| Logic Bomb | <0.001 | 27 | 13 | 14 |

Before changing her own ranking of threats, the Architect reconsidered the outcome attribute values. For many of the threats, the revised frequency estimates were significant changes from the original estimates. In addition, the Architect decided to revise the threat definitions to include only those security incidents that were intentional and malicious. For example, there were very few *Personal Computer Abuse* incidents that really concerned the security staff. Table 5-6 shows the change for each threat.

Table 5 - 5 Architect's Revised Estimates (sorted by frequency)

| Threats | Frequency/Year | | | Lost Revenue (\$\$) | | | Damaged Public Image | | | Customer Relationships | | |
|-------------------------------|----------------|-------|-------|---------------------|-------------|---------|----------------------|-------------|------|------------------------|-------------|------|
| | Low | Exp | High | Low | Most Likely | High | Low | Most Likely | High | Low | Most Likely | High |
| Virus | 2,190 | 4,380 | 7,300 | 2,000 | 4,000 | 250,000 | 1 | 1 | 2 | 1 | 1 | 2 |
| Password Guessing | 4 | 24 | 60 | 0 | 0 | 8,000 | 1 | 1 | 2 | 1 | 1 | 2 |
| System Scanning | 6 | 12 | 48 | 0 | 600 | 2,400 | 1 | 1 | 2 | 1 | 1 | 2 |
| Compromise | 2 | 12 | 36 | 0 | 0 | 30,000 | 1 | 2 | 5 | 1 | 1 | 4 |
| Signal Interception | 4 | 5 | 10 | 0 | 0 | 12,000 | 1 | 1 | 2 | 1 | 1 | 2 |
| Alteration | 1 | 4 | 6 | 0 | 2,000 | 100,000 | 1 | 2 | 5 | 1 | 3 | 7 |
| Internal Vandalism | 2 | 4 | 6 | 100 | 1,000 | 4,000 | 1 | 2 | 4 | 1 | 2 | 4 |
| Denial of Service Attack | 0 | 3 | 6 | 0 | 0 | 35,000 | 1 | 2 | 5 | 1 | 2 | 4 |
| Compromising Emanations | 1 | 3 | 5 | 0 | 0 | 35,000 | 1 | 1 | 2 | 1 | 1 | 2 |
| Trojan Horse | 1 | 2 | 5 | 0 | 0 | 8,000 | 1 | 1 | 2 | 1 | 1 | 2 |
| Procedural Violation | 1 | 2 | 4 | 3,000 | 10,000 | 25,000 | 1 | 1 | 2 | 1 | 1 | 2 |
| Browsing | 1 | 2 | 4 | 0 | 200 | 1,500 | 1 | 1 | 3 | 1 | 1 | 4 |
| Personal Computer Abuse | 1 | 2 | 4 | 1 | 2 | 5 | 1 | 1 | 2 | 1 | 1 | 2 |
| Contamination | 0.5 | 2 | 3 | 100 | 5,000 | 10,000 | 1 | 2 | 5 | 1 | 2 | 6 |
| Vandalism | 0.5 | 1 | 3 | 0 | 1,000 | 25,000 | 2 | 3 | 5 | 2 | 3 | 5 |
| Data Entry Error | 0.5 | 1 | 3 | 0 | 0 | 10,000 | 1 | 2 | 3 | 1 | 3 | 5 |
| Password Nabbing | 0.5 | 1 | 3 | 0 | 0 | 8,000 | 1 | 1 | 2 | 1 | 1 | 2 |
| WEB Page Spoofing | 0.5 | 1 | 3 | 1,000 | 4,000 | 8,000 | 2 | 4 | 6 | 1 | 3 | 5 |
| Electronic Graffiti | 0.5 | 1 | 3 | 0 | 1,000 | 2,000 | 3 | 4 | 6 | 2 | 3 | 5 |
| IP Spoofing | 0.5 | 1 | 3 | 0 | 0 | 500 | 1 | 1 | 2 | 1 | 1 | 2 |
| Distributed Denial of Service | 0 | 1 | 2 | 0 | 0 | 70,000 | 1 | 3 | 6 | 1 | 3 | 5 |
| Theft | 0.5 | 1 | 2 | 250 | 5,000 | 20,000 | 2 | 4 | 5 | 2 | 3 | 4 |
| Trap Door | 0.2 | 0.5 | 2 | 0 | 0 | 8,000 | 1 | 1 | 2 | 1 | 1 | 2 |
| Cryptographic Compromise | 0.2 | 0.5 | 1 | 0 | 0 | 35,000 | 1 | 1 | 2 | 1 | 1 | 2 |
| Fraud/Embezzlement | 0.2 | 0.5 | 1 | 100 | 5,000 | 10,000 | 1 | 2 | 5 | 1 | 2 | 6 |
| Logic Bomb | 0.3 | 0.5 | 1 | 100 | 1,000 | 4,000 | 1 | 1 | 2 | 1 | 1 | 2 |
| Message Stream Modification | 0.2 | 0.3 | 0.5 | 25 | 50 | 1,000 | 1 | 1 | 3 | 1 | 1 | 4 |

Table 5 - 6 Architect's Estimate Change [Architect's Revised Estimates – Initial Estimates]
 [A blank cell indicates value did not change from the initial estimate](Ordered by Threat)

| Threats | Frequency/Year | | | Lost Revenue (\$\$) | | | Damaged Public Image | | | Customer Relationships | | |
|-------------------------------|----------------|-------------|------------|---------------------|-------------|------------|----------------------|-------------|------|------------------------|-----|------|
| | Low | Most Likely | High | Low | Most Likely | High | Low | Most Likely | High | Low | Exp | High |
| Alteration | -155 | -726 | -2,914 | | | | | | | | | |
| Browsing | | | | | | | | | | | | |
| Compromise | 1 | -40 | -329 | | | | | | | | | |
| Compromising Emanations | | -2 | -5 | | | | | | | | | |
| Contamination | -9,999.5 | -364,998 | -5,999,997 | 100 | 5,000 | 2,500 | | 1 | 4 | | | 3 |
| Cryptographic Compromise | -0.8 | -4.5 | -9 | | | | | | | | | |
| Data Entry Error | -1.5 | -5 | -9 | | | | | -1 | -2 | | | |
| Distributed Denial of Service | | 2 | 4 | | | -35,000 | | -1 | -1 | | -1 | -1 |
| Denial of Service | | -2 | -8 | | | 35,000 | | 1 | 1 | | 1 | 1 |
| Electronic Graffiti | | | | | | | | | | | | |
| Fraud/Embezzlement | -1.8 | -4.5 | -49 | 75 | 4,950 | | | | | | | |
| Internal Vandalism | | -2 | -6 | | | | | | | | | |
| IP Spoofing | | | | | | | | | | | | |
| Logic Bomb | -0.03 | | | | | | | | | | | |
| Message Stream Mod | | -0.03 | | | | | | | | | | |
| Password Guessing | -726 | -1,801 | -7,240 | | | | | | | | | |
| Password Nabbing | -0.5 | -3 | -3 | | | | | | | | | |
| Personal Computer Abuse | -182,499 | -364,998 | -729,996 | | | | | | | | | |
| Procedural Violation | -35,999 | -71,998 | -119,996 | 3,000 | 10,000 | | | | | | | |
| Scanning | -2 | -7 | -38 | | -600 | 9,600 | | | | | | |
| Signal Interception | -34 | -38 | -152 | | 600 | -9,600 | | | | | | |
| Theft | -0.5 | -1 | -3 | -24,750 | -495,000 | -1,980,000 | | | | | | |
| Trap Door | | | | | | | | | | | | |
| Trojan Horse | -1 | -2 | -5 | | | | | | | | | |
| Vandalism | | | | | | | | | | | | |
| Virus | | | | | | | | | | | | |
| Web Page Spoofing | -1.5 | -9 | -17 | -4,000 | -6000 | -17,000 | | | | | | |

5.4 Final SAEM Risk Assessment

Table 5 - 7 presents the results from the final SAEM risk assessment. In the final SAEM risk assessment, *Virus* threats ranked the highest, and its relative threat index (100) was significantly greater than those of other threats. The Architect also ranked *Virus* as the most significant threat to the organization.

TABLE 5 - 7 REVISED THREAT PRIORITIES

| Threat | Original Relative Threat Index | Revised Relative Threat Index | Initial SAEM Rank (IS) | Final SAEM Rank (FS) | Initial Architect Rank (IA) | Final Architect Rank (FA) | FS-IS | FA-IA |
|-----------------------------|--------------------------------|-------------------------------|------------------------|----------------------|-----------------------------|---------------------------|------------|------------|
| Virus | 2.0 | 100 | 3 | 1 | 1 | 1 | 2 | 0 |
| Compromise | .1 | 0.243 | 10 | 2 | 3 | 3 | 8 | 4 |
| Password Guessing | .5 | 0.115 | 7 | 3 | 2 | 2 | 4 | 12 |
| Alteration | .9 | 0.102 | 5 | 4 | 5 | 5 | 1 | 4 |
| Denial of Service Attack | .9 | 0.064 | 6 | 5 | 7 | 7 | 1 | 2 |
| Internal Vandalism | 0.006 | 0.051 | 17 | 6 | 6 | 6 | 11 | 16 |
| System Scanning | 0.007 | 0.051 | 15 | 7 | 4 | 4 | 8 | 2 |
| Contamination | 1.9 | 0.038 | 4 | 8 | 10 | 10 | 4 | 2 |
| Distributed DoS | 0.002 | 0.029 | 19 | 9 | 16 | 16 | 10 | 13 |
| Electronic Graffiti | 0.002 | 0.028 | 18 | 10 | 9 | 9 | 8 | 8 |
| WEB Page Spoofing | 0.02 | 0.026 | 11 | 11 | 11 | 11 | 0 | 5 |
| Signal Interception | 0.01 | 0.026 | 12 | 12 | 8 | 8 | 0 | 16 |
| Theft | 0.01 | 0.023 | 13 | 13 | 15 | 15 | 0 | 5 |
| Vandalism | 0.002 | 0.022 | 22 | 14 | 12 | 12 | 8 | 4 |
| Compromising Emanations | 0.002 | 0.022 | 20 | 15 | 13 | 13 | 5 | 2 |
| Browsing | 0.4 | 0.017 | 9 | 16 | 14 | 14 | 7 | 5 |
| Procedural Violation | 23 | 0.013 | 2 | 17 | 18 | 18 | 15 | 3 |
| Data Entry Error | 0.008 | 0.010 | 14 | 18 | 17 | 17 | 4 | 1 |
| Trojan Horse | 0.001 | 0.009 | 23 | 19 | 19 | 19 | 4 | 15 |
| Fraud/Embezzlement | 0.007 | 0.009 | 16 | 20 | 23 | 23 | 4 | 2 |
| Personal Computer Abuse | 100 | 0.008 | 1 | 21 | 21 | 21 | 20 | 1 |
| Password Nabbing | 0.4 | 0.005 | 8 | 22 | 20 | 20 | 14 | 5 |
| IP Spoofing | <0.001 | 0.004 | 24 | 23 | 22 | 22 | 1 | 4 |
| Cryptographic Compromise | 0.002 | 0.004 | 21 | 24 | 25 | 25 | 3 | 2 |
| Message Stream Modification | <0.001 | 0.003 | 25 | 25 | 27 | 27 | 0 | 4 |
| Trap Door | <0.001 | 0.003 | 26 | 26 | 24 | 24 | 0 | 12 |
| Logic Bomb | <0.001 | 0.003 | 27 | 27 | 26 | 26 | 0 | 13 |
| Total | | | | | | | 142 | 162 |

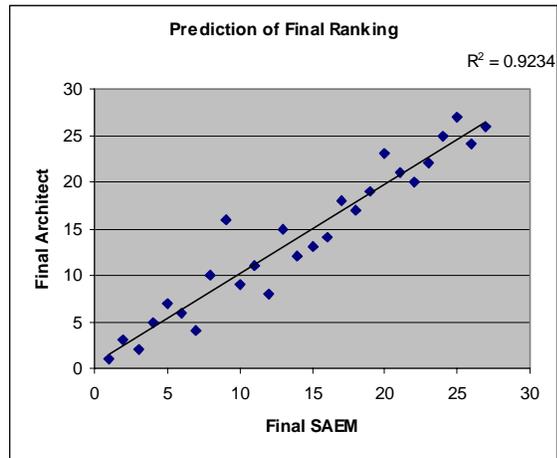
The correlation between the final SAEM risk assessment and the Architect’s revised rankings is .96, indicating a very strong relationship between the final SAEM risk assessment and the Architect’s final rankings. In addition, Table 5 - 7 shows the absolute change in rankings between the initial rankings and the final rankings. The total change for the Architect (162) was greater than the total change for the SAEM risk assessment (142), indicating that the model had a significant influence on the Architect’s final threat ranking.

5.4.1 Regression Analysis

Regression analysis indicates how well the SAEM model can predict the Architect’s final threat ranking. The R²-value, also known as the *coefficient of determination*, measures the percentage of variation in the values of the dependent variable (final Architect ranking) that can be explained by the independent variable (the final SAEM ranking). Figure 5-1 shows the regression line

Figure 5 - 1 Regression Line for Final Results

and the R²-value (.923) for the final SAEM risk assessment model. This high R²-value indicates that over 92% of the variation in Architect’s final ranking can be explained by the final SAEM risk assessment model. The remaining 8% of the variation is due to random or unknown variability.



Although the final SAEM risk assessment appears to be able to closely predict the final Architect’s ranking, it is possible that the Architect’s initial ranking is a better predictor of the final ranking, or that the Architect’s initial rank is a factor in predicting the Architect’s final rank. Table 5 - 8 shows the results of regression analysis using both the Final risk assessment and the Initial Architect’s rankings as predictors of the final rankings. The *Coefficients* column shows the prediction equation¹⁵ and the *t Stat* column shows the ratio between the coefficient and the standard error. The *P-value* is the probability of a t-value this large or larger. Therefore, a P-value less than .05 indicates that the coefficient is significant. Table 5 - 8 shows that only the final SAEM coefficient is significant in predicting the Architect’s final rank.

Table 5 - 8 Regression Analyses for Risk Assessment Results

| | <i>Coefficients</i> | <i>Standard Error</i> | <i>t Stat</i> | <i>P-value</i> |
|------------|---------------------|-----------------------|---------------|----------------|
| Intercept | 1.010 | 0.968 | 1.044 | 0.307 |
| Final SAEM | 1.005 | 0.067 | 14.981 | 0.000 |

¹⁵The prediction equation from the coefficients in the second column is:

$$\text{Threat}_{\text{rank}} = 1.010 + 1.005(\text{SAEM final Rank}) - .078(\text{Architect initial Rank})$$

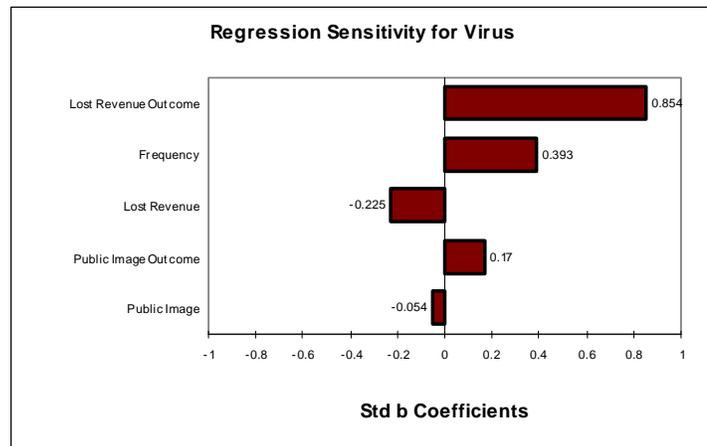
| | | | | |
|--------------|--------|-------|--------|-------|
| Initial Arch | -0.078 | 0.067 | -1.156 | 0.259 |
|--------------|--------|-------|--------|-------|

The final sensitivity analysis shows that the relative threat index for *Virus* attacks is most sensitive to the Lost Revenue consequences of the attack and the frequency of attack. Figure 5-2 shows that the *b* coefficient for Lost Revenue is .85 and the *b* coefficient¹⁶ for Virus Frequency is .39. The Architect estimated that a successful virus attack most likely resulted in 40 hours of lost productivity (\$4,000 in lost revenue) and occurred 72,000 times each year¹⁷, i.e., the most frequently occurring attack. In addition, a *Virus* attack has the highest potential Lost Revenue value (\$250,000). It is not surprising that SAEM and the Architect ranked *Virus* threats as the most significant threat to the organization. Since the *Virus* relative threat index will be most affected by a change in the Lost Revenue Outcome, if the Architect believed that the relative prioritization of *Viruses* to other threats was not correct, then she could re-evaluate the Lost Revenue Outcome first to produce the greatest change in relative threat index.

Figure 5- 2 Virus Tornado Graph

5.4.2 Risk Assessment Insight

One of the key insights that the risk assessment phase highlighted was the security staff's conflict between the organization's relaxed and trusting culture and the risk from internal threats. At first, the Architect estimated a very high number of minor security policy violations, such as procedural violations and personal computer abuse. Although the Architect revised these estimates to include only those incidents that were intentional or malicious, the risk assessment suggested a significant loss of productivity because of these minor violations. In addition, unintentional or non-malicious violations put the organization at greater risk for viruses and other external attacks.



5.5 Benefit Analysis

After completing the risk assessment, the Architect and a technical advisor completed the benefit analysis phase of SAEM. These benefit-analysis interviews took approximately 5 hours to complete. The Architect thought the process of estimating the effectiveness of security technologies was valuable, but tedious.

¹⁶ Recall from chapter 3 that the threat beta coefficients estimate the linear relationship of the attribute values to the threat index.

¹⁷ Recall that this is the estimated consequence if anti-virus software were not installed.

5.5.1 Initial Iteration

Initially, the Architect and technical advisor selected ten security technologies that they believed would be the most effective in reducing the organization’s overall security risks. Table 5 - 9 shows their ordered selection of security technologies. Anti-virus Software is at the top of the list because the organization had rated viruses as the most significant threat to the information system.

Next, the Architect and technical assistant identified risk-mitigating security technologies and estimated their effectiveness against each of the threats. The BENEFIT ANALYSIS DATA section at the back of this chapter shows their estimates for each threat. On average, the Architect and technical advisor selected seven security technologies per threat. Recall from Chapter 4 that the benefit analysis step uses the risk assessment to calculate a threat index change for each technology. The overall threat index change is determined based on the cumulative effectiveness of the security technology against all threats. Therefore, SAEM rates security technologies that are effective against threats with high threat indexes more highly than technologies that are effective against threats with low threat indexes.

Table 5 - 10 shows the results of computing the effectiveness for each security technology in Table 5 – 9. SAEM analysis showed that Hardening the Operating System was the most effective risk-mitigation measure because it reduced the overall threat index by 99.4%, more than any other technology. Anti-virus Software and Forensic Software were identified as the next most effective security technologies in reducing the organization’s risks. Anti-virus Software reduced the threat index by 99.2% and Forensic Software reduced the threat index by 49.7%. Notice that the remaining security technologies in Table 5 - 10 changed the total threat index less than 1%. Although these technologies were highly ranked relative to other security technologies, their overall effectiveness is estimated to be considerably less than the top three security technologies. Their effectiveness ratings are low because the Architect identified them as being effective against threats with low threat indexes.

Table 5 - 9 Architect’s Most Effective Security Technologies

| Initial Rank | Security Technology |
|--------------|-------------------------|
| 1 | Anti-virus Software |
| 2 | Hardened OS |
| 3 | Line Encryption |
| 4 | Secure ID/Password |
| 5 | Automatic Logout |
| 6 | Host-based IDS |
| 7 | Virtual Private Network |
| 8 | Proxy Firewall |
| 9 | Packet Filter Firewall |
| 10 | Net-based IDS |

Table 5 - 10 SAEM Most Effective Security Technologies

| Rank | Security Technology | Threat Index % Change |
|------|------------------------------|-----------------------|
| 1 | Hardened OS | 99.4% |
| 2 | Antivirus Products | 99.2% |
| 3 | Forensic Software | 49.7% |
| 4 | DB Encrypted Data Storage | .40% |
| 5 | Auditing Tools | .31% |
| 6 | Secure E-mail | .29% |
| 7 | Authorization Policy Servers | .29% |
| 8 | Authentication Tokens | .24% |
| 9 | Smart Card Products | .24% |
| 10 | Biometrics | .21% |

5.5.2 Analysis of Results

Since Hardening the OS is effective against ten threats, SAEM ranked it the most effective overall. Although Anti-virus software is highly effective against viruses, it was ranked second because it was not identified as effective against other threats. Interestingly, SAEM computed Forensic Analysis software as the third most effective risk-mitigation measure for the organization. The Architect and technical advisor identified Forensic Analysis software as effective in mitigating 12 different threats, especially virus threats. The Architect agreed that Forensic software should be included in the top ten list, and that the Virtual Private Network should be removed since line encryption was already included and performed a similar function.

Although SAEM determined that several of the authentication mechanisms, such as biometrics, smart cards, one-time passwords, and user ID/Password ranked high in effectiveness relative to other security technologies, these mechanisms all mitigate the same threats, and the participants considered them almost equally effective. For example, the

Architect considered biometrics only five percent more effective in mitigating threats than other authentication mechanisms. The organization identified Secure ID/Password authentication mechanisms as one of their top most effective security technologies, consistent with SAEM results.

From the initial list presented by the analyst, the Architect did not identify any threats mitigated by several of the technologies. These technologies are shown in Table 5 - 11. Although the Architect agreed that there was little benefit in the technologies that SAEM identified as having little or no benefit, other security managers felt several of the technologies, such as SPAM filters and URL blockers, reduced the risk of viruses. I identified in my final report to the Architect that other security managers had identified SPAM filters and URL Blockers as risk-mitigating technologies.

Table 5 - 11 Not Identified as Effective Technologies

| Security Technologies |
|-----------------------|
| SPAM Filters |
| URL Blockers |
| PKI/Cert Products |
| E-Mail Filters |
| Sniffer Detection |
| Mobile Code Scanners |
| Anti-SPAM Filters |
| Modem Access Control |
| Software Lockout |
| Key Stroke Replicator |

5.6 Coverage Evaluation

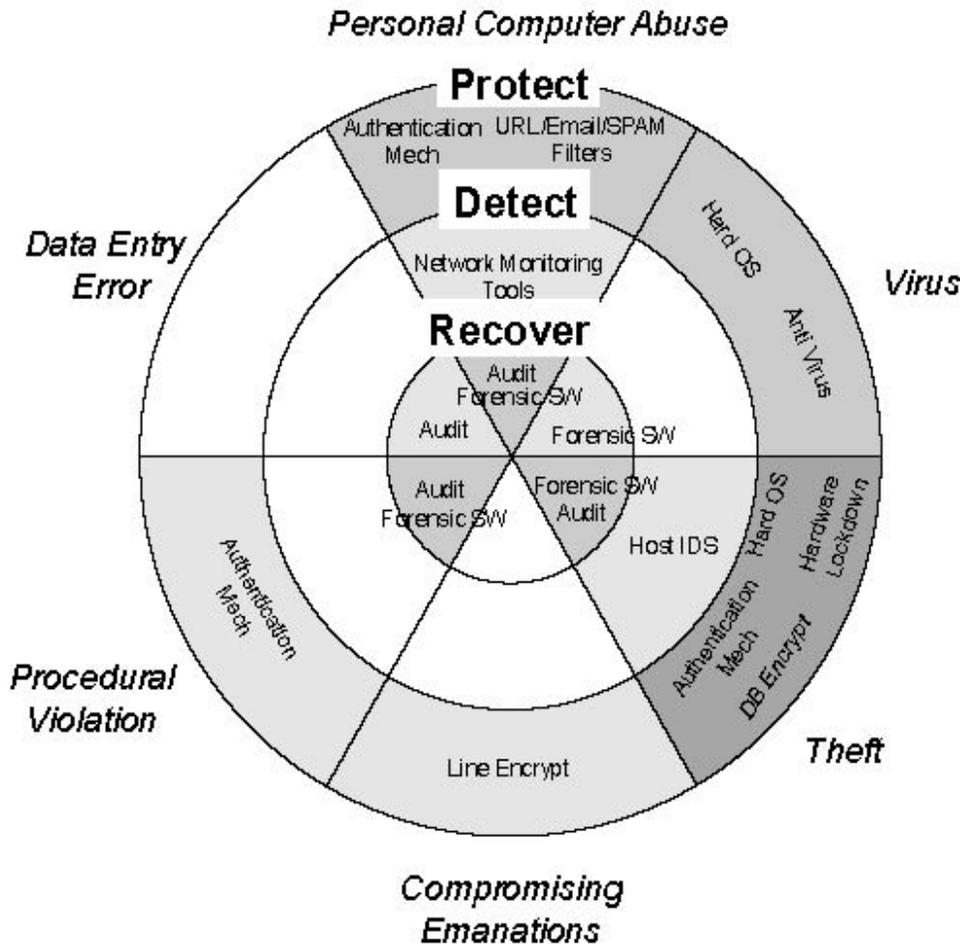
The Architect identified technologies currently in the organization’s security architecture in order to complete coverage evaluation. Table 5 - 12 shows a list of these technologies. Figure 5-3 is the coverage model, which shows how the organization’s security technologies mitigate the risk from the top six threats, as computed by SAEM. In the coverage model, the shaded areas indicate security technologies are present and the darker the shading the more security technologies. Figure 5-3 shows that the organization is weak in detection mechanisms and that mitigation of data entry errors occurs only through the Auditing Tools—a recovery mechanism. The Architect did not identify many security technologies for data entry errors, because application-level edit checks detect and prevent these types of errors. Application-level edit checks were not considered a security technology for the purpose of this study so would not have been identified in the coverage model.

Although the organization appears weak in detection mechanisms for its top threats, the organization already has network-monitoring tools, host-based IDS and network IDS -- the three most important detection mechanisms. The Architect did not identify these mechanisms as risk mitigating to the organization for the threats illustrated in the coverage model; however, other case-study security managers identified these security technologies as risk mitigating in their organizations. The Architect found Figure 5-3 interesting and felt that it gave a “clear view” of what was missing in the security architecture, but she did not make changes to the benefit-analysis inputs on the basis of this figure

TABLE 5 - 12 SECURITY ARCHITECTURE COMPONENTS

| Security Technology |
|-----------------------------------|
| Anti-virus Software |
| Auditing Tools |
| Automatic Logout |
| DB Security Access Controls |
| DB Encrypted Data Storage |
| Email Filters |
| Forensic Software |
| Hardened OS |
| Hardware Lockdown |
| Host-based IDS |
| Line Encryption |
| Load Balancers |
| Modem Access Control |
| Network Monitoring |
| One Time Passwords |
| Packet Filter Firewalls |
| Penetration Testing Tools |
| PKI/Cert Products |
| Proxy Firewalls |
| Secure Email |
| Secure User ID/Password |
| SPAM Filters |
| URL Blockers |
| Virtual Private Network |
| Vulnerability Assessment Scanners |

Figure 5 - 3 Coverage Model for Top Six Threats
 (Darker shading indicates greater security technology coverage)



5.7 Security Technology Tradeoff Analysis

In the final phase of SAEM, the Architect selected and ranked four technologies for comparison and identified several objectives that she used to assess whether to employ a security technology for inclusion in the security architecture. This section describes the results of using multi-attribute analysis techniques to help the Architect decide which security technologies best meet her objectives.

5.7.1 Initial Iteration

First, the Architect selected and ranked four security technologies: 1) Host Intrusion Detection, 2) Network-based Intrusion Detection, 3) Smart Cards, and 4) Biometrics (in order). Next, she identified six objectives that she uses to select security technologies:

1. Ease of Maintenance
2. Purchase Cost
3. Global Deployment¹⁸
4. Effectiveness
5. Business Alignment¹⁹
6. Cultural Impact.²⁰

Next, the Architect determined the weights of each of the objectives. She allocated 100 points across each of the objectives to establish the relative weight of each objective. After establishing the weights, the Architect provided her assessment of each technology in relation to each objective. For example, she ranked the most effective technology, Smart Card, with 100 points and the others relative to it. Table 5 - 13 shows the security technologies and the weights of the objectives. Initially, the analyst would have used the effectiveness ratings from the benefit analysis phase, but the effectiveness ratings for these technologies were all less than 1% and not significantly distinguished from each other.

5.7.2 Refinement

Table 5 - 14 shows the multi-attribute analysis results, which depict a higher ranking for Smart Cards than that offered by the Architect. In fact, the correlation between the multi-attribute analysis and the Architect's initial ranking is 0.2. When the Architect reviewed the results, she eliminated the Cultural Impact objective and redistributed the weights according to Table 5 - 15. The effect of eliminating the objective was that Smart Cards and Biometrics were ranked even more highly than before. Further discussions with the Architect disclosed that she perceived User ID/Passwords as a fairly effective risk- mitigation mechanism, and the added value from Biometrics and Smart Cards would have been minimal compared to the benefits she expected to achieve from integrating more intrusion detection mechanisms into the security architecture.

¹⁸ Global deployment refers to how well the security technology will fit within foreign information system operations. Security components that will operate in non-USA countries must have USA export approvals, and have native-language interfaces as a minimum.

¹⁹ Business alignment refers to the degree that a technology is consistent with the strategic business plan.

²⁰ Cultural Impact refers to the degree that a security technology will have an impact on the organization's operational environment.

Table 5 - 13 Initial Security Technology Tradeoff Assessments

| | | Tradeoff Attribute | | | | | |
|----------------------------|-------------------|---------------------|---------------|-------------------|---------------|--------------------|-----------------|
| | | Ease of Maintenance | Purchase Cost | Global Deployment | Effectiveness | Business Alignment | Cultural Impact |
| Objective Weight | | 25 | 20 | 15 | 20 | 5 | 15 |
| Security Technology | Network Based IDS | 25 | 60 | 100 | 80 | 100 | 80 |
| | Host Based IDS | 30 | 25 | 80 | 71 | 80 | 100 |
| | Biometrics | 100 | 75 | 15 | 60 | 10 | 10 |
| | Smart Card | 75 | 100 | 55 | 100 | 20 | 30 |

Table 5 - 14 Initial Security Technology Tradeoff Analysis

| | | Tradeoff Attribute | | | | | | $\Sigma w_i v_i(x_i)$ |
|----------------------------|-------------------|---------------------|---------------|-------------------|---------------|--------------------|-----------------|-----------------------|
| | | Ease of Maintenance | Purchase Cost | Global Deployment | Effectiveness | Business Alignment | Cultural Impact | |
| Objective Weight | | 0.25 | 0.20 | 0.15 | 0.20 | 0.05 | 0.15 | |
| Security Technology | Network Based IDS | 0.11 | 0.23 | 0.40 | 0.26 | 0.48 | 0.36 | 0.26 |
| | Host Based IDS | 0.13 | 0.10 | 0.32 | 0.23 | 0.38 | 0.45 | 0.23 |
| | Biometrics | 0.43 | 0.29 | 0.06 | 0.19 | 0.05 | 0.05 | 0.22 |
| | Smart Card | 0.33 | 0.38 | 0.22 | 0.32 | 0.10 | 0.14 | 0.28 |

Table 5 - 15 Revised Multi-attribute Analysis

| | | Tradeoff Attribute | | | | | $\Sigma w_i v_i(x_i)$ |
|----------------------------|-------------------|---------------------|---------------|-------------------|---------------|--------------------|-----------------------|
| | | Ease of Maintenance | Purchase Cost | Global Deployment | Effectiveness | Business Alignment | |
| Objective Weight | | 0.30 | 0.25 | 0.15 | 0.25 | 0.05 | |
| Security Technology | Network Based IDS | 0.11 | 0.23 | 0.40 | 0.26 | 0.48 | 0.21 |
| | Host Based IDS | 0.13 | 0.10 | 0.32 | 0.23 | 0.38 | 0.16 |
| | Biometrics | 0.43 | 0.29 | 0.06 | 0.19 | 0.05 | 0.22 |
| | Smart Card | 0.33 | 0.38 | 0.22 | 0.32 | 0.10 | 0.26 |

The final correlation of the revised tradeoff analysis was -0.8. If the Security Architect had not had any authentication mechanism in place, then she certainly would have selected authentication mechanisms over intrusion detection technologies. If her original rankings are reordered to show

authentication mechanisms over intrusion technologies, then the correlation is .8 between multi-attribute analysis and a re-ordered ranking based on authentication mechanisms as a priority over intrusion detection mechanisms. Table 5 - 16 compares the original rank with SAEM's rank and a rank based on a hypothetical preference for authentication mechanisms over intrusion detection mechanisms.

TABLE 5 - 16 RANK COMPARISONS

| | | Architect's Original Rank | SAEM Rank | Rank Assuming Authentication as a Priority |
|---------------------|-------------------|---------------------------|-----------|--|
| Security Technology | Host Based IDS | 1 | 4 | 3 |
| | Network Based IDS | 2 | 3 | 4 |
| | Smart Card | 3 | 1 | 1 |
| | Biometrics | 4 | 2 | 2 |

5.7.3 Analysis of Results

Although the security tradeoff analysis did not result in a clear security technology recommendation, the commercial case study tradeoff analysis did highlight a problem in constructing a hypothetical comparison of security technologies. The SAEM security tradeoff analysis assesses the Security Manager's preferences regardless of the existing security technologies. During case study interviews, the Architect indicated that she chose intrusion detection mechanisms because she needed more detection mechanisms and felt that switching to Smart Cards or Biometrics would not have resulted in any perceived added value. The Security Tradeoff analysis can be modified to compare security technologies *given the current security architecture*, which would probably make the comparison more meaningful to the Architect.

5.8 Summary

The satisfaction survey indicated that the Architect was very pleased with the results of the risk assessment, benefit analysis, and coverage evaluation. The Architect seemed to feel that most of the benefit of SAEM was in the risk assessment, and the analysis of the risk-assessment results confirmed that the SAEM risk-assessment process significantly influenced the Architect's final prioritization of threats. Although the benefit analysis process was tedious, the Architect felt the analysis validated her perceptions about the effectiveness of the corporation's security. More importantly, the analysis helped point out security technologies, such as URL Blockers and SPAM filters, which the security organization had overlooked as potentially important to mitigating their most important threat: *Viruses*. The Architect appeared to be most pleased with the coverage evaluation. She felt that this evaluation gave her the most insight of SAEM's four phases. Finally, although the security tradeoff analysis did not show clear results, it did show the need to revise the process to assess preferences based on the organization's security architecture.

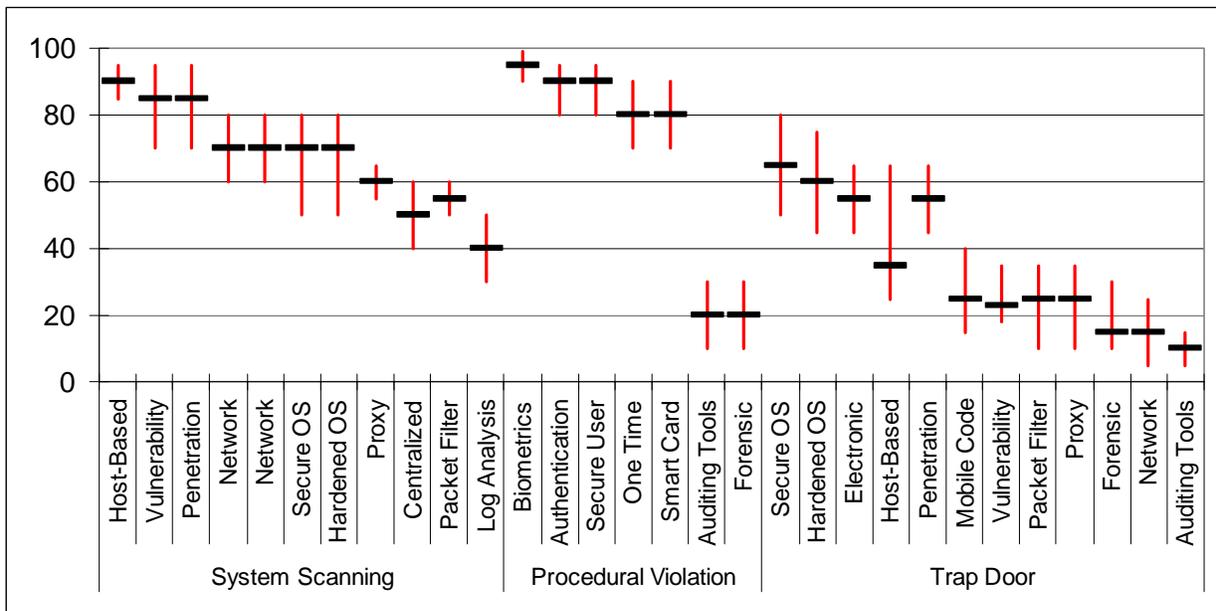
Six months after I delivered the final report to the Architect, she reported that the results of the analysis had been used to justify purchase of intrusion detection mechanisms that would address the

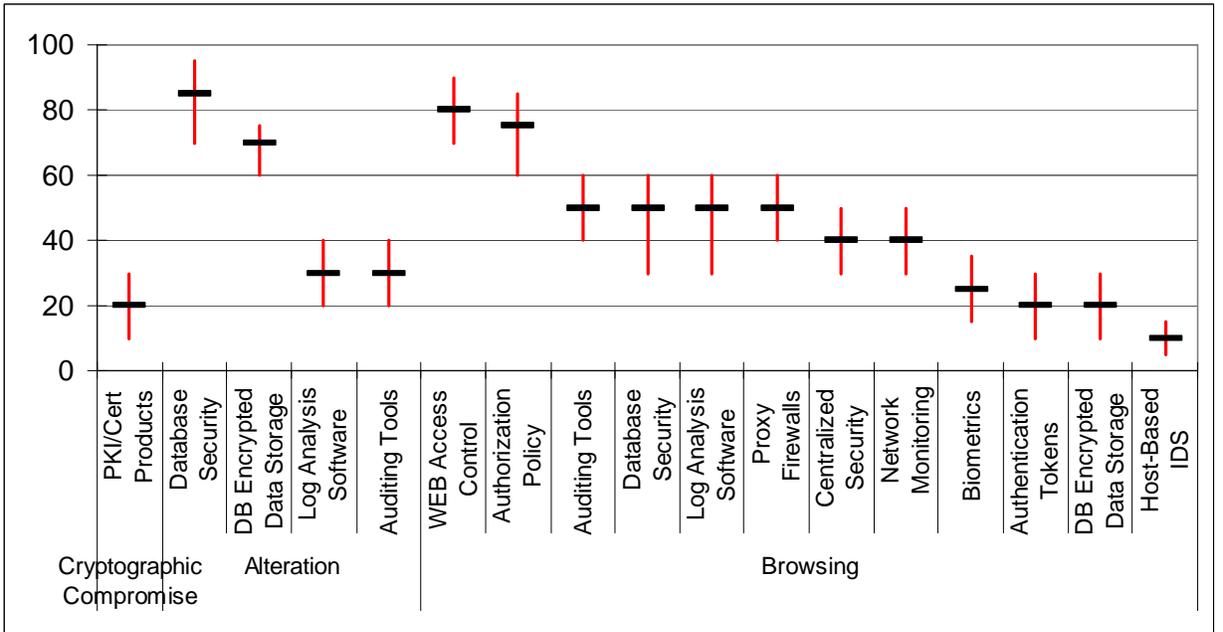
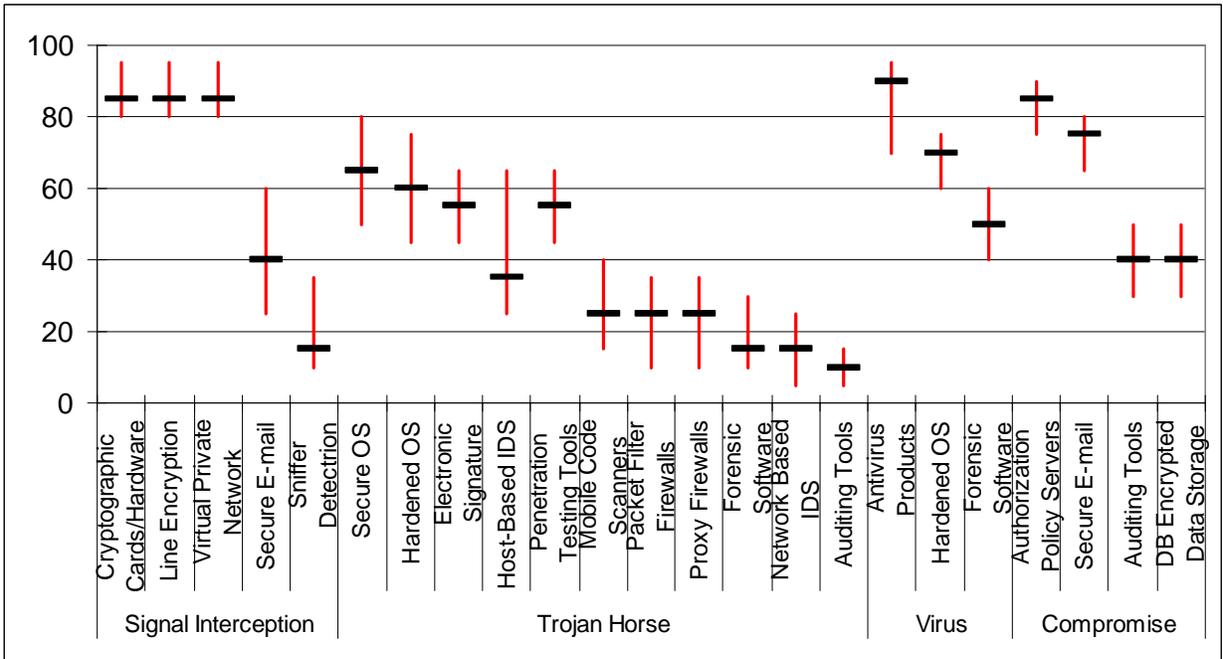
weaknesses identified during the coverage analysis. In addition, the Architect intended to use the results of the analysis in preparation and defense of the organization's next fiscal year budget.

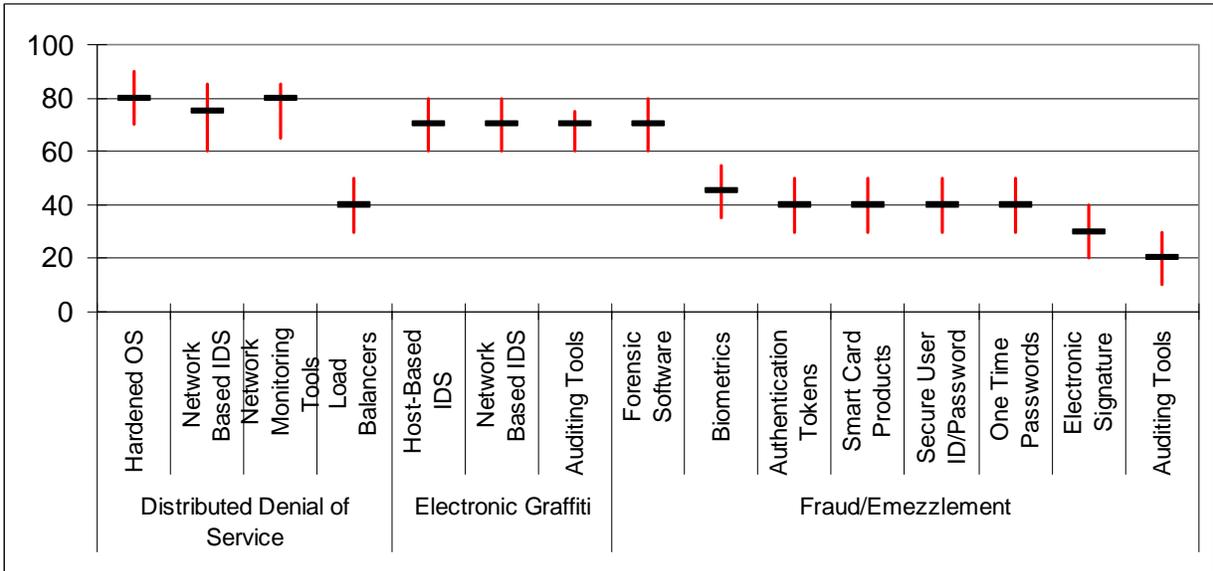
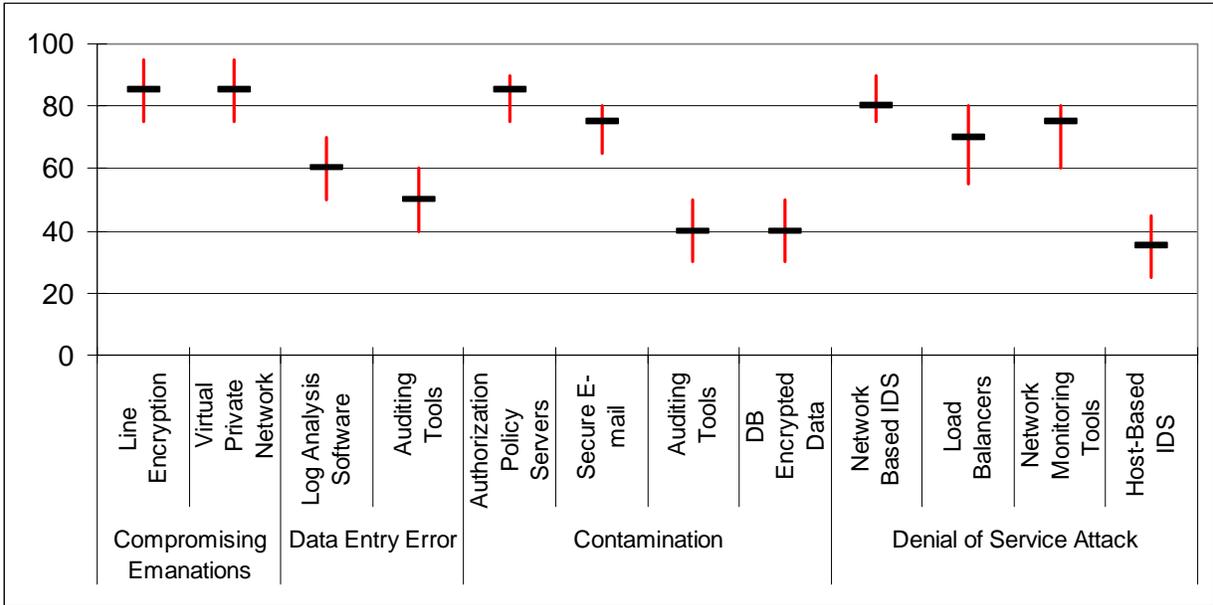
BENEFIT ANALYSIS DATA

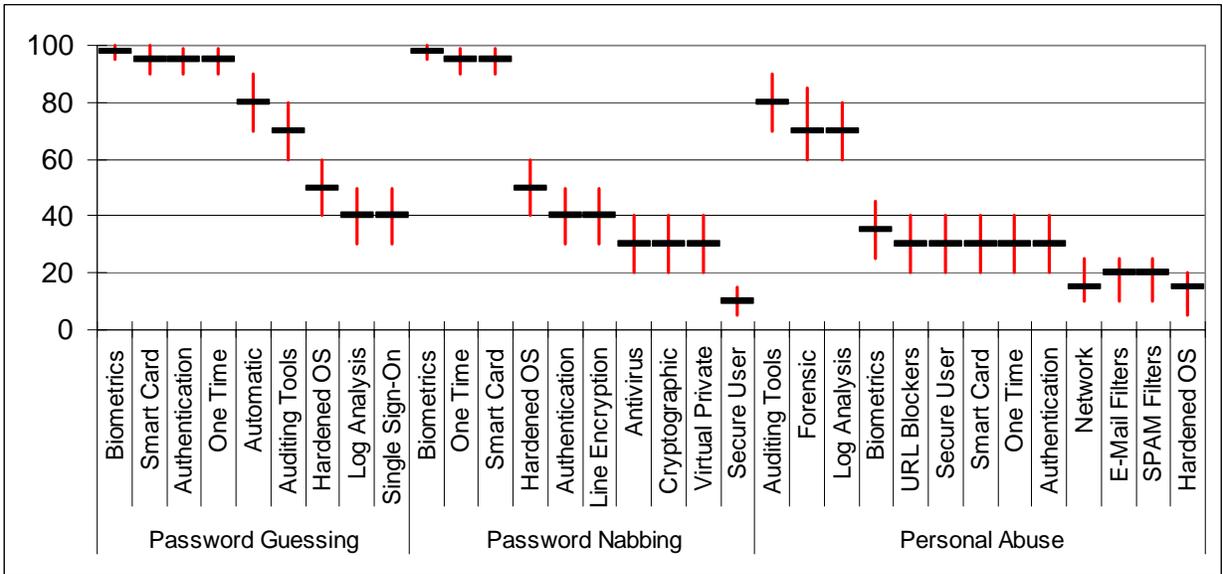
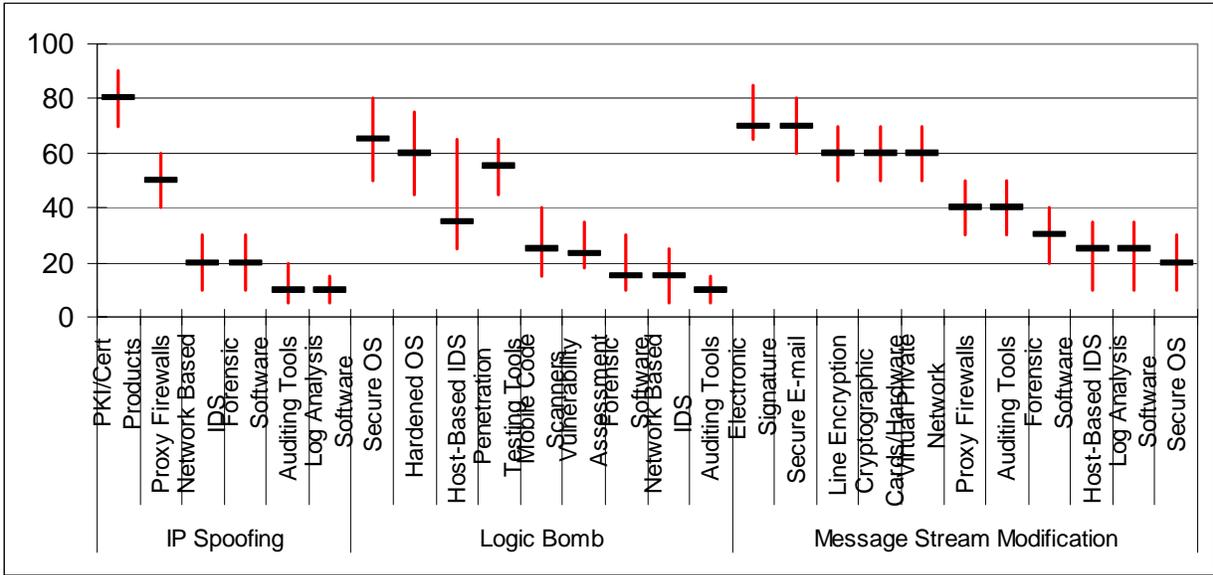
(Commercial Case Study)

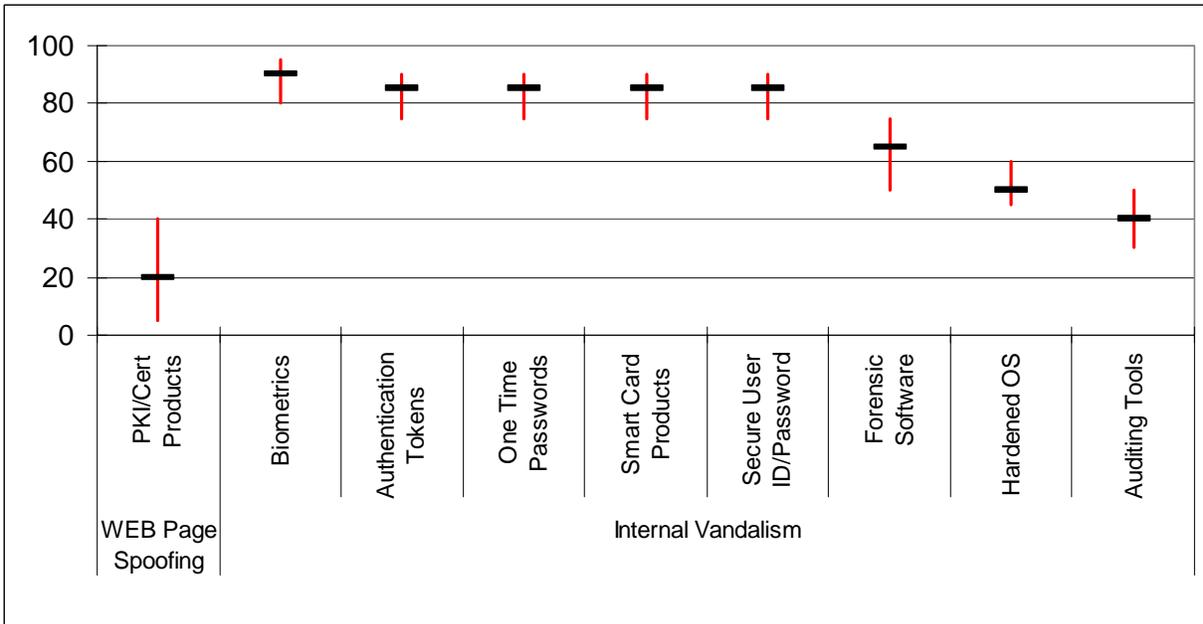
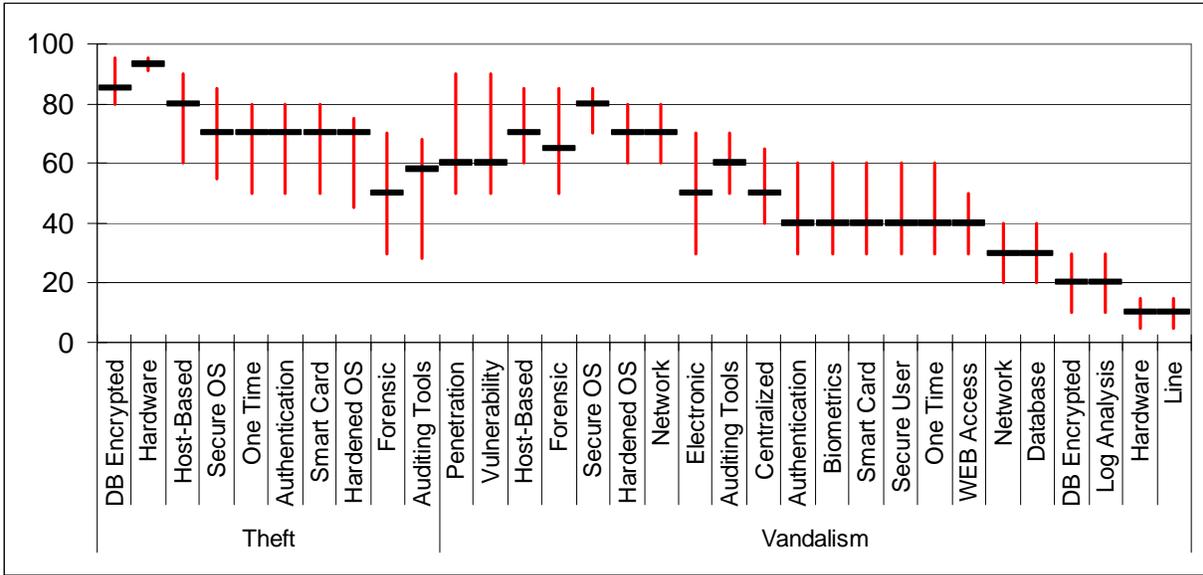
The following charts show the effectiveness ratings for each threat and the security technologies that the Architect identified as risk-mitigating. The Y-axis is the percentage effectiveness, with the black tick mark at the most-likely effectiveness for the technology. The red vertical lines show the range of estimated effectiveness. For example, The Architect estimated that a Host-based IDS is most likely to be 90% effective against System Scanning attacks, but could range from 85% effective to 95% effective. The technologies are sorted according to the most-likely effectiveness within each threat.











SATISFACTION SURVEY

(Commercial Case Study)

The Global Security Architect completed this survey at the end of the SAEM process. The comments at the end of each section are the Architect's unedited comments.

I. Risk Assessment: During the Risk Assessment, participants identified the organization's threats and estimated the frequency and outcome of attacks. The Risk Assessment resulted in a prioritization of the organization's threats.

| How difficult was it ... | <i>Not at all difficult</i> | | <i>Somewhat</i> | | | <i>Very difficult</i> | |
|---|-----------------------------|---|-----------------|---|----------|-----------------------|---|
| | 0 | 1 | 2 | 3 | 4 | 5 | 6 |
| ... to identify and initially rank the threats? | 0 | 1 | 2 | 3 | 4 | X | 6 |
| ... to estimate the frequency of attacks? | 0 | 1 | 2 | 3 | X | 5 | 6 |
| ... to estimate the outcomes? | 0 | 1 | 2 | 3 | 4 | X | 6 |
| | | | | | | | |
| | <i>None at all</i> | | <i>Somewhat</i> | | | <i>A great deal</i> | |
| | 0 | 1 | 2 | 3 | 4 | 5 | 6 |
| How much insight did you gain about the organization's threats from the Risk Assessment? | 0 | 1 | 2 | 3 | X | 5 | 6 |
| How much insight did you gain about the organization's outcomes from the Risk Assessment? | 0 | 1 | 2 | 3 | 4 | X | 6 |
| How much did the risk assessment change your perception of the organization's risks? | 0 | 1 | 2 | 3 | X | 5 | 6 |
| How much easier would it be to explain the organization's risk priorities using the SAEM Risk Assessment than previous assessments? | 0 | 1 | 2 | 3 | 4 | X | 6 |
| | | | | | | | |

| | | | | | | | |
|---|----------------------------|----|----------------|---|----------|-------------------------|---|
| How strongly would you approve or disapprove of submitting the risk assessment ranking to your CIO for use in making decisions about risk management? | <i>Strongly disapprove</i> | | <i>Neutral</i> | | | <i>Strongly approve</i> | |
| | -3 | -2 | -1 | 0 | X | 2 | 3 |
| | | | | | | | |
| How satisfied are you with the threat rankings? | <i>Very dissatisfied</i> | | <i>Neutral</i> | | | <i>Very satisfied</i> | |
| | -3 | -2 | -1 | 0 | 1 | X | 3 |

What did you like about the Risk Assessment?

- Better understanding of what are the greatest risks to management,
- Better understanding of how threats are perceived differently depending upon a person’s role in the organization,
- Better understand of the political factors (such as perceived ROI, personal biases, etc.) that impact decisions
- Better understanding of our security tools coverage gaps

What would you like to see improved in the Risk Assessment process or results?

- Coming up with more practical ways to determine the impact of the threat. As you know, we struggled with identifying actual \$ lost or revenue lost.
- Also, I would try to reduce the number of threats offered up to the client. We fell into the trap of creating more than you offered and I think that many of the threats could have been placed in certain classifications.
- I also see this tool being something that would require a well trained facilitator to walk a client through. There is a great deal of discussion and debate that goes on through the process and it’s important to the tool facilitator help to capture those issues, walk the client through their questions, and ensure the data is consistent. Your background and understanding about the meaning of the various questions was critical to getting to the heart of the assessment.
- If I were to do it again, I would consider two different models – one for assessing an entire organization’s risks (like we tried to do) and one for looking at a segment of the business or of the security zone. There may be some value in creating separate assessment methodologies depending upon what the client needs – perhaps a ‘fast track high-level’ assessment for broad organizational reviews and more detailed assessments for specific business zones.

II. Benefit Analysis: In the benefit analysis phase, participants estimated the effectiveness of security technologies against the threats resulting in a prioritization of security technologies.

| | |
|---|---|
| <p>How difficult was it....</p> <p>to identify and initially rank the security technologies?</p> <p>to estimate the effectiveness of the technologies?</p> | <p><i>not at all difficult</i></p> <p><i>very difficult</i></p> <p>0 1 2 3 X 5 6</p> <p>0 1 2 3 4 X 6</p> |
| <p>How much insight did you gain about the value that security technologies provide?</p> <p>How much did the benefit analysis change your perception of the organization's security technologies?</p> <p>How much easier would it be to explain why a particular security technology should be purchased?</p> | <p><i>none at all</i></p> <p><i>very much</i></p> <p>0 1 X 3 4 5 6</p> <p>0 1 2 X 4 5 6</p> <p>0 1 2 3 X 5 6</p> |
| <p>How strongly would you approve or disapprove of submitting the benefit analysis results to your CIO for use in making decisions about spending financial resources?</p> | <p><i>strongly disapprove</i></p> <p><i>strongly approve</i></p> <p>-3 -2 -1 0 X 2 3</p> |

| | | | | | | | |
|--|--------------------------|-----------|-----------|----------|----------|----------|-----------------------|
| How satisfied are you with the security technology rankings? | <i>very dissatisfied</i> | | | | | | <i>very satisfied</i> |
| | -3 | -2 | -1 | 0 | 1 | X | 3 |

What did you like about the Benefit Analysis?

- I feel that the best part about this section was really validating what many of us already believed. It was also a very good exercise for our staff to discuss our opinions on actual effectiveness of various controls. Depending upon personal experience, technical knowledge, etc, the answers can vary.

What would you like to see improved in the Benefit Analysis process or results?

- This is a very tedious process. Again, if there's a way to reduce the number technologies or to categorize them differently, it might be easier to see patterns and make faster calls.

II. Coverage Analysis: The coverage analysis showed the security technology defense-in-depth coverage (protect, detect, and recover) of the top six threats as determined by SAEM.

| | |
|---|---|
| <p>How much insight did you gain about the overall defense-in-depth coverage that your organization's current security technologies provide?</p> | <p><i>none at all</i> <i>very much</i></p> <p>0 1 2 3 4 X 6</p> |
| <p>How much did the coverage analysis change your perception of the organization's security status?</p> | <p>0 1 2 3 X 5 6</p> |
| <p>How much easier would it be to explain why a particular security technology should be purchased if the coverage analysis showed a gap?</p> | <p>0 1 2 3 4 X 6</p> |
| | |
| <p>How strongly would you approve or disapprove of submitting the coverage analysis results to your CIO for use in making decisions about spending financial resources?</p> | <p><i>strongly disapprove</i> <i>strongly approve</i></p> <p>-3 -2 -1 0 1 X 3</p> |
| | |
| <p>How satisfied are you with the coverage analysis?</p> | <p><i>very dissatisfied</i> <i>very satisfied</i></p> <p>-3 -2 -1 0 1 X 3</p> |

What did you like about the Coverage Analysis?

- As you know, this was the area that really provided a clear view of what is missing in our security architecture. I feel that that the reason this was so well received is that is a very easy to understand, visually appealing and quickly communicates the issue.

What would you like to see improved in the Coverage Analysis process or results?

- I was very satisfied with this section and don't have any solid recommended changes. I feel that this is the area where you can really map out what you have, what you care about, and where the gaps are. I would recommend that this portion of the assessment process be very flexible to the client so that they can perform their own 'what if' scenarios

CHAPTER 6. Hospital Case Study

6.1 Introduction

This chapter describes the participants, activities and results from the hospital case study. The multi-attribute analyst elicited the data for this case study through a series of interviews and questionnaires during a two-month period. The Hospital's Technical Director (Director) and a technical advisor provided the input for each SAEM step, but the Director made the final decisions during each refinement step. The SAEM risk assessment process helped the Director identify organizational risks that he had not previously considered. The benefit analysis phase showed that there were several technologies, which the Director had not previously considered, that could help reduce risk from attacks. SAEM coverage analysis showed the security architecture was weak in detection mechanisms, and the security-tradeoff analysis showed mixed results when compared to the Director's priorities. The Director ranked intrusion detection technologies more highly than authentication security technologies because the organization's security architecture was weak in detection mechanisms. Overall, the Director reported that the analysis prepared him for the Health Insurance Portability and Accountability (HIPAA) security requirements and made him more aware of weaknesses in the security architecture.

6.2 Case Study Description

The second case study is for a small local hospital that is connected to larger medical facilities through virtual private networks. The hospital has a small staff to run its information system and does not have anyone dedicated full-time to security. The primary security responsibility falls on the Director, who handles the day-to-day operation of the information system. The Director must balance security requirements with medical-staff productivity. The Director noted that medical staff, especially doctors, find some security procedures and policies irritating and complain that security interferes with their ability to work efficiently, thus adversely impacting patient care.

The hospital's security budget is very limited, but the hospital is required to follow the new Health Insurance Portability and Accountability Act (HIPAA) guidelines concerning privacy and security. These guidelines place additional pressure on limited security budgets, so the Director tries to find technologies that balance the objectives of effectively implementing security policies, minimizing the impact on medical staff, and improving productivity.

The hospital's culture is relaxed and informal. Policies governing use of corporate computing resources for personal use, such as personal email, games, and web browsing, are not very restrictive. Most recently, the hospital information system personnel focused on reducing the risk from external attacks, but the organization has experienced few security incidents. The organization did not track security incidents, operate an incident response center, or have a pre-existing risk assessment; therefore, the Director made all estimates based on his subjective assessment or consultation with his advisor.

6.2.1 Participants

The Director was the organization’s primary participant in this case study. The Director was responsible for the hospital’s overall security architecture, security requirements, and security policy changes, and preparation of the organization’s security budget. The Director performed those duties in addition to ensuring the day-to-day operation of the hospital’s information system. One individual assisted the Director by providing advice and expertise about threats and countermeasures. This individual participated in the risk assessment and benefit analysis steps of the case study. Generally, the participants were not familiar with many of the threats and security technologies that the analyst presented.

6.3 The Risk Assessment

This section outlines the sequence of events and presents the data collected during the risk assessment step of the case study. The Director and an advisor provided their initial threat ranking, but the analyst collected the threat frequency and outcome data through a questionnaire. After completing the questionnaire, the analyst presented the SAEM results to the Director, who made revisions to his initial risk priorities and his estimates for the Risk Assessment.

Table 6 - 1 Initial Threat Rankings

| Order | Threat |
|-------|-----------------------------|
| 1 | Virus |
| 2 | Compromise |
| 3 | Alteration |
| 4 | Message Stream Modification |
| 5 | Compromising Emanations |
| 6 | System Scanning |
| 7 | Denial of Service |
| 8 | Signal Interception |
| 9 | Vandalism |
| 10 | Fraud/Embezzlement |
| 11 | Theft |
| 12 | Cryptographic Compromise |
| 13 | WEB Page Spoofing |
| 14 | Electronic Graffiti |
| 15 | IP Spoofing |

6.3.1 Initial Iteration

Table 6 - 1 shows the Director’s initial ranking of threats. The participants eliminated several threats from the analyst’s initial threat list. In most cases, they eliminated threats that were not applicable to their information system environment or were redundant. For example, the Director felt that *Logic Bombs* were essentially the same as *Viruses*. The organization is mostly concerned about protecting patient privacy and maintaining the integrity of patient data, but it has experienced few security compromises each year.

6.3.1.1 Outcome Attributes

The Director identified three attack consequences (or attributes) that were important to the organization: diminished Quality of Patient Care, negative Physician’s Perception, and deteriorated Community Relationships. In addition to minimizing violations of patient privacy, the hospital wanted to avoid attacks that diminished the quality of patient care, which included ensuring the integrity and availability of patient information. In addition, the hospital was concerned with attracting physicians, because physicians may not want to work at the hospital if security compromises are frequent or have serious consequences. The Director also noted that if

physicians perceived security procedures and technologies as too encumbering, then they may not want to work at the hospital. Finally, the hospital was also concerned about its image in the community, since damaged community relationships would keep patients away. Overall, the Director felt that he must balance the risks from a security compromise with the medical staff's perception that security prevents them from doing their job.

Table 6 - 2 represents the results of the Director's ranking of the outcome attributes using the swing-weight method. In this case study, the Director wanted to allocate 100 points among the outcome attributes, rather than give the most important attribute and rank the others relative to 100 points. He found the allocation technique cognitively appealing and the results are similar.

Table 6 - 2 Outcome Attribute Weights²¹

| Attribute | Rank | Weight |
|-------------------------|------|--------|
| Quality of Patient Care | 1 | 0.80 |
| Physician Perceptions | 3 | 0.05 |
| Community Relationships | 2 | 0.15 |

6.3.1.2 Outcome and Frequency Estimates

The participants used the 7-point scale described in Chapter 4 to assess the impact of a security compromise on the Quality of Patient Care, Physician's Perceptions, and Community Relationships. Table 6 - 3 shows all of the Director's estimates for attack outcome values and frequencies. Notice that although the Director expects *Compromising Emanations* to occur more frequently than any other threat, these would most likely have little impact on the Quality of Patient Care.

In contrast to the commercial case study described in Chapter 5, the hospital participants did not rely on previous experiences to guide them in their outcome estimates since they had not experienced many security compromises. In addition, the analyst used a questionnaire to elicit specific threat frequency and outcome values rather than elicit these values through interviews. The Director took approximately two weeks to complete the questionnaire.

²¹ The table is organized according to expected frequency with the most frequently occurring attacks listed first.

Table 6 - 3 Initial Estimated Outcome and Frequency Values
(Sorted by Frequency)

| Threats | Frequency/year | | | Quality of Patient Care | | | Community Relationships | | | Physician Perceptions | | |
|--------------------------|----------------|-------------|------|-------------------------|-------------|------|-------------------------|-------------|------|-----------------------|-------------|------|
| | Low | Most Likely | High | Low | Most Likely | High | Low | Most Likely | High | Low | Most Likely | High |
| Compromising Emanations | 60 | 84 | 180 | 1 | 1 | 2 | 1 | 3 | 5 | 1 | 2 | 3 |
| Alteration | 24 | 60 | 120 | 1 | 4 | 7 | 2 | 4 | 7 | 3 | 5 | 7 |
| Virus | 36 | 60 | 360 | 2 | 3 | 7 | 1 | 5 | 7 | 1 | 4 | 7 |
| Scanning | 24 | 36 | 144 | 1 | 1 | 2 | 1 | 2 | 3 | 1 | 2 | 3 |
| Compromise | 0 | 24 | 120 | 1 | 1 | 3 | 3 | 5 | 7 | 2 | 4 | 7 |
| Signal Interception | 0 | 12 | 36 | 1 | 1 | 3 | 1 | 2 | 4 | 1 | 2 | 4 |
| Theft | 12 | 12 | 60 | 1 | 2 | 4 | 2 | 4 | 6 | 3 | 5 | 7 |
| IP Spoofing | 12 | 12 | 60 | 1 | 3 | 6 | 2 | 4 | 6 | 2 | 3 | 6 |
| Vandalism | 1 | 5 | 12 | 1 | 4 | 5 | 1 | 3 | 5 | 1 | 3 | 5 |
| Denial of Service | 2 | 5 | 15 | 1 | 2 | 4 | 1 | 3 | 4 | 1 | 2 | 4 |
| Cryptographic Compromise | 1 | 3 | 5 | 1 | 3 | 7 | 1 | 4 | 7 | 1 | 3 | 5 |
| Electronic Graffiti | 0 | 1 | 5 | 1 | 1 | 2 | 1 | 2 | 4 | 1 | 1 | 3 |
| Web Page Spoofing | 0 | 1 | 5 | 1 | 1 | 2 | 1 | 2 | 3 | 1 | 2 | 3 |
| Fraud/Embezzlement | 0 | 0.2 | 3 | 1 | 1 | 2 | 2 | 4 | 5 | 2 | 4 | 6 |
| Message Stream Mod | 0 | 0.1 | 1 | 1 | 5 | 7 | 3 | 5 | 7 | 2 | 4 | 6 |

Table 6 - 4 SAEM and Director Threat Rankings

6.3.1.3 Initial Results

Table 6 - 4 compares the results from SAEM using the Director’s initial estimated values. Table 6-4 also shows the Director’s initial ranking for comparison. In general, SAEM and the Director differed significantly in only a few their rankings. The correlation between the SAEM Rank and the Director’s Initial Rank is .54, which indicates that the two ranks are moderately correlated. The last column of Table 6 - 4 shows that two threats, *IP Spoofing* and *Message Stream Modification*, have the greatest difference between SAEM’s rank and the Director’s rank.

| Threat | Relative Threat Index | SAEM Rank (S) | Initial Rank (I) | S-I |
|-----------------------------|-----------------------|---------------|------------------|-----|
| Virus | 100 | 1 | 1 | 0 |
| Alteration | 98 | 2 | 3 | 1 |
| Compromising Emanations | 60 | 3 | 5 | 2 |
| Compromise | 26 | 4 | 2 | 2 |
| System Scanning | 23 | 5 | 6 | 1 |
| IP Spoofing | 19 | 6 | 15 | 9 |
| Theft | 16 | 7 | 11 | 4 |
| Signal Interception | 9 | 8 | 8 | 0 |
| Vandalism | 7 | 9 | 9 | 0 |
| Denial of Service Attack | 5 | 10 | 7 | 3 |
| Cryptographic Compromise | 5 | 11 | 12 | 1 |
| Electronic Graffiti | 0.7 | 12 | 14 | 2 |
| WEB Page Spoofing | 0.7 | 13 | 13 | 0 |
| Fraud/Embezzlement | 0.2 | 14 | 10 | 4 |
| Message Stream Modification | 0.2 | 15 | 4 | 11 |

6.3.2 Refinement

The Director reviewed the results of SAEM and revised his initial rankings, but did not change his inputs to the risk assessment. Initially, upon reviewing the risk-assessment results, the Director changed his rank of *Message Stream Modification* to 11th and reduced the rank of *IP spoofing* to 10th. In addition, the analyst presented to the Director a comparison of some of the risk assessment inputs based on noticeable inconsistencies between the Director’s estimates and ranks. Table 6 - 5 shows the analyst’s comparison of four threats: *Compromise*, *Alteration*, *IP Spoofing* and *Denial of Service*. This type of comparison was useful to the Director because it quickly highlighted inconsistencies between threat rankings and the input values.

Table 6 - 5 Threat Comparisons

| | | Outcome Attributes | | | |
|--------|-------------------|-------------------------|-------------------------|-------------------------|-----------------------|
| | | Expected Frequency/year | Quality of Patient Care | Community Relationships | Physician Perceptions |
| Threat | Compromise | 24 | None | Moderately Severe | Moderate |
| | Alteration | 60 | Moderate | Moderate | Moderately Severe |
| | IP Spoofing | 12 | Moderately Mild | Moderate | Moderately Mild |
| | Denial of Service | 5 | Mild | Moderately Mild | Mild |

The Director ranked *Compromise* second, but SAEM ranked it 4th. In addition, the Director ranked *Alteration* lower than *Compromise*, but when the Director reviewed the risk assessment inputs for *Alteration* and *Compromise*, the comparison shows that the inputs for *Alteration* are more damaging and are likely to occur more often than *Compromise* incidents. Therefore, the Director changed his rankings of *Alteration* and *Compromise* to 2nd and 3rd respectively.

Table 6 - 6 Revised Threat Priorities

Next, the Director compared the inputs for *IP Spoofing* and *Denial of Service* threats. The Manager had just changed the *IP spoofing* ranking to 10th. However, when he reviewed the inputs between the two threats, he saw that the frequency and consequences of an *IP Spoofing* incident were greater than those of a *Denial of Service* attack. Based on this review, he revised his ranking of the two threats, re-ranking *IP Spoofing* lower than *Denial of Service* attacks. The Director's final ranking is shown in Table 6 - 6.

| Threat | Relative Threat Index | SAEM Rank (S) | Director's Initial Rank (DI) | Director's Final Rank (DR) | DR-I |
|-----------------------------|-----------------------|---------------|------------------------------|----------------------------|------|
| Virus | 100 | 1 | 1 | 1 | 0 |
| Alteration | 98 | 2 | 3 | 2 | 1 |
| Compromising Emanations | 60 | 3 | 5 | 4 | 1 |
| Compromise | 26 | 4 | 2 | 3 | 1 |
| System Scanning | 23 | 5 | 6 | 5 | 1 |
| IP Spoofing | 19 | 6 | 15 | 6 | 9 |
| Theft | 16 | 7 | 11 | 12 | 1 |
| Signal Interception | 9 | 8 | 8 | 7 | 1 |
| Vandalism | 7 | 9 | 9 | 8 | 1 |
| Denial of Service Attack | 5 | 10 | 7 | 10 | 3 |
| Cryptographic Compromise | 5 | 11 | 12 | 13 | 1 |
| Electronic Graffiti | 0.7 | 12 | 14 | 15 | 1 |
| WEB Page Spoofing | 0.7 | 13 | 13 | 14 | 1 |
| Fraud/Embezzlement | 0.2 | 14 | 10 | 9 | 1 |
| Message Stream Modification | 0.2 | 15 | 4 | 11 | 7 |
| Total Change | | | | | 30 |

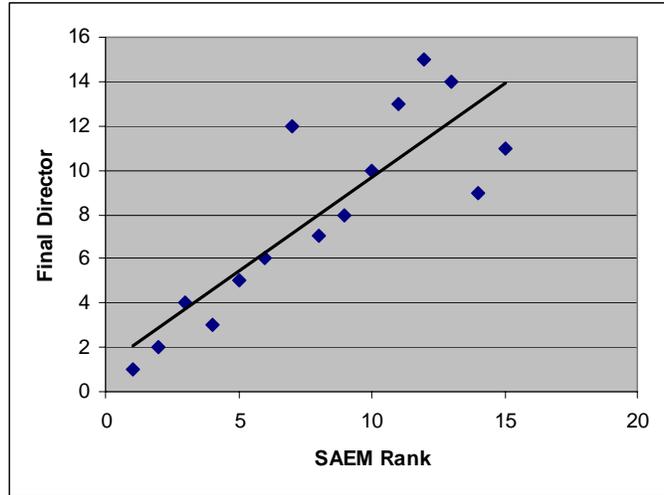
The SAEM risk assessment influenced the Director's final threat ranking, since the Director changed his final rankings and did not change the inputs to SAEM. In addition, the rank correlation between SAEM and the Director's final prioritization was .85, which meant there was a strong correlation between the two rankings.

6.3.3 Regression Analysis

Regression analysis indicates how well the SAEM model can predict the Director's final threat ranking. The R²-value, also known as the coefficient of determination, measures the percentage of variation in the values of the dependent variable (final Director's ranking) that can be explained by the independent variable (the SAEM rank). Figure 6 - 1 shows the regression line and the R²-value (.72) for the SAEM Risk Assessment model. This R²-value indicates that 72% of the variation in the

Director’s final ranking can be explained by the SAEM risk assessment model. The remaining 28% of the variation is due to random or unknown variability.

Figure 6 - 1 Regression Line for Final Results



6.3.3.1 Prediction Regression Analysis

Although the final SAEM Risk Assessment appears to be a significant predictor of the final Director’s ranking, it is possible that the Director’s initial ranking is a better predictor of the final ranking, or that the Director’s initial rank is a significant factor in predicting the final rank. Table 6 - 7 shows the results of regression analysis using both the Final SAEM risk assessment and the Director’s Initial (Initial TD) rankings as predictors of the final rankings. The *Coefficients* column in Table 6 - 7 shows the prediction equation and the *t Stat* column shows the ratio between the coefficients and the standard error. The P-value is the probability of a t-value this large or larger. Therefore, a p-value less than .05 indicates that the coefficient is significant. Table 6 - 7 Regression Analysis for Risk Assessment Results, shows that both the Final SAEM and the Initial coefficients are significant in predicting the Director’s final rank.

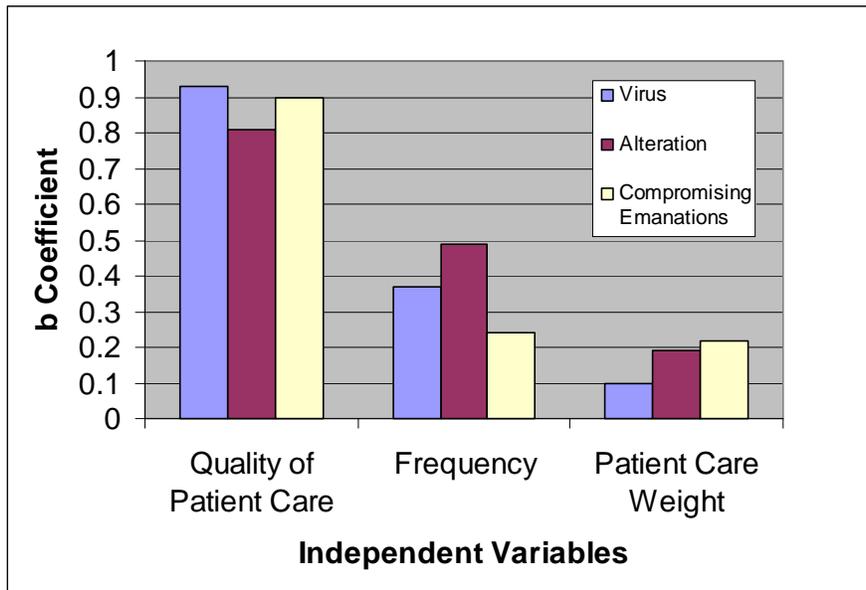
Table 6 - 7 Regression Analysis for Risk Assessment Results

| | <i>Coefficients</i> | <i>Standard Error</i> | <i>t Stat</i> | <i>P-value</i> |
|------------|---------------------|-----------------------|---------------|----------------|
| Intercept | -0.22 | 1.21 | -0.18 | 0.86 |
| Final SAEM | 0.64 | 0.14 | 4.52 | 0.00 |
| Initial TD | 0.39 | 0.14 | 2.72 | 0.02 |

6.3.3.2 Risk Assessment Insight

One of the threat prioritization key insights that the risk assessment step highlighted for the hospital was the impact that an attack has on of the Quality of Patient Care. The top three threats are most sensitive to the consequence of Quality of Patient Care, as shown in Figure 6-2. Recall from Chapter 3 that the regression equation for a threat index helps predict the change in threat index given a percentage change in the standard deviation of an independent variable. The coefficients of the regression equation indicate the percentage change in the threat index, i.e., the dependent variable. For example, in Figure 6 - 2, the Quality of Patient Care consequence is at least .8 for each threat; therefore a one standard deviation change in the Quality of Patient Care consequence for any of the top three threats will result in at least a .8 (or 80%) increase in the threat index.

Figure 6 - 2 Regression Sensitivity of Threat Indexes²²



6.4 Benefit Analysis

After completing the risk assessment, the Director and an advisor completed the benefit analysis step of SAEM. In this step of SAEM, the analyst chose to use an elicitation interview for some preliminary information about security technologies and threats, followed by a questionnaire to elicit the security technology effectiveness estimates from the Director. The questionnaire allows the participants to make the estimates as their time permits and is not quite as tedious as the interview process. The benefit analysis phase elicitation interview took approximately one hour to complete, but the Director took more than two weeks to complete the security technology effectiveness questionnaire. Overall, The Director found it very difficult to estimate the effectiveness of the security technologies, but he claimed that he gained significant insight about the value that some of the security technologies could provide his organization.

6.4.1 Initial Iteration

Initially, the Director identified ten security technologies that he believed would be the most effective in reducing the organization's risks from the threats identified during the risk assessment phase. Table 6 - 8 shows his selection of security technologies. He considered Anti-virus Software the most important security technology, which is consistent with the Director's identification of *Viruses* as the organization's most significant threat.

²² Figure 6-1 only shows independent variables that have at least a b coefficient > .1

Table 6 - 8 Director's Choice for Most Effective Security Technologies for the Organization

After identifying the top ten security technologies, the Director and an assistant identified risk-mitigating security technologies for each threat. On average, they identified four security technologies per threat. During the elicitation interview, the analyst became aware that the participants were not very familiar with many of the security technologies presented, which could partially explain why so few security technologies were identified for each threat.

After the elicitation interview, the analyst gave the Director a questionnaire that required him to estimate the effectiveness of the security technologies against each of the organization's threats. Appendix 6 - A shows his estimates for each threat. On average, he estimated that security technologies are 94% effective against a threat, which was a higher average estimate than those given by other case study security managers²³.

Table 6 - 9 shows the results of computing the effectiveness for each security technology. SAEM analysis showed that Hardware Lockdown was the most effective risk-mitigation measure because it reduced the threat index 40%. The Director identified Hardware Lockdown at least 80% effective for five out of fifteen threats. SAEM Log Analysis Software and Hardening OS were identified as the next most effective security technologies, reducing the organization's risks 38% and 37% respectively. Although the Director had initially selected Anti-virus Software as the most effective risk-reducing technology, SAEM ranked it 8th because it reduced the threat index only 29%. SAEM ranked it lower than the Director's ranking because the Director identified Antivirus Software as effective against only one threat—*Viruses*.

The Director did not identify several security technologies as mitigating any of the organization's threats. For example, the Director identified Secure User ID/Passwords as risk-mitigating, but didn't identify Smart Cards Products or One Time Passwords. These technologies perform the same function as Secure User ID/Passwords, so the Director should have selected them.

| Initial Order | Security Technology |
|---------------|-----------------------------|
| 1 | Anti-virus Software |
| 2 | Authorization Policy Server |
| 3 | Auditing Tools |
| 4 | Centralized Management |
| 5 | Cryptographic Cards |
| 6 | Hardened OS |
| 7 | Hardware Lockdown |
| 8 | Log-on Limits |
| 9 | Network Monitoring Tools |
| 10 | Virtual Private Network |

TABLE 6 - 9 SECURITY TECHNOLOGY THREAT INDEX

| Rank | Security Technology | Threat Index % Change |
|------|-----------------------------------|-----------------------|
| 1 | Hardware Lockdown | 40% |
| 2 | Log Analysis Software | 38% |
| 3 | Hardened OS | 37% |
| 4 | Single Sign-On Apps | 34% |
| 5 | Database Security Access Controls | 34% |
| 6 | Software Lockout | 32% |
| 7 | Host-Based IDS | 32% |
| 8 | Antivirus Products | 29% |
| 9 | E-Mail Filters | 29% |
| 10 | Secure User ID/Password | 28% |
| 11 | Virtual Private Network | 23% |
| 12 | Line Encryption | 23.4% |
| 13 | Secure E-mail | 23.0% |
| 14 | Network Based IDS | 10.4% |

²³ Differences among case studies will be discussed in chapter 8.

6.4.2 Refinement

When the Director compared the SAEM benefit analysis results with his initial selection, he revised his initial list. Interestingly, the Director selected Authorization Policy Servers as one of the top most effective security technologies, but did not specifically identify them as effective against any threat. When questioned, he noted that the Authorization Policy Server was actually an efficient mechanism for allowing users into different systems; therefore, he eliminated it from his original list. The Director also removed Centralized Management technologies from the list for similar reasons, and added Database Security Access Controls and Email Filters. Table 6 - 10 is the Director's revised list based on a review of the SAEM results. He did not change any of the effectiveness estimates shown at the back of this Chapter—Benefit Analysis Data.

Table 6 - 10 Director's Revised List of Organization's Most Effective Security Technologies

6.4.3 Analysis of Results

When compared to the Director's revised list of the organization's most effective security technologies, SAEM identified seven out of ten from the Director's revised list. In addition, SAEM identified Auditing Software, but the Director identified Log Analysis Software in his final analysis; however, these technologies perform overlapping functions. The Director also identified Virtual Private Networks as the 10th most effective technology for the organization and SAEM benefit analysis ranked it 11th.

| Revised Order | Security Technology |
|---------------|-----------------------------------|
| 1 | Anti-virus Software |
| 2 | Database Security Access Controls |
| 3 | Auditing Tools |
| 4 | Host-base IDS |
| 5 | Email Filters |
| 6 | Hardened OS |
| 7 | Hardware Lockdown |
| 8 | Log-on Limits |
| 9 | Single Sign-on |
| 10 | Virtual Private Network |

Since the Director changed his initial assessment based on a review of the SAEM benefit analysis, and there is significant overlap between the Director's assessment and SAEM analysis, the benefit analysis process influenced the Director's prioritizations of security technologies. The Satisfaction Survey results confirm that the Director thought the process was insightful and useful.

Although the Director carefully reviewed each security technology when identifying risk-mitigation technologies for the organization's threats, he failed to select all of the authentication mechanisms when selecting at least one authentication mechanism. As previously mentioned, Smart Cards were not chosen as a risk-mitigation technology for any threat, even though Secure User/ID was selected for *Alteration*. During a follow-up interview, the Director indicated that his organization was considering buying Smart Cards because they would be more efficient for medical staff to use, rather than logging in with a User/ID Password. Overall, the Director selected authentication mechanisms as risk-mitigating technologies in only a few instances, whereas most security managers would consider authentication mechanisms an important part of their security architectures.

6.5 Coverage Analysis

In the first step of coverage analysis, the Director identified technologies currently in the organization's security architecture. Table 6 - 11 shows a list of these technologies. In the next step, SAEM coverage analysis determined how the organization's security technologies mitigate the risk from the top six threats. Figure 6 - 3 shows that the organization is weak in detection mechanisms and that *Scanning* threats are only mitigated through the recover mechanism: auditing tools. Recall from the benefit analysis phase that the Director did not identify many security technologies. Other security managers identified other security technologies, such as Firewalls, as mitigating risks from *Scanning* threats. Since coverage analysis depends on the Director's identification of risk-mitigating security technologies in the benefit analysis phase, managers that are unfamiliar with security technology capabilities will see more gaps in the coverage analysis phase and there will be fewer technologies identified to fill the gaps.

The organization appears weak in detection mechanisms for its top threats; the hospital Corporate Information Officer (CIO) and the Director were considering adding an Intrusion Detection System to the Security Architecture. Figure 6 - 4 shows how Network and Host-based IDS mechanisms would fit into the security architecture, from a defense-in-depth perspective. The Operations Manager found Figures 6-3 and 6-4 interesting and informative, and felt that these figures supported the organization's need to integrate detection mechanisms into its architecture.

Table 6 - 11 Current Security Architecture Components

| Security Technology |
|------------------------------|
| Anti-virus Software |
| Auditing Tools |
| Cryptographic Cards/Hardware |
| Hardened OS |
| Hardware Lockdown |
| Log-on Limits |
| Modem Access Control |
| Packet Filter Firewalls |
| Proxy Firewalls |
| Software Lockout |
| URL Blockers |
| Virtual Private Network |

Figure 6 - 3 Current Coverage

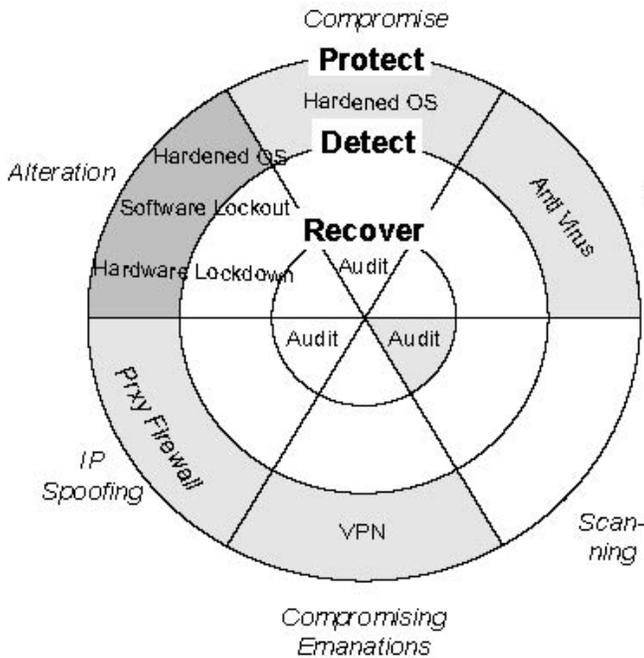
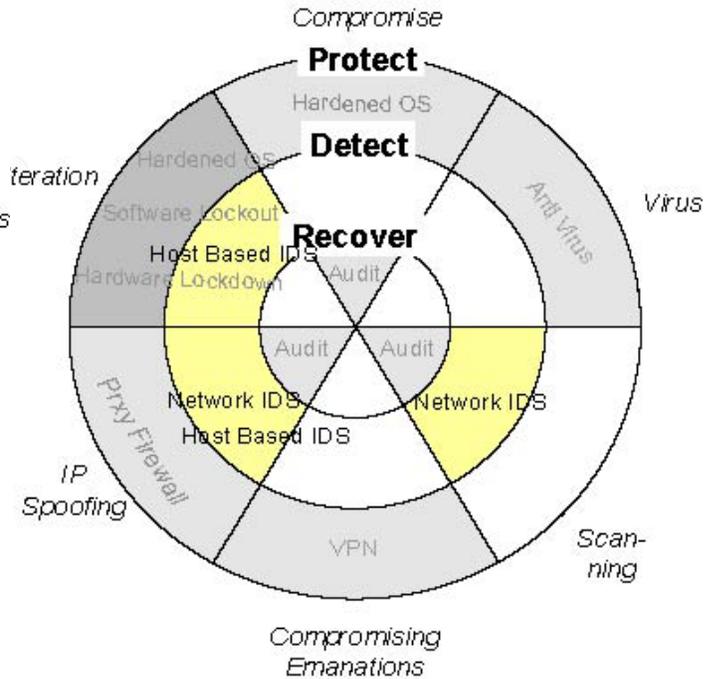


Figure 6 - 4 Current Coverage and IDS



6.6 Security Technology Tradeoff Analysis

In the final step of SAEM, the Director selected and ranked four technologies. He selected these technologies because he had been considering them for future integration into the security architecture. In addition, he identified several objectives, which he used to assess whether to select a security technology for inclusion in the security architecture. This section describes the results of using multi-attribute analysis techniques to help the Director decide which security technologies best meet his objectives.

6.6.1 Initial Iteration

First, the Director selected and ranked four security technologies for comparison: Vulnerability Assessment Scanner, Secure Email, Smart Cards, and Electronic Signature (in order). Ironically, he was considering Smart Cards even though he did not often identify them as risk-mitigating. During the interview, he noted that he thought Smart Cards would make the staff more efficient, but not necessarily provide any additional security.

Next, he identified four objectives which he used to select security technologies:

1. Ease of Maintenance
2. Purchase Cost
3. Impact to Productivity
4. Vulnerability

Ease of Maintenance refers to how complex the security technology is to implement and maintain. The operations staff has limited time and skills, so complex tools would be a drain on personnel resources. The Impact to Productivity refers to an increase of productivity for the medical staff. Therefore, technologies that help increase productivity are rated more highly than those that have little increase in productivity.

Next, he ranked each of the objectives. Table 6 - 12 shows the security technologies and the weights of the objectives. In addition, he ranked each of the technologies with respect to the objectives. The Director felt that the Vulnerability Assessment Scanner and Secure Email would not any impact to hospital staff productivity so those technologies received zero points for the Impact to Productivity objective.

Table 6 - 12 Initial Security Technology Tradeoff Assessment

| | | Tradeoff Attributes | | | | Tradeoff Ranking |
|---------------------|----------------------------------|---------------------|---------------|---------------|---------------------|---------------------|
| | | Ease of Maintenance | Purchase Cost | Vulnerability | Productivity Impact | |
| Security Technology | Rank | w = .10 | w = .25 | w = .35 | w = .30 | $\sum w_i v_i(x_i)$ |
| | Vulnerability Assessment Scanner | 25 | 25 | 40 | 0 | .20 |
| | Secure Email | 40 | 35 | 20 | 0 | .24 |
| | Smart Card | 25 | 15 | 30 | 60 | .34 |
| | E-Signature | 10 | 25 | 10 | 40 | .22 |

Table 6 - 12 also shows the multi-attribute analysis results, which depict a higher ranking for Smart Cards than ranked by the Director. In fact, the correlation between the multi-attribute analysis and the Manager's initial ranking is -.4.

When the Director reviewed the results, the analyst reminded him that he should prioritize these technologies as if he did not have other security technologies in place, specifically Secure User ID/Passwords. Based on the analyst's comments, he re-ranked his initial ordering. He did not change the tradeoff attribute values or weights. Table 6 - 13 shows the final results of the security tradeoff analysis. Most notably, the Director re-ranked Smart Cards number one and Secure Email number four. The final correlation was .8, which indicates a strong correlation

Table 6 - 13 Security Technology Ranking Comparisons

| RANK | SAEM Ranking | Manager’s Initial Ranking | Manager’s Revised Ranking |
|-------------|----------------------------------|----------------------------------|----------------------------------|
| 1 | Smart Card | Vulnerability Assessment Scanner | Smart Card |
| 2 | Secure Email | Secure Email | Secure Email |
| 3 | Vulnerability Assessment Scanner | Smart Card | Electronic Signature |
| 4 | Electronic Signature | Electronic Signature | Vulnerability Assessment Scanner |

6.7 Summary

Overall, the Director was very pleased with the results. Although the participants were not as familiar with the threats and security technologies as were other case-study participants, the Director reported that:

“The time we spent made me aware of attacks/security holes that I was not aware of. Your assessment prepared me for the Red Siren assessment and HIPAA. Before you showed I was behind as far as network security goes, and now I have a plan created that will get the hospital [sic] where we need to be for HIPAA”.

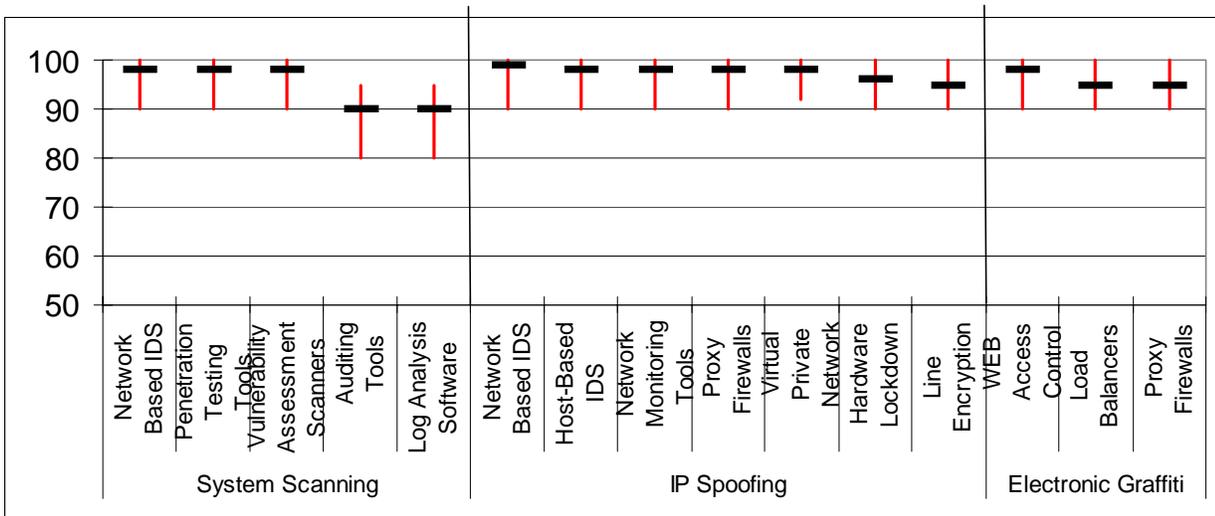
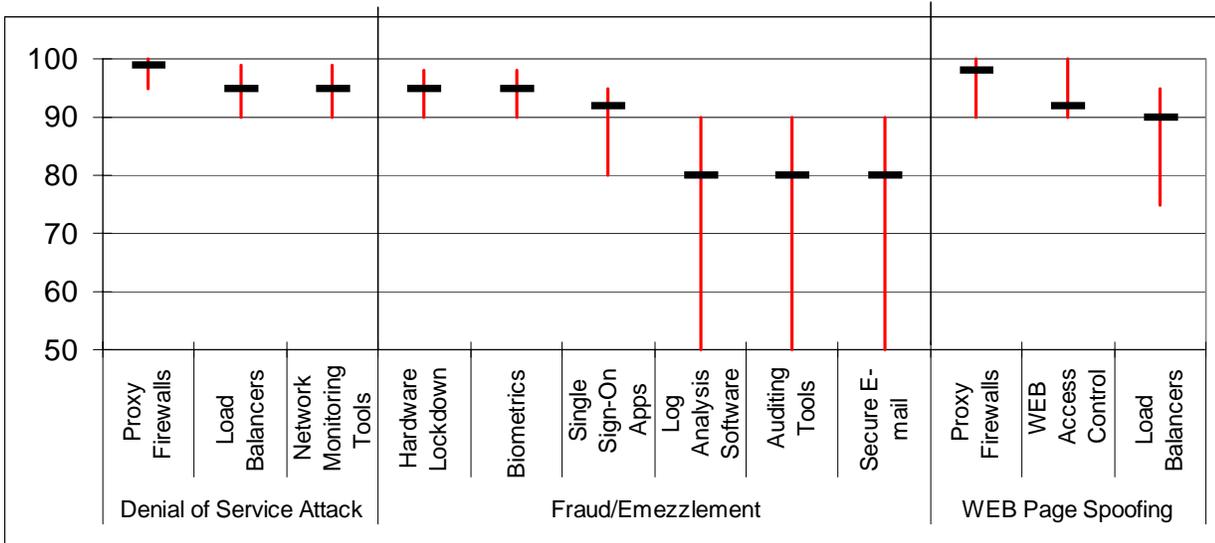
A review of the SAEM results shows that the SAEM risk assessment process moderately influenced the Director’s final prioritization of threats, and the satisfaction survey indicated that the Director gained significant insight about the value of some security technologies. In addition, he thought that the coverage analysis would make it much easier to explain why a particular security technology should be purchased. Finally, although the security tradeoff analysis ultimately showed a high positive correlation, the analyst had to remind the Director to make the initial assessment based on an absent security architecture.

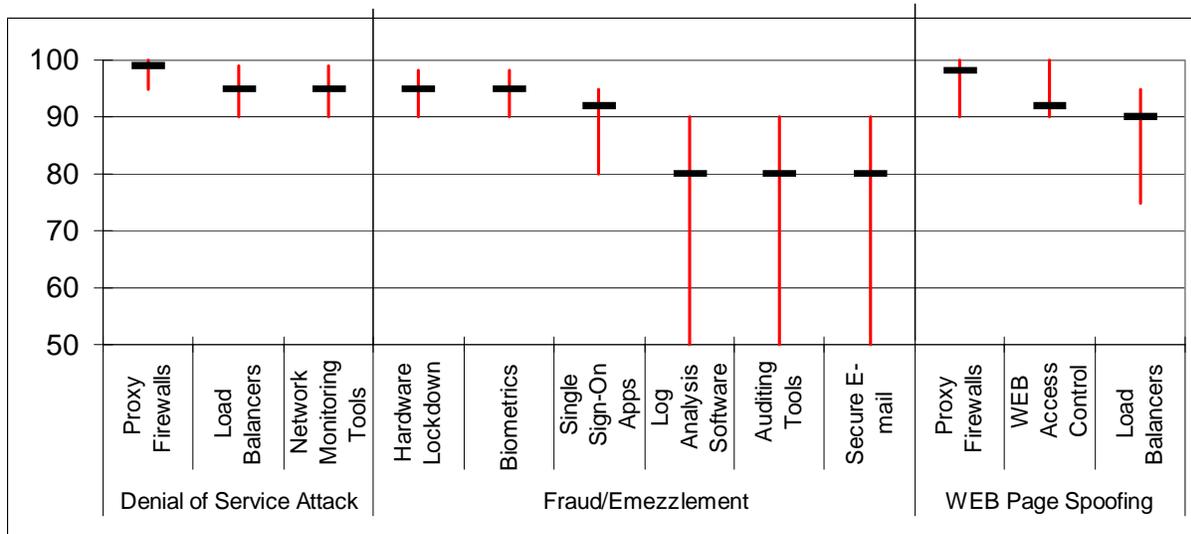
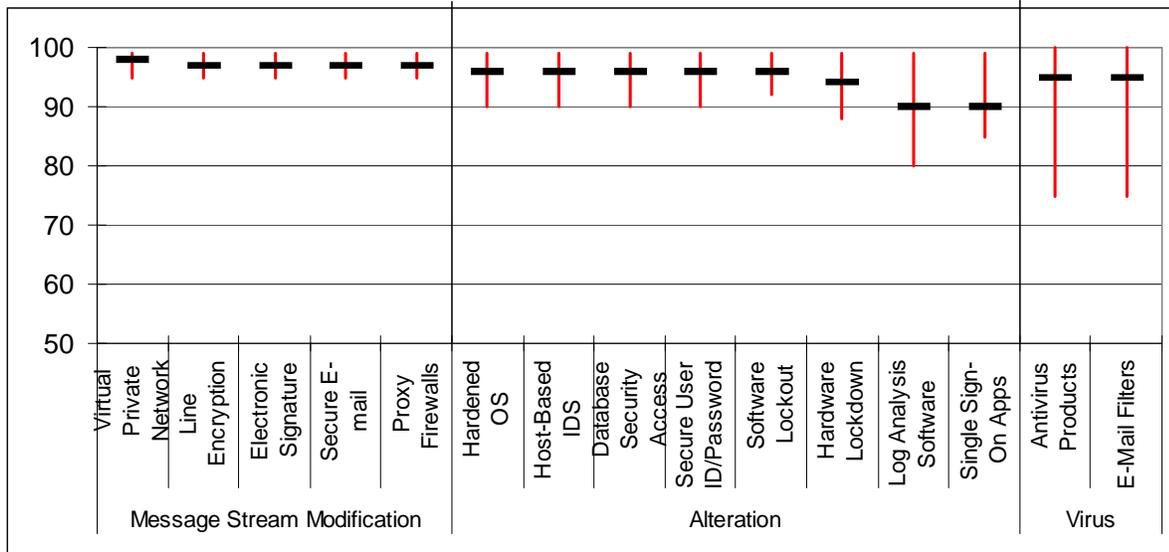
In this case study, the security tradeoff analysis highlighted the disadvantage of conducting the tradeoff analysis with the goal of evaluating the overall value of security technologies, rather than the added value that a particular security technology could provide if selected for an existing architecture. In reality, when security managers select security technologies, they consider the added value that a specific security technology provides given the organization’s security architecture. If an existing security technology already provides a particular function, such as authentication or authorization, then it is taken for granted in the security architecture and security technologies that provide new and different functionality are more appealing. Future research should modify the tradeoff analysis process to be able to capture the added benefit that a security technology provides.

Seven months after I delivered the final report, the Director reported that he had used the results of the coverage analysis to justify additional intrusion detection technologies for the security architecture. He spent \$85,000 on Network- and Host-based Intrusion Detection technologies. Furthermore, the Director purchased a System Scanning tool to help reduce the risk against scanning attacks.

BENEFIT ANALYSIS DATA

(Hospital Case Study)





SATISFACTION SURVEY

(Hospital Case Study)

The Technical Director completed this survey at the end of the SAEM process. He declined to make specific comments at the end of each section, but he did provide the overall assessment that I used in the Section 6.Summary.

I. Risk Assessment

| | | | | | | | |
|---|---------------------------------|----|----|-----------------|-----------------|-----------------|-----------------------------|
| How difficult was it.... | <i>not at all difficult</i> | | | | | | <i>very difficult</i> |
| to identify and initially rank the threats? | 0 | 1 | 2 | 3 | 4 | <u>X</u> | 6 |
| to estimate the frequency of attacks? | 0 | 1 | 2 | 3 | 4 | <u>X</u> | 6 |
| to estimate the outcomes? | 0 | 1 | 2 | <u>3</u> | 4 | 5 | 6 |
| | | | | | | | |
| How much insight did you gain about the organization's threats? | <i>none at all</i> | | | | | | <i>very much</i> |
| How much insight did you gain about the organization's outcomes? | 0 | 1 | 2 | 3 | <u>X</u> | 5 | 6 |
| How much did the risk assessment change your perception of the organization's risks? | 0 | 1 | 2 | 3 | <u>X</u> | 5 | 6 |
| How much easier would it be to explain the organization's risk priorities? | 0 | 1 | 2 | 3 | <u>X</u> | 5 | 6 |
| | | | | | | | |
| How strongly would you approve or disapprove of submitting the risk assessment ranking to your CIO for use in making decisions about risk management? | <i>strongly disapprove</i> | | | | | | <i>strongly approve</i> |
| | -3 | -2 | -1 | 0 | <u>X</u> | 2 | 3 |

| | | | | | | | |
|---|--------------------------------|----|----|---|---|-----------------|----------------------------|
| How satisfied are you with the risk rankings? | <i>very dissatisfied</i> -3 | -2 | -1 | 0 | 1 | <u>X</u> | <i>very satisfied</i> 3 |
|---|--------------------------------|----|----|---|---|-----------------|----------------------------|

Security Manager did not respond to survey questions

II. Benefit Analysis

| | | | | | | | |
|---|---------------------------------|----|----|---|----------|----------|-----------------------------|
| How difficult was it.... | <i>not at all difficult</i> | | | | | | <i>very difficult</i> |
| to identify and initially rank the security technologies? | 0 | 1 | 2 | 3 | 4 | 5 | <u>X</u> |
| to estimate the effectiveness of the technologies? | 0 | 1 | 2 | 3 | 4 | <u>X</u> | 6 |
| | | | | | | | |
| How much insight did you gain about the value that some security technologies provide? | <i>none at all</i> | | | | | | <i>very much</i> |
| | 0 | 1 | 2 | 3 | 4 | <u>X</u> | 6 |
| How much did the benefit analysis change your perception of the organization's security technologies? | 0 | 1 | 2 | 3 | <u>X</u> | 5 | 6 |
| How much easier would it be to explain why a particular security technology should be purchased? | 0 | 1 | 2 | 3 | <u>X</u> | 5 | 6 |
| | | | | | | | |
| How strongly would you approve or disapprove of submitting the benefit analysis results to your CIO for use in making decisions about spending financial resources? | <i>strongly disapprove</i> | | | | | | <i>strongly approve</i> |
| | -3 | -2 | -1 | 0 | <u>X</u> | 2 | 3 |
| | | | | | | | |
| How satisfied are you with the security technology rankings? | <i>very dissatisfied</i> | | | | | | <i>very satisfied</i> |
| | -3 | -2 | -1 | 0 | 1 | <u>X</u> | 3 |

Security Manager did not respond to survey questions.

III. Coverage Analysis

| | |
|---|--|
| <p>How much insight did you gain about the overall defense-in-depth coverage that your organization's current security technologies provide?</p> | <p><i>none at all</i> <i>very much</i></p> <p>0 1 2 3 <u>X</u> 5 6</p> |
| <p>How much did the coverage analysis change your perception of the organization's security status?</p> | <p>0 1 2 3 <u>X</u> 5 6</p> |
| <p>How much easier would it be to explain why a particular security technology should be purchased if the coverage analysis showed a gap?</p> | <p>0 1 2 3 4 <u>X</u> 6</p> |
| | |
| <p>How strongly would you approve or disapprove of submitting the coverage analysis results to your CIO for use in making decisions about spending financial resources?</p> | <p><i>strongly disapprove</i> <i>strongly approve</i></p> <p>-3 -2 -1 0 <u>X</u> 2 3</p> |
| | |
| <p>How satisfied are you with the coverage analysis?</p> | <p><i>very dissatisfied</i> <i>very satisfied</i></p> <p>-3 -2 -1 0 <u>X</u> 2 3</p> |

Security Manager did not respond to survey questions

CHAPTER 7. Government Case Study

7.1 Introduction

This chapter describes the participants, activities and results from the government case study. The multi-attribute analyst elicited the data for this case study through a series of interviews and questionnaires during a four-month period. The Government Agency's Deputy Director for the computer-incident response center and a staff of technical advisors provided the input for the risk-assessment phase, but only the Deputy Director provided the input for the benefit analysis and coverage assessment phases. The Director of Mission Assurance, the Deputy Director's supervisor, provided additional information for the risk assessment. Although she did not participate in the risk assessment process, she provided additional insight into the Deputy Director's threat prioritizations and outcome attribute weights. Overall, the participants in this case study demonstrated the most technical expertise and experience of all case study participants.

Although SAEM did not reveal any significant flaws in the organization's security architecture, the process helped the security staff validate previous decisions about the security architecture and the Director thought that the process would improve communication of their security requirements. The risk assessment highlighted the differences between the Deputy Director's expectations and the Director's beliefs about the organization's risks. The benefit analysis showed that the security technologies that helped reduce the organization's internal threats were ranked more highly, but were not considered as important to the Deputy Director as those technologies that helped him reduce his top threats. Although the coverage analysis did not reveal weaknesses in the security architecture, the Director of Mission Assurance found it very useful in communicating to other executives within the organization the value of past security investments. Tradeoff analysis supported the Deputy Director's preference for investing in Modem Access Control Mechanisms. From the Satisfaction Surveys (end of the chapter), the Director thought SAEM was useful, but would have liked to have seen more automation of the analysis.²⁴ The Deputy Director found it less helpful, but gained some insights about the organization's threats.

7.2 Case Study Description

The third case study is from a large civilian government organization that has almost 100,000 employees. The organization operates several large mainframe computers, which store sensitive information. Employees access applications running on the mainframes and have limited connections to the Internet. This government organization has a large staff to run the information system and dedicated staff to ensure the system's security. In addition, the

²⁴ The Director did not fully participate in each of the processes so did not experience ASESS, the semi-automated tool for SAEM.

organization has an incident response center, which collects statistics on selected security incidents.²⁵

The day-to-day security responsibility primarily falls on the Deputy Director, who handles the operation of the computer incident response center and advises the Director on security policies, technologies, and risk-mitigation procedures. The Deputy Director's sole focus is security, and the Director of Mission Assurance is more broadly responsible for balancing security requirements with customer mission assurance requirements.

Although the security budget is very large relative to the other case studies, the Deputy Director must justify expenditures. Each year, the incident response center collects security incident statistics that enable the Director of Mission Assurance to justify additional funds to strengthen its security architecture. Currently, of the security technologies presented during the SAEM benefit analysis process, the government organization had all but two technologies in its security architecture.

The government has strict policies that prohibit the use of computers for personal use, such as for games, personal web maintenance and web email. In addition, employees browsing through sensitive data are fired if they are not explicitly authorized to view the data. Most recently, the Deputy Director had been focused on reducing the risk from identity theft, but the organization has experienced many security incidents from all types of threats. The organization began tracking security incidents about two years ago, but did not have a previous risk assessment for the information system.

7.2.1 Participants

The Deputy Director was the organization's primary participant in this case study. The Deputy Director was responsible for developing security requirements and recommending changes to the security architecture, such as the purchase of new security technologies, and recommending security policy changes. The Deputy Director was also responsible for the organization's incident response center, which collects and investigates information about security incidents. The Deputy Director used the data collected by the incident response center to estimate the frequency and outcomes of some of the threats. Individuals from the incident response center assisted the Deputy Director by providing technical advice and expertise about threats. Generally, the participants were very familiar with all of the threats and security technologies presented.

7.3 The Risk Assessment

This section outlines the details, sequence of events, and data collected during the risk-assessment phase of the case study. The Deputy Director provided his initial threat ranking at the start of the risk-assessment phase, but the analyst collected the threat frequency and outcome estimates using a questionnaire. After completing the questionnaire, the analyst presented the SAEM results to the Deputy Director, who used them to inform revisions to his initial risk priorities and the risk assessment input data.

²⁵ The incident response center did not have statistics on all threats presented at the time of the risk assessment.

Table 7 - 1 Initial Threat Rankings

| Rank | Threat |
|------|-------------------------------|
| 1 | Theft |
| 2 | Contamination |
| 3 | Compromise |
| 4 | Password Guessing |
| 5 | Signal Interception |
| 6 | Message Stream Modification |
| 7 | Alteration |
| 8 | Logic Bomb |
| 9 | Procedural Violation |
| 10 | Browsing |
| 11 | Virus |
| 12 | Trojan Horse |
| 13 | Data Entry Error |
| 14 | Personal Abuse |
| 15 | System Scanning |
| 16 | Password Nabbing |
| 17 | Trap Door |
| 18 | IP Spoofing |
| 19 | Electronic Graffiti |
| 20 | Distributed Denial of Service |
| 21 | Denial of Service |
| 22 | Fraud/Embezzlement |
| 23 | WEB Page Spoofing Vandalism |
| 24 | Cryptographic Compromise |

7.3.1 Initial Iteration

Table 7 - 1 shows the Deputy Director’s initial ranking of threats. He eliminated two threats, Vandalism and Compromising Emanations, from the analyst’s initial threat list because he thought they were redundant with other threats on the analyst’s list. The organization experiences many security compromises each year, but the Deputy stated that he was mostly concerned about security compromises that could lead to identity theft.

7.3.1.1 Outcome Attributes

The Deputy identified three attack consequences that were important to the organization: damaged *public reputation*, *increased oversight*, and inability to conduct *administration* functions. Since this government agency had been struggling for several years to improve its public image, it was concerned about security incidents that could affect its public reputation. As with many government agencies, security compromises can spark increased oversight from the Government Accounting Office (GAO) or Congress. Such additional oversight can be time consuming for key personnel and disruptive to the organization’s operations.

Finally, security incidents can prevent the organization from conducting its primary administration function, which is the reason for its existence.

The participants used the 7-point scale described in Chapter 4 to assess the impact of a security compromise on their concerns related to Public Reputation, Increased Oversight, and Administration. Table 7 - 2 represents the results of the Deputy Director’s ranking of the outcome attributes using the swing weight method. In this case study, the Deputy allocated 100 points among each of the outcome attributes, rather than allot 100 points to the most important attribute and rank the others relative to the most important attribute.

Table 7 - 2 Outcome Attribute Weights

| Attribute | Rank | Weight |
|---------------------|------|--------|
| Public Reputation | 2 | 0.25 |
| Increased Oversight | 3 | 0.25 |
| Administration | 1 | 0.5 |

7.3.1.2 Outcome and Frequency Estimates

Table 7 - 3 is the Deputy Director's estimates of frequency and outcome values for each of the 24 threats. The table is sorted according to expected frequency. The Deputy Director estimated that the first three threats-- *Browsing*, *Personal Computer Abuse*, and *Procedural Violations*-occur, on average, hourly within the organization. These three threats are internal threats, which he estimated based on the number of employees in the organization and how often he believes an employee violates the organization's security policies.

He expects all other threats to occur with considerably less frequency, but if they occur, they are *most likely* to result in more severe consequences. For example, the average estimate for the public reputation consequence is 4.7, which means that the Security Manger estimated that the *most likely* impact to the organization's public reputation from any attack would be between moderate and moderately severe damage using the Likert Scale. In contrast, he estimated that the average impact from any of the three most frequently occurring attack would result in mild to moderately mild damage to the organization's public image (an average of 2.7 on the Likert Scale).

Table 7 - 3 Estimated Outcome and Frequency Values

| Threats | Frequency/Year | | | Public Reputation | | | Increased Oversight | | | Administration | | |
|-------------------------------|----------------|-----------|------------|-------------------|-------------|------|---------------------|-------------|------|----------------|-------------|------|
| | Low | Exp | High | Low | Most Likely | High | Low | Most Likely | High | Low | Most Likely | High |
| Personal Computer Abuse | 268,800 | 2,688,000 | 26,880,000 | 1 | 3 | 7 | 3 | 4 | 6 | 1 | 2 | 4 |
| Browsing | 268,800 | 2,688,000 | 26,880,000 | 1 | 2 | 3 | 1 | 3 | 3 | 1 | 1 | 1 |
| Procedural Violation | 13,440 | 53,760 | 268,800 | 2 | 3 | 4 | 3 | 4 | 5 | 3 | 4 | 5 |
| Compromise | 104 | 156 | 520 | 4 | 6 | 7 | 3 | 6 | 7 | 3 | 4 | 6 |
| Contamination | 104 | 156 | 520 | 2 | 5 | 7 | 3 | 5 | 7 | 3 | 6 | 7 |
| Alteration | 24 | 60 | 120 | 3 | 6 | 7 | 3 | 6 | 7 | 3 | 4 | 6 |
| Theft | 12 | 60 | 120 | 2 | 5 | 7 | 2 | 5 | 7 | 2 | 6 | 7 |
| Trojan Horse | 0 | 52 | 156 | 2 | 5 | 7 | 3 | 5 | 6 | 2 | 3 | 5 |
| Virus | 0 | 52 | 156 | 1 | 2 | 3 | 2 | 3 | 4 | 1 | 2 | 4 |
| Password Guessing | 24 | 36 | 60 | 3 | 6 | 7 | 3 | 6 | 7 | 3 | 4 | 6 |
| System Scanning | 12 | 24 | 36 | 2 | 5 | 7 | 3 | 4 | 6 | 1 | 2 | 4 |
| Electronic Graffiti | 0 | 24 | 36 | 2 | 5 | 7 | 3 | 5 | 7 | 1 | 3 | 6 |
| IP Spoofing | 0 | 24 | 36 | 2 | 4 | 6 | 3 | 4 | 5 | 2 | 3 | 6 |
| Data Entry Error | 0 | 12 | 36 | 2 | 6 | 7 | 3 | 6 | 7 | 2 | 3 | 6 |
| Password Nabbing | 0 | 12 | 36 | 2 | 5 | 6 | 3 | 5 | 6 | 2 | 3 | 5 |
| Trap Door | 0 | 12 | 36 | 2 | 5 | 7 | 3 | 5 | 7 | 2 | 3 | 6 |
| Fraud/Embezzlement | 0 | 12 | 36 | 2 | 5 | 7 | 3 | 5 | 7 | 3 | 4 | 6 |
| WEB Page Spoofing | 0 | 12 | 24 | 2 | 5 | 6 | 3 | 5 | 6 | 0 | 3 | 6 |
| Cryptographic Compromise | 0 | 12 | 24 | 2 | 5 | 7 | 3 | 5 | 7 | 2 | 3 | 6 |
| Distributed Denial of Service | 0 | 12 | 24 | 2 | 4 | 6 | 3 | 5 | 6 | 2 | 3 | 6 |
| Denial of Service Attack | 0 | 12 | 24 | 2 | 4 | 6 | 3 | 5 | 6 | 2 | 3 | 6 |
| Signal Interception | 0 | 3 | 5 | 3 | 5 | 7 | 3 | 5 | 7 | 2 | 3 | 6 |
| Message Stream Modification | 0 | 2 | 5 | 2 | 5 | 6 | 3 | 5 | 7 | 2 | 3 | 6 |
| Logic Bomb | 0 | 2 | 3 | 3 | 6 | 7 | 3 | 6 | 7 | 3 | 4 | 6 |

7.3.1.3 Initial Results

Table 7 - 4 compares the results from SAEM using the Deputy Director's initial estimated values. SAEM highly ranked the three most frequently occurring threats, i.e., *Personal Computer Abuse*, *Browsing*, and *Procedural Violations*. In addition, the Deputy Director ranked *Personal Computer Abuse* considerably lower (13th) than SAEM. The correlation between the SAEM's rank (S) and the Deputy's initial rank (I) was .61, which indicates moderate correlation between the two ranks. The last column in Table 7-4 highlights the threats with the greatest difference between SAEM's rank and the Deputy's rank.

7.3.2 Refinement

The Deputy Director reviewed the risk assessment results. In addition, the analyst presented a comparison of some of the threats. Table 7 - 5 table shows the Deputy Director's initial estimates for some of the threats that had rankings significantly different than SAEM's ranking. During the review of the SAEM results and the data in Table 7 - 5, the analyst pointed out that the Deputy Director had initially ranked *Theft* number one, but that *Contamination* incidents had similar consequences and occurred more frequently than *Theft*. The Deputy Director had initially ranked *Contamination* as the 2nd highest threat. Table 7 - 5 also shows that *Fraud/Embezzlement* occur with the same frequency and consequence as *Password Guessing*, but *Fraud/Embezzlement* was ranked 22nd and *Password Guessing* was ranked 4th.

Table 7 - 4 SAEM and Director's Threat Rankings

| Threat | Relative Threat Index ²⁶ | SAEM Rank (S) | Initial Rank (I) | S-I |
|-------------------------------|-------------------------------------|---------------|------------------|-----|
| Personal Computer Abuse | 100 | 1 | 14 | 13 |
| Browsing | 73 | 2 | 10 | 8 |
| Procedural Violation | 26 | 3 | 9 | 6 |
| Compromise | 1.1 | 4 | 2 | 2 |
| Contamination | 0.005 | 5 | 3 | 2 |
| Alteration | 0.004 | 6 | 7 | 1 |
| Theft | 0.002 | 7 | 1 | 6 |
| Trojan Horse | 0.002 | 8 | 12 | 4 |
| Password Guessing | 0.001 | 9 | 4 | 5 |
| Virus | 0.001 | 10 | 11 | 1 |
| System Scanning | 0.001 | 11 | 15 | 4 |
| Data Entry Error | 0.001 | 12 | 22 | 10 |
| Fraud/Embezzlement | <0.001 | 13 | 13 | 0 |
| Trap Door | <0.001 | 14 | 17 | 3 |
| Password Nabbing | <0.001 | 15 | 16 | 1 |
| Signal Interception | <0.001 | 16 | 5 | 11 |
| Logic Bomb | <0.001 | 17 | 8 | 9 |
| Electronic Graffiti | <0.001 | 18 | 6 | 12 |
| Message Stream Modification | <0.001 | 19 | 19 | 0 |
| IP Spoofing | <0.001 | 20 | 18 | 2 |
| Cryptographic Compromise | <0.001 | 21 | 24 | 3 |
| WEB Page Spoofing | <0.001 | 22 | 21 | 1 |
| Distributed Denial of Service | <0.001 | 23 | 20 | 3 |
| Denial of Service Attack | <0.001 | 24 | 23 | 1 |

²⁶ For each threat, the relative threat index is the threat index value normalized between 0 and 100.

Table 7 - 5 Threat Comparisons

| | | Outcome Attributes | | | | | |
|------------------------------------|---------------------------------------|--------------------|-----------|------------|--------------------|---------------------|-----------------|
| | | Frequency/Year | | | Public Reputation | Increased Oversight | Administration |
| | | low | exp | high | Most Likely Values | | |
| Threat (SAEM Rank-Initial Rank) | Browsing (2-10) | 268,800 | 2,688,000 | 26,880,000 | Mild | Moderately Mild | Moderately Mild |
| | Procedural Violations (3-9) | 13,440 | 53,760 | 268,800 | Moderately Mild | Moderate | Moderate |
| | Contamination (4-2) | 104 | 156 | 520 | Moderately Severe | Moderately Severe | Severe |
| | Theft (7-1) | 12 | 60 | 120 | Moderately Severe | Moderately Severe | Severe |
| | Password Guessing (9-4) | 24 | 36 | 60 | Severe | Severe | Moderate |
| | Fraud/Embezzlement (12-22) | 0 | 12 | 36 | Moderately Severe | Moderately Severe | Moderate |

Based on a review of the SAEM results, the Deputy Director made the following changes:

- *Fraud/Embezzlement* initial rank changed to 3rd.
- *Browsing* Frequency changed to 2-low/200-most likely/1000-high per hour²⁷
- *Procedural Violation* “most likely” consequences were changed to:
 - Public Reputation - None (1²⁸)
 - Increased Oversight - Moderate (4)
 - Administration - None (1)
- *Procedural Violation* “high” consequences were changed to:
 - Public Reputation - Mild (2)
 - Increased Oversight - Moderately Severe (5)
 - Administration - Mild (2)
- *Virus* Frequency changed to 50-low/90-most likely/500-high per day

In addition to the changes mentioned above, the Deputy Director was uncertain about the ranking of *Personal Computer Abuse* since his ranking differed significantly from SAEM’s ranking, 14th and 1st respectively. The Deputy refined the definition of *Personal Computer Abuse* to consist of six different types of security incidents. Table 7-6 shows the six new security incidents and their estimated frequencies²⁹ and outcomes.

²⁷ These values are converted to yearly rates in all of the tables.

²⁸ Likert Scale rating

²⁹ Notice that the frequencies are estimated in incidents per day.

Table 7 - 6 Personal Computer Abuse Threats

| | | Outcome Attribute Values | | | | | | | | | | | |
|---------|--------------------------|--------------------------|-------------|------|---------------------|-------------|------|----------------|-------------|------|---------------|--------|--------|
| | | Public Reputation | | | Increased Oversight | | | Administration | | | Frequency/Day | | |
| | | low | most likely | high | low | most likely | high | low | most likely | high | low | exp | high |
| THREATS | Web Administration | 1 | 1 | 3 | 1 | 1 | 3 | 1 | 1 | 1 | 300 | 1,000 | 3,000 |
| | Pornographic Downloads | 1 | 2 | 5 | 1 | 2 | 4 | 1 | 1 | 3 | 100 | 200 | 1,000 |
| | Games | 1 | 1 | 2 | 1 | 1 | 3 | 1 | 1 | 3 | 50 | 100 | 500 |
| | Chat | 1 | 1 | 2 | 1 | 1 | 3 | 1 | 1 | 3 | 80 | 200 | 1,000 |
| | Web-based Personal Email | 1 | 1 | 3 | 1 | 1 | 5 | 1 | 1 | 2 | 2,000 | 10,000 | 30,000 |
| | Web Surfing | 1 | 1 | 3 | 1 | 1 | 2 | 1 | 1 | 2 | 1,000 | 3,000 | 7,000 |

Using the Deputy Director’s new information concerning *Personal Computer Abuse* incidents, the analyst used ASESS to re-compute the threat indexes of the six *Personal Computer Abuse* threats. **Table 7 - 7** shows the threat indexes for these threats. *Personal Web-based Email* ranked first with a threat index of 603,875 – a relative threat index of 100, still the highest threat index of all threats, so *Personal Computer Abuse* threats remained the organization’s number one threat³⁰.

Table 7 - 7 Threat RTI's

| Threat | RTI |
|--------------------------|-------|
| Web-based Personal Email | 100 |
| Web Surfing | 12 |
| Web Administration | 6 |
| Chat | 1.0 |
| Games | 0.500 |
| Pornographic Downloads | 0.006 |

Table 7 - 8 shows the results of SAEM’s analysis with the Deputy Director’s changes. SAEM continued to rank *Personal Computer Abuse* as the most significant threat to the organization based on the Deputy Director’s revised input. In addition, SAEM ranked *Viruses* as the fourth most significant threat. After reviewing the results of this iteration, the Deputy Director continued to believe that *Theft* was still the most significant threat to the organization. The correlation between the final SAEM risk assessment and the Deputy Director’s revised rankings is .57, slightly lower than the SAEM’s initial ranking.

³⁰ Since SAEM rated one of the Personal Computer Abuse threats as the organization’s most significant threat, I continued using Personal Computer Abuse, rather than the sub-threats, since using the sub-threats would have required the Deputy Director to re-rank all of the threats and his availability was extremely limited.

Table 7 - 8 Revised SAEM Threat Rankings

| Threat | Relative Threat Index | SAEM Initial Rank (SI) | SAEM Final Rank (SF) | Deputy's Initial Rank (MI) | Deputy's Final Rank (MF) | SF-SI | MF-MI |
|-------------------------------|-----------------------|------------------------|----------------------|----------------------------|--------------------------|-------|-------|
| Personal Computer Abuse | 100 | 1 | 1 | 14 | 15 | 0 | 1 |
| Browsing | 12 | 2 | 2 | 10 | 11 | 0 | 1 |
| Procedural Violation | 2 | 3 | 3 | 9 | 10 | 0 | 1 |
| Virus | 0.9 | 10 | 4 | 11 | 12 | 7 | 1 |
| Compromise | 0.011 | 4 | 5 | 2 | 4 | 1 | 2 |
| Contamination | 0.010 | 5 | 6 | 3 | 2 | 1 | 1 |
| Alteration | 0.004 | 6 | 7 | 7 | 8 | 1 | 1 |
| Theft | 0.004 | 7 | 8 | 1 | 1 | 1 | 0 |
| Trojan Horse | 0.003 | 8 | 9 | 12 | 13 | 1 | 1 |
| Password Guessing | 0.003 | 9 | 10 | 4 | 5 | 1 | 1 |
| System Scanning | 0.001 | 11 | 11 | 15 | 16 | 0 | 1 |
| Data Entry Error | 0.001 | 12 | 12 | 22 | 14 | 0 | 8 |
| Fraud/Embezzlement | 0.001 | 13 | 13 | 13 | 3 | 0 | 10 |
| Trap Door | 0.001 | 14 | 14 | 17 | 18 | 0 | 1 |
| Password Nabbing | 0.001 | 15 | 15 | 16 | 17 | 0 | 1 |
| Signal Interception | 0.0 | 16 | 16 | 5 | 6 | 0 | 1 |
| Logic Bomb | 0.0 | 17 | 17 | 8 | 9 | 0 | 1 |
| Electronic Graffiti | 0.0 | 18 | 18 | 6 | 20 | 0 | 14 |
| Message Stream Modification | 0.0 | 19 | 19 | 19 | 7 | 0 | 12 |
| IP Spoofing | 0.0 | 20 | 20 | 18 | 19 | 0 | 1 |
| Cryptographic Compromise | 0.0 | 21 | 21 | 24 | 24 | 0 | 0 |
| WEB Page Spoofing | 0.0 | 22 | 22 | 21 | 23 | 0 | 2 |
| Distributed Denial of Service | 0.0 | 23 | 23 | 20 | 21 | 0 | 1 |
| Denial of Service Attack | 0.0 | 24 | 24 | 23 | 22 | 0 | 1 |
| Total | | | | | | 13 | 64 |

7.3.3 Additional Refinement

Although the Deputy Director did not offer any additional information that might have helped resolve the differences between the SAEM ranks and his revised ranks, the Director of Mission Assurance reviewed the two final rankings (SAEM final ranking and the Deputy Director's final ranking) and confirmed that SAEM had more accurately ranked the organization's top five threats. The Director of Mission Assurance explained that recently the Deputy Director had been most concerned about identity theft, fraud, and privacy compromises. These are threats for which he had not achieved an acceptable level of risk mitigation. Furthermore, the Director believed that if he had not been able to bring the internal threats to an acceptable level of risk mitigation, then he would be more concerned about the *Personal Computer Abuse*, *Browsing* and *Procedural Violations* than *Theft* and *Fraud/Embezzlement*.

The Director of Mission Assurance also disagreed with some of the Deputy Director’s estimates. The Director asked that we re-evaluate the risk assessment using generally higher attribute weights and lower frequency estimates for *Personal Computer Abuse* and *Browsing*. She also changed the outcome values for these two threats. Table 7-9 shows the changes she requested and the difference between the Deputy Director’s estimates and the Director’s estimates. Overall, she stated that if the Deputy Director’s estimated frequencies were correct, then the agency would have stopped operating. Therefore, she lowered the frequency estimates.

Table 7 - 9 Director’s Changes

| THREAT | Outcome Attribute Values | | | | | | | | | | | |
|-----------------------------------|--------------------------|-------------|------|---------------------|-------------|------|----------------|-------------|------|----------------|------------|-------------|
| | Public Reputation | | | Increased Oversight | | | Administration | | | Frequency/Year | | |
| | low | most likely | high | low | most likely | high | low | most likely | high | low | exp | high |
| Personal Computer Abuse | 3 | 5 | 7 | 3 | 5 | 7 | 1 | 2 | 3 | 700 | 1300 | 1,820 |
| Change in Personal Computer Abuse | 2 | 2 | 0 | 0 | 1 | 1 | 0 | 0 | -1 | -268100 | -2686700 | -26878180 |
| Browsing | 3 | 5 | 7 | 3 | 5 | 7 | 1 | 2 | 3 | 12 | 60 | 120 |
| Change in Browsing | 2 | 3 | 4 | 2 | 2 | 4 | 0 | 1 | 2 | -268,788 | -2,687,940 | -26,879,880 |

Using the Director’s estimates, the outcomes for *Personal Computer Abuse* and *Browsing* are now more severe and the frequency of *Browsing* was reduced. These changes resulted in a different prioritization of the threats as shown in Table 7 - 10. Most notably, *Personal Computer Abuse* dropped to third, and *Browsing* dropped in ranking to 7th. In addition, *Procedural Violations* and *Viruses* rose to first and second respectively.

Although the Director of Mission Assurance did not closely review all of the threats and the Deputy Director’s estimates, she felt that her changes resulted in a more accurate prioritization of the organization’s most significant threats.

7.3.4 Regression Analysis

Regression analysis indicates how well the SAEM model can predict the Deputy Director’s final threat ranking. The R²-value, also known as the *coefficient of determination*, measures the percentage of variation in the values of the dependent variable (Deputy’s Final Ranking) that can be explained by the independent variable (the SAEM Final Ranking). The R²-value for the final SAEM risk assessment model is .34. This low R²-value indicates that only 34% of the variation in Deputy’s final ranking can be explained by the final SAEM risk assessment model. The remaining 66% of the variation is due to random or unknown variability.

Figure 7-1 shows that the Deputy Director changed his final ranking very little from his initial ranking. The correlation between his initial ranking and his final ranking is .9, which means that his initial ranking is very highly correlated with his final ranking. In contrast, the correlation of the SAEM Final Ranking with the Deputy's Final Ranking is .57, which means that the SAEM Final Ranking is significantly less correlated with the Deputy's Final Ranking than the Deputy's initial ranking. Finally, since the Deputy changed Fraud/Embezzlement from 22nd to 13th ranked, the threats ranked 13th or higher were shifted higher in rank. This slight shifting of these threats caused the correlation to be slightly worse than the initial rank correlation, .61 down to .57, but this is not a significant drop in the correlation.

Although the final SAEM risk assessment appears to be a moderate predictor of the final Deputy's ranking, it is possible that the Deputy's initial ranking is a better predictor of his final ranking, or that the Deputy's initial rank is a significant factor in predicting the Deputy's final rank. Table 7 - 11 shows the results of regression analysis using both the Final risk assessment and the Initial Deputy's rankings as predictors of the final rankings. The *Coefficients* column shows the prediction equation³¹ and the *t Stat* column shows the ratio between the coefficient and the standard error. The *P-value* is the probability of a t-value this large or larger. Therefore, a P-value less than .05 indicates that the coefficient is significant. Table 7 - 11 shows that only the Deputy Director's initial ranking coefficient is significant in predicting the Deputy's final rank.

Table 7 - 10 SAEM Results Using Director's Changes

| Threat | Relative Threat Index | SAEM Final Ranking | Deputy's Final Ranking |
|-------------------------------|-----------------------|--------------------|------------------------|
| Procedural Violation | 100 | 1 | 9 |
| Virus | 54 | 2 | 11 |
| Personal Computer Abuse | 26 | 3 | 14 |
| Compromise | 18.3 | 4 | 4 |
| Contamination | 0.350 | 5 | 2 |
| Alteration | 0.306 | 6 | 7 |
| Browsing | 0.136 | 7 | 10 |
| Theft | 0.136 | 8 | 1 |
| Trojan Horse | 0.121 | 9 | 12 |
| Password Guessing | 0.097 | 10 | 5 |
| System Scanning | 0.083 | 11 | 15 |
| Data Entry Error | 0.049 | 12 | 13 |
| Trap Door | 0.027 | 13 | 17 |
| Fraud/Embezzlement | 0.026 | 14 | 22 |
| Password Nabbing | 0.026 | 15 | 16 |
| Signal Interception | 0.023 | 16 | 6 |
| Logic Bomb | 0.007 | 17 | 9 |
| Electronic Graffiti | 0.005 | 18 | 19 |
| Message Stream Modification | 0.004 | 19 | 6 |
| IP Spoofing | 0.004 | 20 | 18 |
| Cryptographic Compromise | 0.003 | 21 | 24 |
| WEB Page Spoofing | 0.002 | 22 | 23 |
| Denial of Service Attack | 0.002 | 23 | 21 |
| Distributed Denial of Service | 0.002 | 24 | 21 |

³¹The prediction equation from the coefficients in the second column is:
 $Threat_{Rank} = .94 + .07(Initial\ SAEM\ Rank) - .086(Initial\ DD)$

Figure 7 - 1 Initial Ranking Compared to Final Rankings

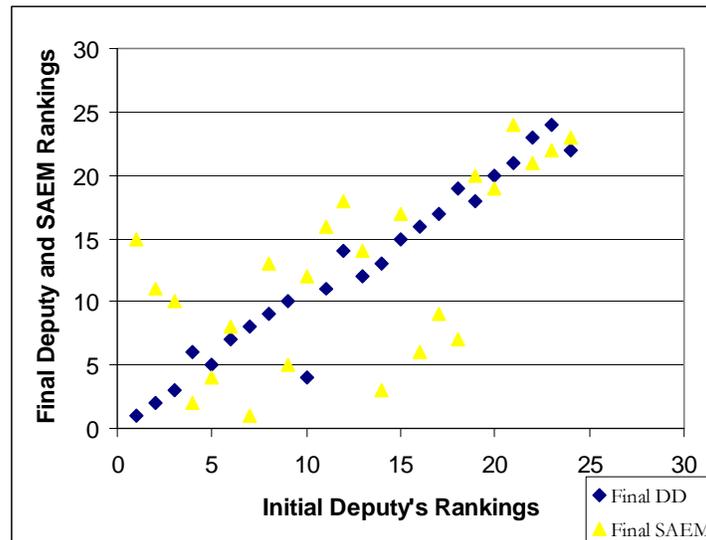


Table 7 - 11 Regression Analyses for Risk Assessment Results

| | Coefficients | Standard Error | t Stat | P-value |
|--------------|--------------|----------------|--------|---------|
| Intercept | 0.94 | 1.52 | 0.62 | 0.54 |
| Initial SAEM | 0.07 | 0.45 | 0.15 | 0.88 |
| Final SAEM | 0.00 | 0.44 | 0.00 | 1.00 |
| Initial DD | 0.86 | 0.13 | 6.79 | 0.00 |

7.3.5 Risk Assessment Insight

One of the key insights that the risk-assessment phase highlighted for the government organization was the discrepancy between the Deputy Director’s risk assessment and the Director of Mission Assurance’s risk assessment estimates. The greatest difference between the Deputy Director’s risk assessment and Director’s risk assessment was the frequency of internal threats. Interestingly, during the final presentation of the SAEM results, the Director of Mission Assurance and the Deputy Director said that they had just completed a recent analysis of the organization’s incident response center data. The data confirmed that SAEM had correctly identified four out of the top five organizational threats: *Procedural Violations*, *Viruses*, *Compromise*, and *Contamination*. They indicated that *Personal Computer Abuse* incidents were significant, but did not rank within the top five threats. They did not recall the ranking of any other threats.

7.3.6 Risk Assessment Summary

Although the risk assessment did not influence the Deputy Director’s assessment of the organization’s threat priorities, the Director of Mission Assurance felt that the assessment was insightful. According to the Director’s comments in the Satisfaction Survey, the process encouraged her to think more broadly about the organization’s threats. In contrast, the Deputy Director did not find it as useful, most likely because he had difficulty estimating threat frequencies and outcome values

under the assumption that the organization had not established a security architecture.³² In the future, it might be more meaningful to the Deputy Director to conduct the risk-assessment phase given the existing security architecture.

It is difficult to determine the why there was such a large discrepancy between the Director’s and Deputy Director’s estimates. The Director had been with the organization several years longer than the Deputy Director, but the organization had only recently started collecting incident statistics. I believe that this discrepancy highlights a typical problem between security specialists and information system executives, i.e., these executives are often skeptical about risk estimates. A key advantages that SAEM offers is that the basis for these estimates is made explicit, executives and security managers can see the impact these estimates have on the organization’s threat priorities, and make decisions about whether to collect better information so that they are more confident in the results.

7.4 Benefit Analysis

After completing the risk assessment, the Deputy Director completed the benefit analysis phase of SAEM. The benefit analysis interviews took approximately two hours to elicit the Deputy Director’s initial selection of the organization’s most effective security technologies and to determine which security technologies mitigated each of the threats. The Deputy Director completed a questionnaire, which provided the percentage effectiveness against threats for each security technology. Although the Deputy found it very tedious to estimate the effectiveness of the security technologies for each threat, the Director of Mission Assurance found it useful.

7.4.1 Initial Iteration

Initially, the Deputy Director selected ten security technologies that he believed were the most effective in mitigating the organization’s overall security risks. Table 7 - 12 shows his selection of security technologies. Hardened OS and Secure Operating Systems are ranked first and second respectively. These two security technologies are very similar in function and were selected as risk-mitigating for almost every threat.

Table 7 - 12 Deputy Director’s Most Effective Security Technologies

| Initial Order | Security Technology |
|---------------|------------------------------------|
| 1 | Hardened OS |
| 2 | Secure Operating Systems |
| 3 | Packet Filter Firewall |
| 4 | Proxy Firewall |
| 5 | Database Security Access Control |
| 6 | Anti-virus Software |
| 7 | Hardware Lockout |
| 8 | Network Intrusion Detection System |
| 9 | Log Analysis Software |
| 10 | Vulnerability Assessment Scanner |

Next, the Deputy Director identified risk-mitigating security technologies and estimated their effectiveness against each of the threats. Appendix A shows their estimates for each threat. On average, the Deputy Director selected twenty-two security technologies per threat, the highest average among all of the case studies. Recall from Chapter 4 that the Benefit Analysis phase uses the risk assessment to calculate a threat index change for each technology. The overall threat index change is determined based on the cumulative effectiveness of the security technology against each threat. Therefore, SAEM ranks

³² Recall from Chapter 4, that all participants were asked to estimate frequencies and outcomes as if the current security architecture did not exist.

security technologies that are effective against threats with high threat indexes more highly than technologies that are effective against threats with low threat indexes. Table 7 - 13 shows the benefit analysis results for the most highly ranked security technologies. The 15 security technologies in Table 7 – 13 resulted in at least a 70% reduction of the total threat index. The remaining security technologies, i.e. the technologies not listed, resulted in less than a 16% reduction in the total threat index.

Table 7 - 13 Security Technology Threat Index Changes

7.4.2 Analysis of Results

SAEM analysis showed that Hardening the Operating System and Secure Operating Systems are the two most effective countermeasures among all the security technologies. Not surprisingly, SAEM gave high rankings to Email Filters, URL Blockers, and WEB Access Control Products because they provide the greatest protection against the internal threats that SAEM ranked highly, such as *Personal Computer Abuse* and *Procedural Violations*. The Deputy Director would not have given these security technologies high rankings since he did not consider these internal threats to be as significant as other threats. Although the Deputy Director ranked Network-based IDS and Log Analysis Software in the top ten, SAEM ranked them only near the top ten, 12th and 14th respectively.

The Deputy Director reviewed the results of the Benefit Analysis phase and did not revise his initial list of the top most effective security technologies. Although he ranked Hardware Lockdown and DB Secure Access Controls as some of the most effective security technologies for the organization, SAEM ranked Hardware Lockdown 38th and DB Secure Access Controls 20th. The Deputy Director identified Hardware Lockdown as effective against twelve threats, but they were threats that he had highly ranked, but SAEM had ranked low. In addition, the Deputy Director identified DB Secure Access Controls as effective against several threats, but only *Browsing* had a high relative threat index, 12. Again, DB Secure Access Controls are effective against threats that the Deputy Director ranked high, such as *Theft* and *Contamination*.

| Rank | Security Technology | Threat Index % |
|------|-----------------------------------|----------------|
| 1 | Hardened OS | 100% |
| 2 | Secure OS | 100% |
| 3 | E-Mail Filters | 99% |
| 4 | URL Blockers | 89% |
| 5 | WEB Access Control Products | 88% |
| 6 | Proxy Firewalls | 88% |
| 7 | Penetration Testing Tools | 88% |
| 8 | Packet Filter Firewalls | 88% |
| 9 | Vulnerability Assessment Scanners | 88% |
| 10 | Forensic Software | 76% |
| 11 | Host-Based IDS | 73% |
| 12 | Network Based IDS | 72% |
| 13 | Centralized Security Management | 72% |
| 14 | Log Analysis Software | 71% |
| 15 | Auditing Tools | 70% |

7.4.3 Analysis of Director’s Changes

The Director of Mission Assurance’s changes in the risk assessment resulted in a different ranking of security technologies (See Table 7-14). Most notably, Virus Protection Software was ranked 5th in contrast to the initial SAEM ranking of 18th. A higher ranking of the technology is more consistent with the risk ranking. In addition, Database Security Access Controls were ranked higher (14th), but Log Analysis Software was rated lower (19th). Both the Deputy Director and SAEM ranked Antivirus Products in the top then most effective security technologies.

**Table 7 - 14 Security Technology Threat Index Changes
using Director's Risk Assessment**

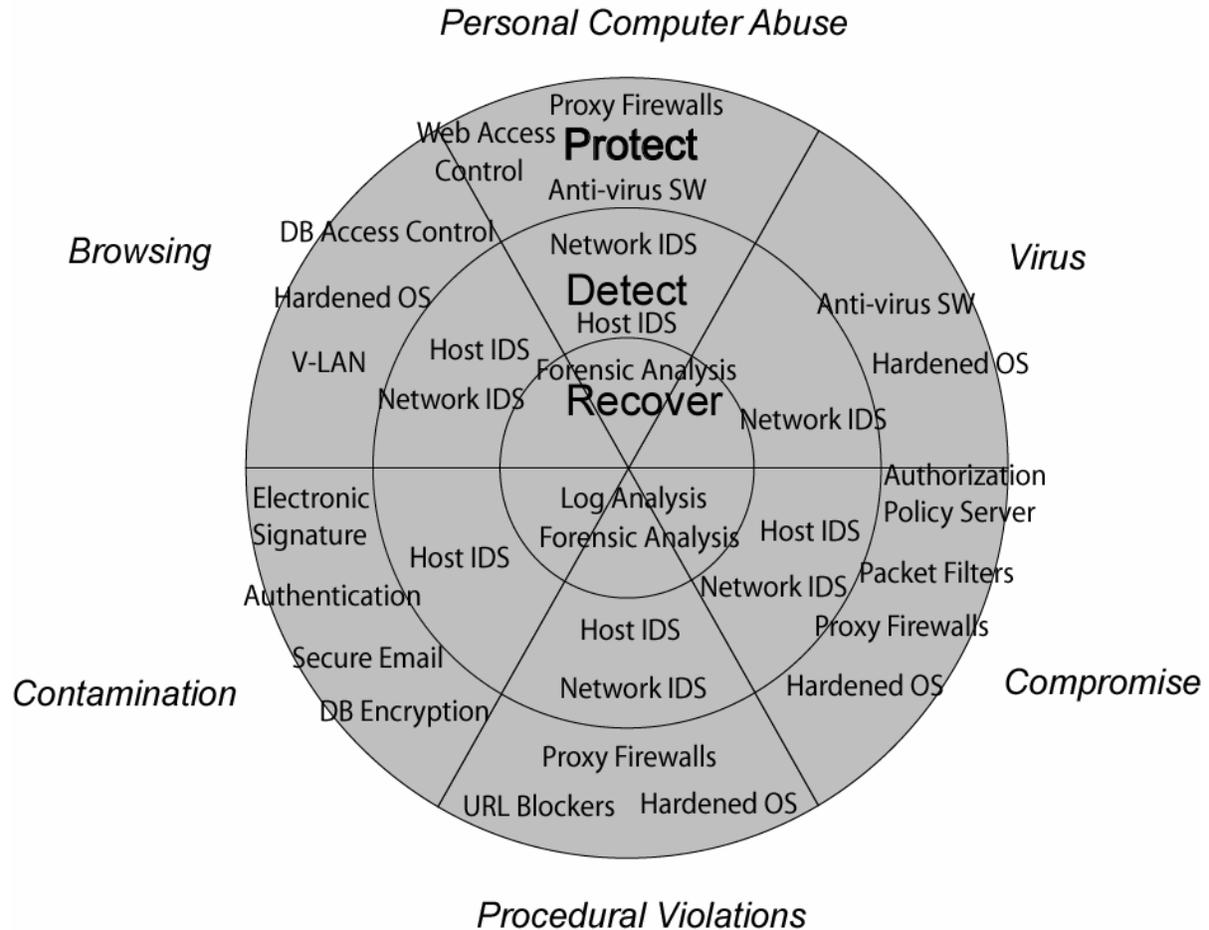
| Security Technology | Relative Threat Index % | Director's Rank | Deputy's Rank | Relative Threat Index % |
|-----------------------------------|-------------------------|-----------------|---------------|-------------------------|
| Hardened OS | 99% | 1 | 1 | 100% |
| Secure OS | 99% | 2 | 2 | 100% |
| E-Mail Filters | 99% | 3 | 3 | 99% |
| URL Blockers | 99% | 4 | 4 | 89% |
| Antivirus Products | 81% | 5 | >15 | |
| WEB Access Control Products | 73 % | 6 | 5 | 88% |
| Penetration Testing Tools | 73% | 7 | 7 | 88% |
| Proxy Firewalls | 73% | 8 | 6 | 88% |
| Packet Filter Firewalls | 73% | 9 | 8 | 88% |
| Vulnerability Assessment Scanners | 73% | 10 | 9 | 88% |
| Network Based IDS | 62% | 11 | 12 | 72% |
| Host-Based IDS | 58% | 12 | 11 | 73% |
| Software Lockout | 55% | 13 | >15 | |
| Database Security Access Controls | 55% | 14 | >15 | |
| Authorization Policy Servers | 55% | 15 | >15 | |

7.5 Coverage Analysis

In order to complete Coverage Analysis, the Deputy Director identified technologies currently employed in the organization's security architecture. The Deputy Director identified three technologies which the organization had *not* integrated into the security architecture: 1) Key Stroke Replicators, Mobile Code Scanners, and 3) Biometrics. Figure 7 -2 shows how some of the organization's security technologies mitigate the risk from the top six threats, as computed by SAEM. The security technologies depicted in the defense-in-depth model are those that the Deputy Director estimated as the most effective against a threat in each category of Protect, Detect, and Recover. Notice that the segments in Figure 7-2 are uniformly grayed, since the security technologies depicted only represent a small sample of the technologies employed.

In this case study, the coverage analysis process did not reveal significant gaps or weaknesses in the security architecture. However, the participants agreed that the process validated actions they had taken to improve security, while providing a tool that could be helpful in persuading budget decision-makers to fund additional security technologies. The Director's comments in the Satisfaction Survey (at the end of this chapter) indicated she would find it very useful to justify purchase of a security technology.

Figure 7 - 2 Defense-in-Depth Coverage of Selected Security Technologies



7.6 Security Technology Tradeoff Analysis

In the final phase of SAEM, the Deputy Director selected and ranked three technologies for comparison and identified several objectives, which he used to assess whether to recommend a security technology for inclusion in the security architecture. This section describes the results of using multi-attribute analysis techniques to help the Deputy decide which security technologies best meet his objectives.

First, the Deputy Director selected and ranked three security technologies: Mobile Code Scanners, Personal Firewalls, and Host-based Intrusion Detection Systems (in order). Next, he identified four objectives that he used to select security technologies:

1. Threat
2. Effectiveness of the Technology
3. Coverage
4. Maintenance.

Threat refers to the degree of risk that the threat poses to the organization. Effectiveness refers to the difficulty of defeating the technology’s security function. Ideally, the Deputy Director wants to find technologies that cover a breadth of threats and plug as many security gaps as possible. Maintenance refers to the cost of implementation and continued annual maintenance of the technology. The organization is large, so the Deputy Director would like to find risk-mitigating technologies that can be implemented and maintained from a central location rather than requiring several thousand configurations. Table 7 - 15 shows the security technologies, the Deputy’s ranks and weights with respect to each of the objectives.

Table 7 - 15 Initial Security Technology Tradeoff Assessment

| | | Tradeoff Attributes | | | | | | | | Tradeoff Ranking |
|---------------------|---------------------|---------------------|--------|---------------|--------|----------|--------|-------------|--------|------------------|
| | | Threat | | Effectiveness | | Coverage | | Maintenance | | |
| | | Rank | Weight | Rank | Weight | Rank | Weight | Rank | Weight | |
| Security Technology | Mobile Code Scanner | 1 | 60 | 1 | 60 | 1 | 70 | 1 | 80 | 1 |
| | Host-based IDS | 3 | 5 | 3 | 5 | 3 | 10 | 2 | 10 | 3 |
| | Personal Firewall | 2 | 35 | 2 | 35 | 2 | 20 | 3 | 10 | 2 |

Since Mobile Code Scanner clearly dominates (i.e. it is rated highest in all objectives) in each tradeoff attribute, Mobile Code Scanners is ranked first. Personal Firewalls dominates in three tradeoff attributes and is the same in the fourth attribute, so it dominates Personal Firewall and is immediately after Mobile Code Scanners. Therefore, Personal Firewalls are ranked third. The ranking of Table 7-16 is consistent with the Deputy Director’s initial ranking.

When the Director of Mission Assurance reviewed the results, she believed that one additional criterion should be used to evaluate security technologies. She wanted to add “Customer Support,” which was an indication of how well the security technology met the mission assurance needs of the customer. Customers within this government organization are other departments. She gave each security technology a Customer Support weight as follows:

- Mobile Code Scanner – 80
- Personal Firewall – 10
- Host-base IDS – 10

Since Mobile Code Scanner has the highest ranking within the Customer Support Attribute, it still dominates the other security technologies. In addition, since Personal Firewall and Host-base IDS are ranked equally, Personal Firewall is still rated second overall and Host-based IDS rated third. Therefore, the addition of another attribute did not change the overall rankings of the security technologies.

7.7 Summary

Despite the fact that SAEM did not appear to significantly influence the Deputy Director's prioritization of threat, the participants found the results insightful. The Deputy Director's prioritization of threats reflected his present concerns for the organization and SAEM may have been more useful if other assumptions were used during the risk assessment process. In addition the benefit analysis resulted in highly ranking security technologies that mitigated threats that the risk assessment had determined were important to the organization. The organization had purchased all of the security technologies for those threats so the benefit analysis validated their previous selections. Finally, although the Deputy Director did not comment on the coverage analysis, the Director of Mission Assurance felt that the coverage analysis was a very important for communicating the organization's needs.

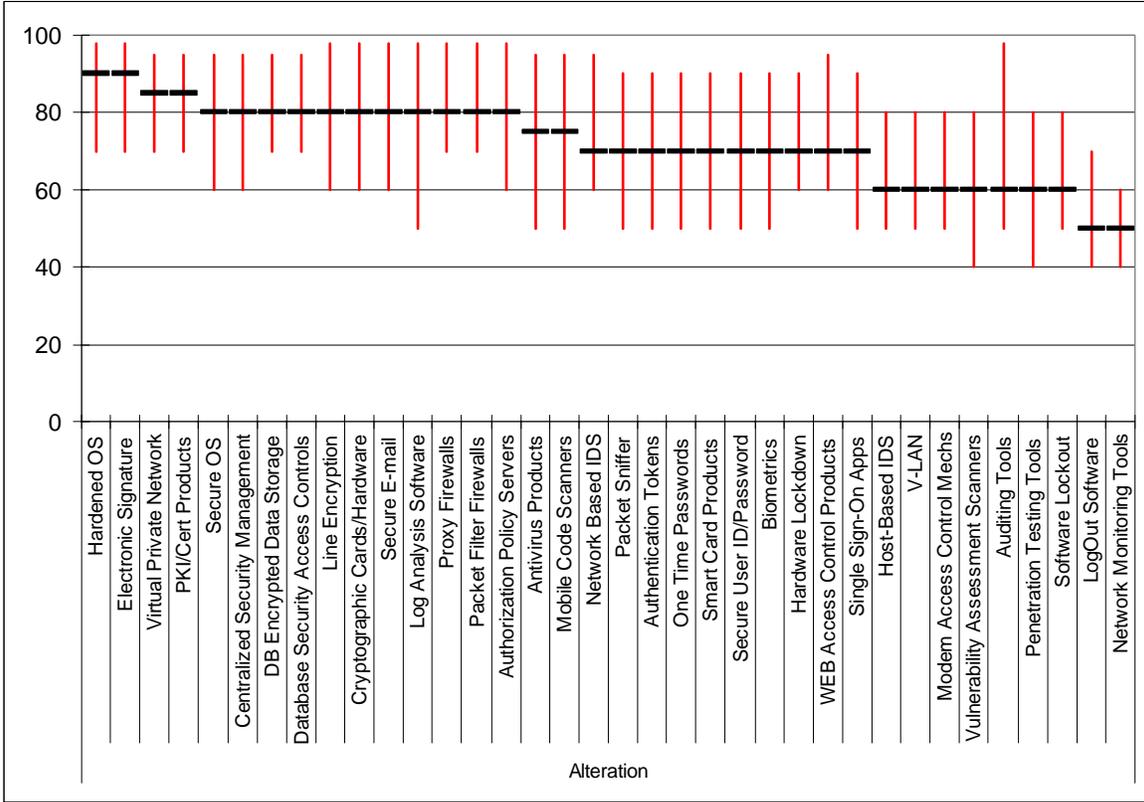
Overall, this case study provided two key insights to SAEM. First, this case study highlighted the *point of reference* issue when asking for frequency and outcome estimates. If the security manager wants to know whether a security architecture technology is the most effective, then he or she might ask how many attacks would be occurring and what would be the outcomes if that security technology were not in place. In contrast, the security manager might want to know which security technology should be added to the architecture, in which case he or she might estimate the frequency and outcomes of security incidents given the current security architecture.

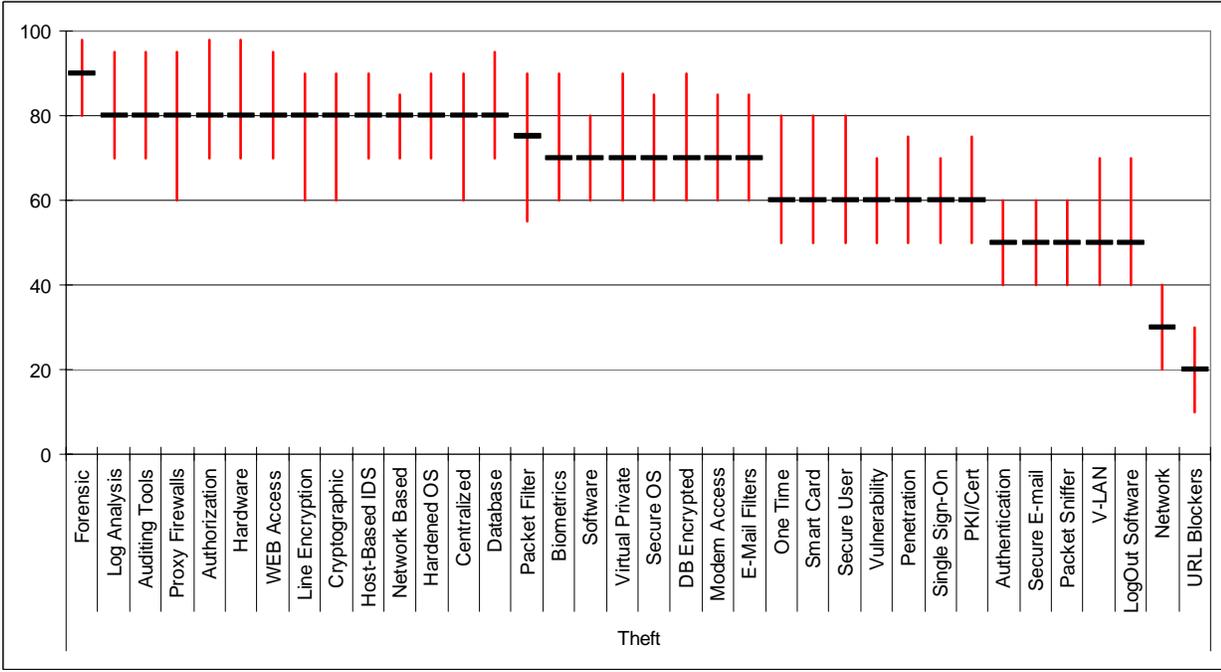
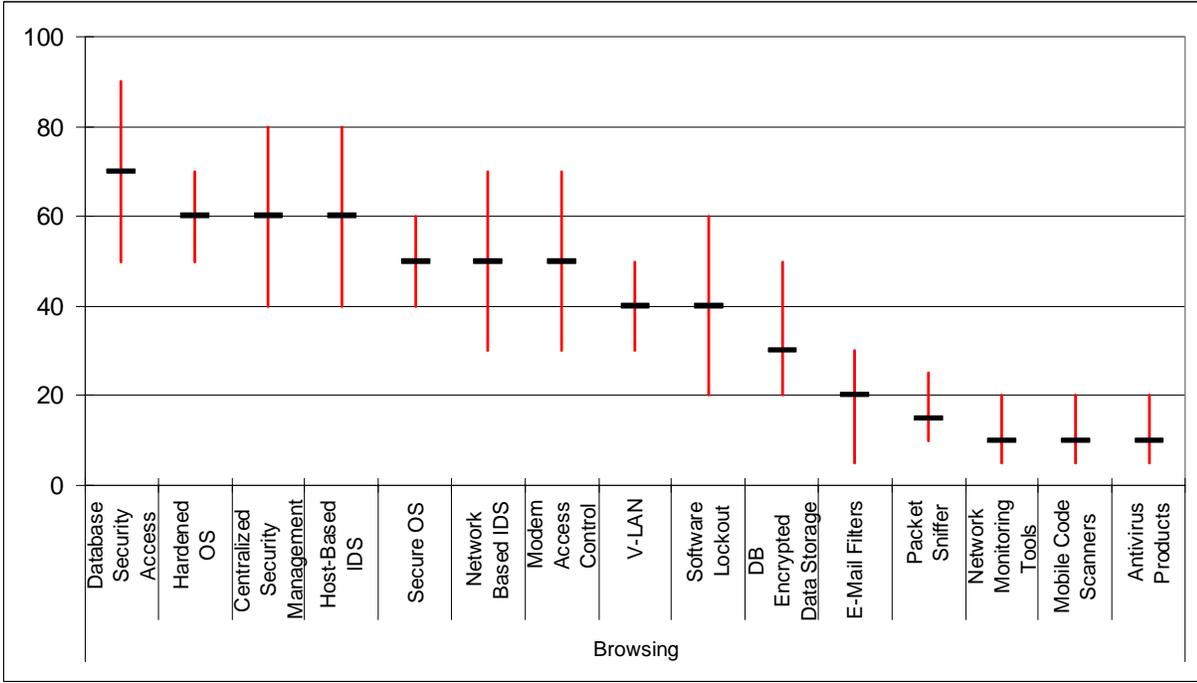
The point-of-reference perspective determines whether the analyst elicits information about all attack attempts or only elicits information about successful security compromises. The frequencies and outcome estimates would be very different between the two points of reference, i.e., estimating attack attempts would result in much higher frequency and outcome estimates for threats that have been significantly reduced by the organization's security architecture. In this thesis, the point of reference for each of the case studies was to evaluate the security technologies against all security compromise attempts.

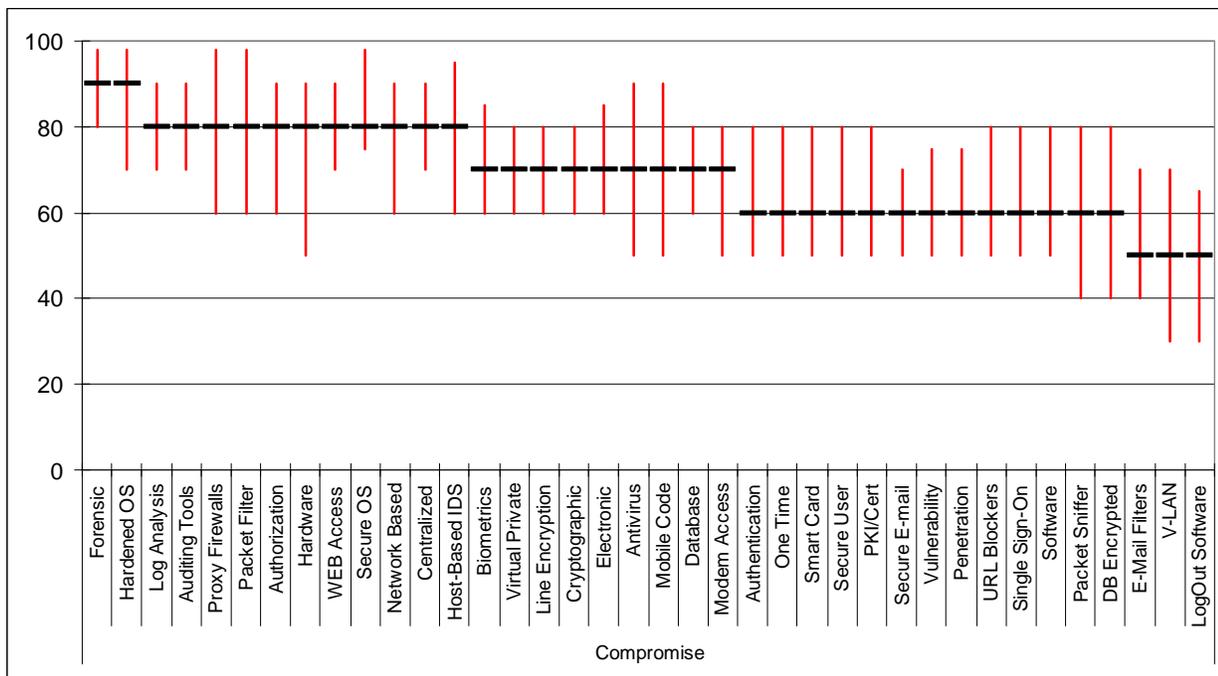
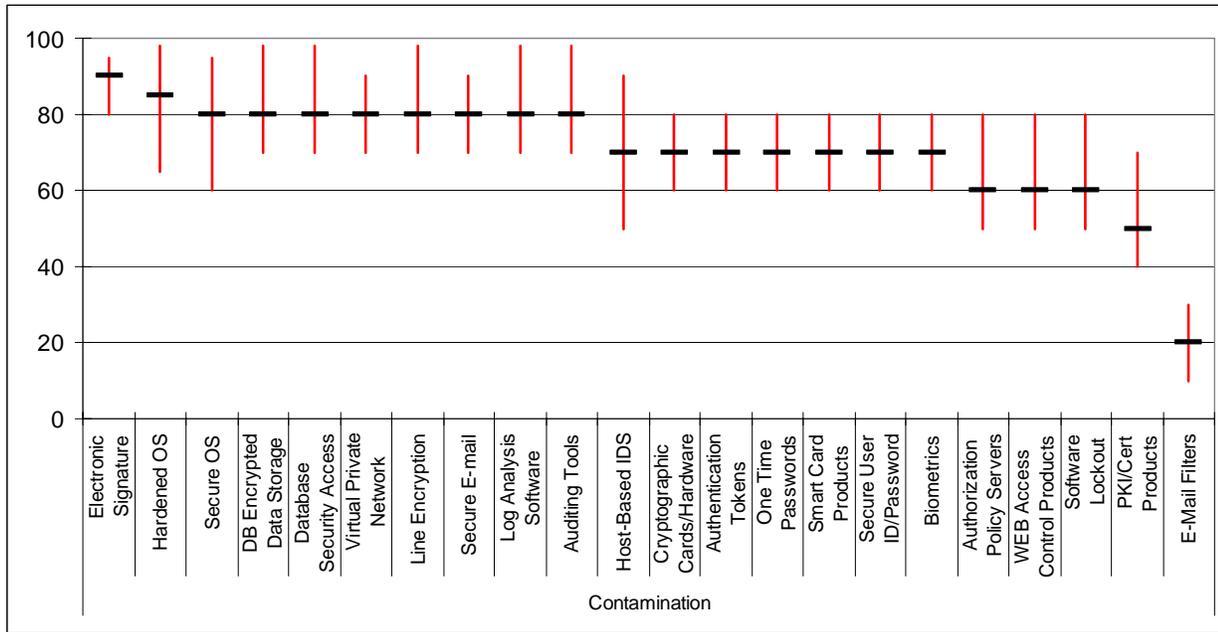
Secondly, the participants in this case study were very experienced and technically proficient. Not surprisingly, the processes did not uncover any security weaknesses. Organizations that do not have such experienced security professionals may find SAEM is more useful in guiding decisions. However, the main contribution of SAEM in this case study was in the ability of the Director to communicate with non-security executives the state of their security environment. In her Satisfaction Survey, the Director strongly indicated that she needed to automate SAEM, a clear signal that she valued the method.

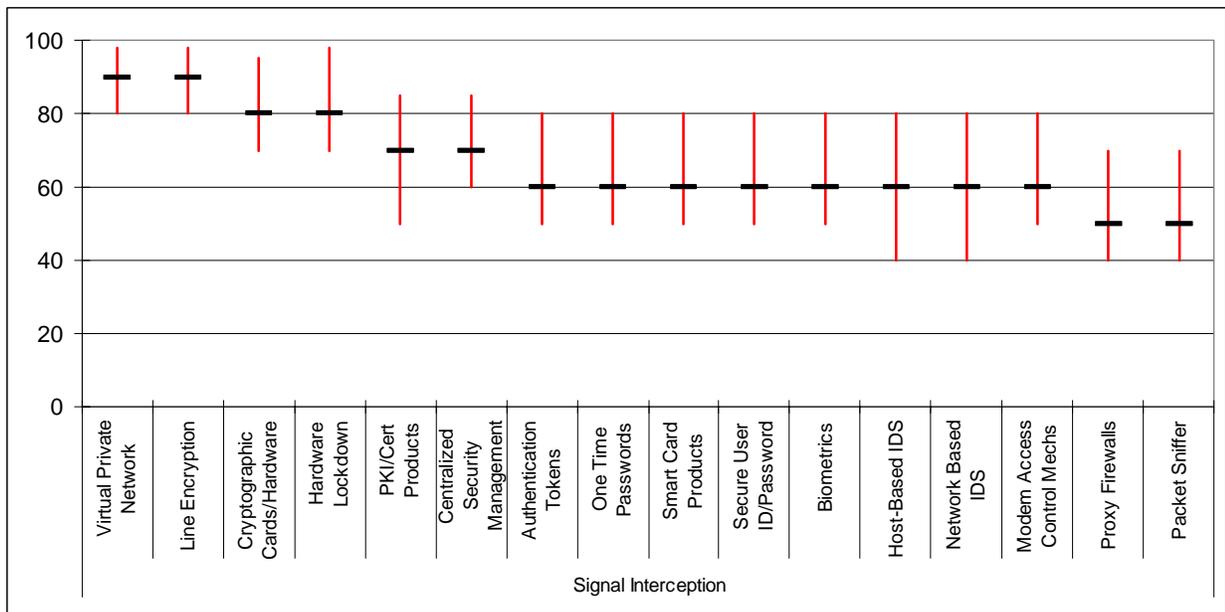
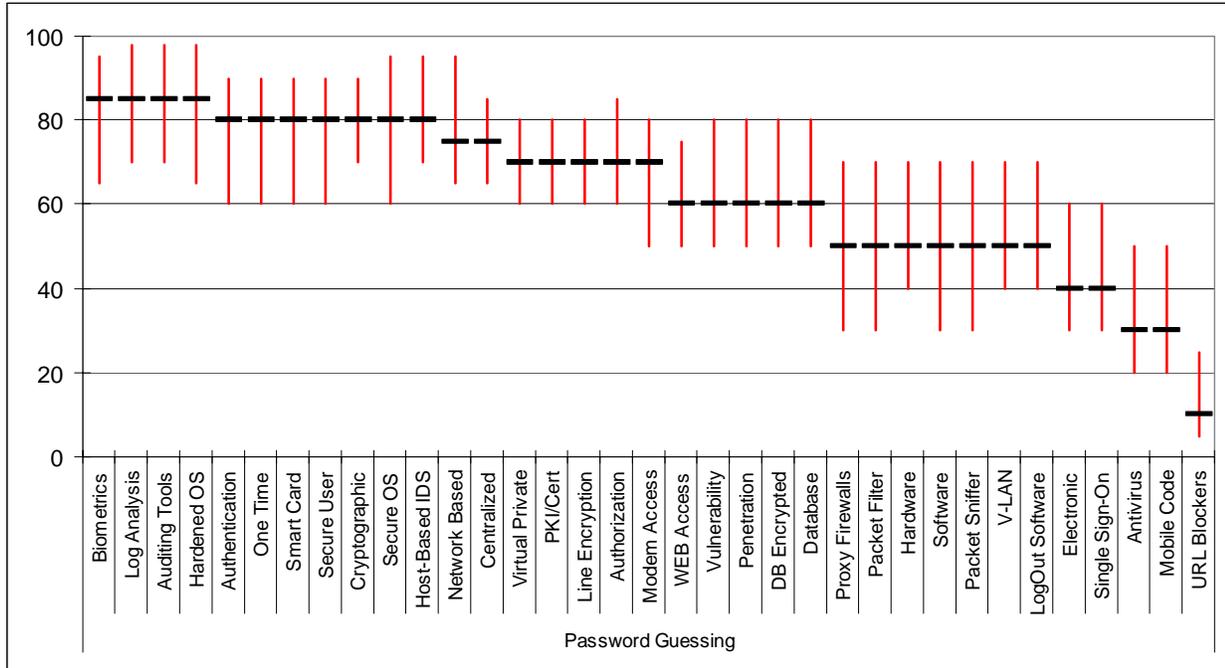
BENEFIT ANALYSIS DATA (Government Case Study)

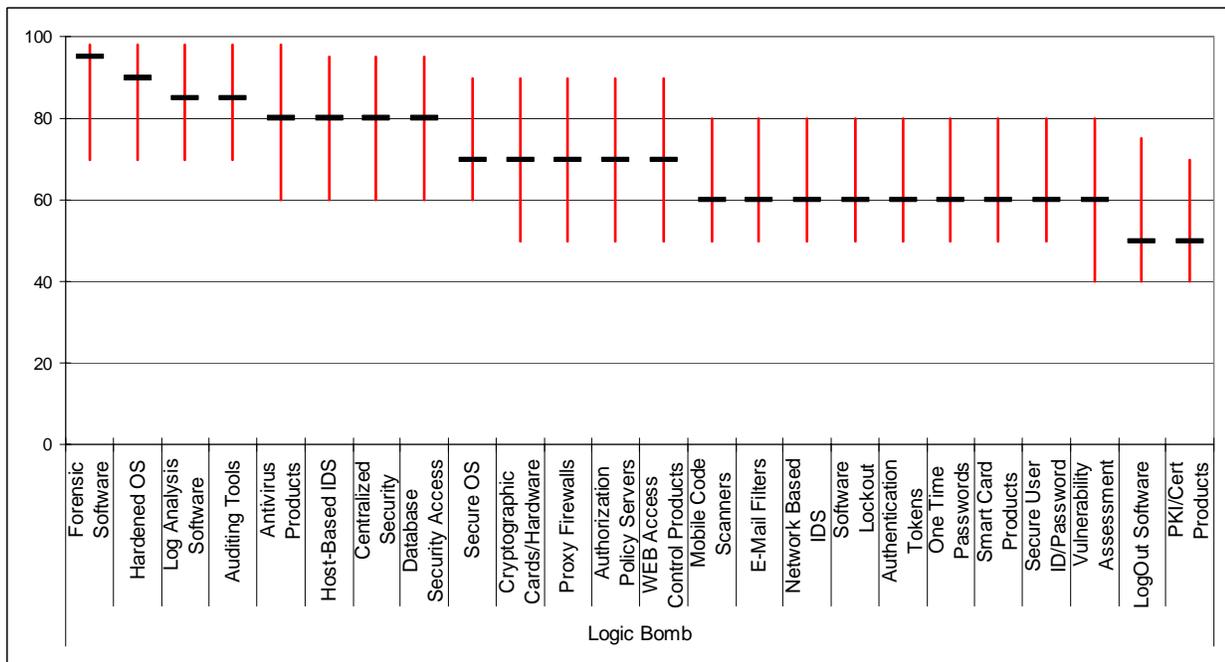
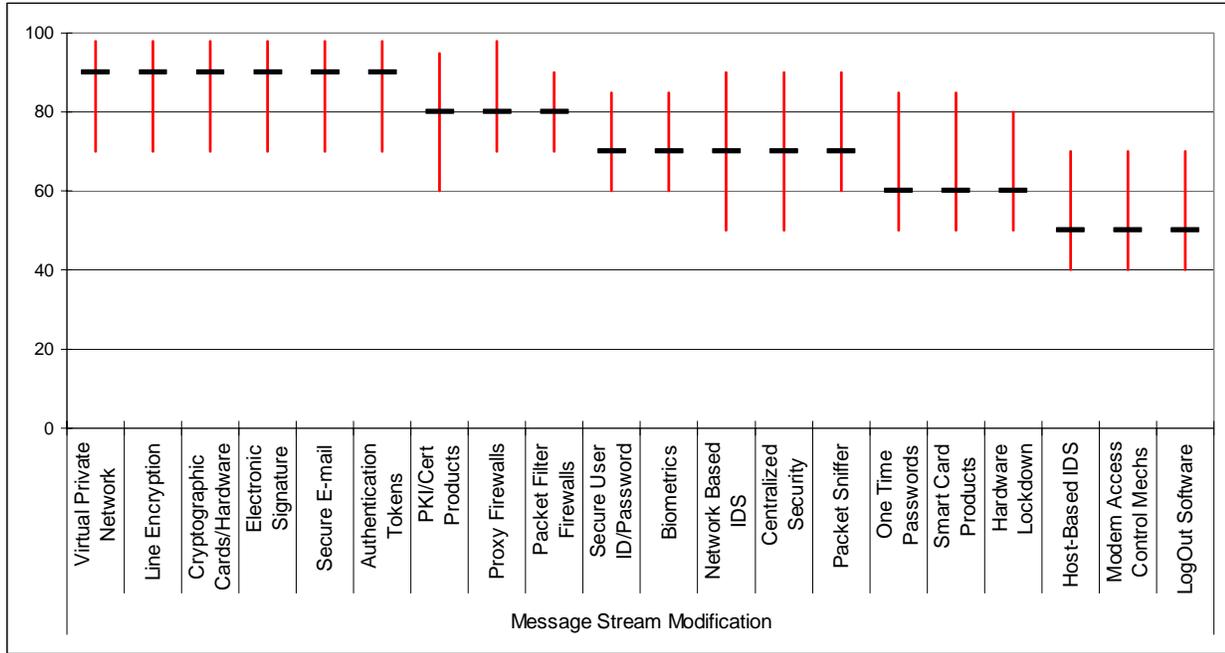
(Note: Y-Axis is Percentage Effectiveness on all graphs)

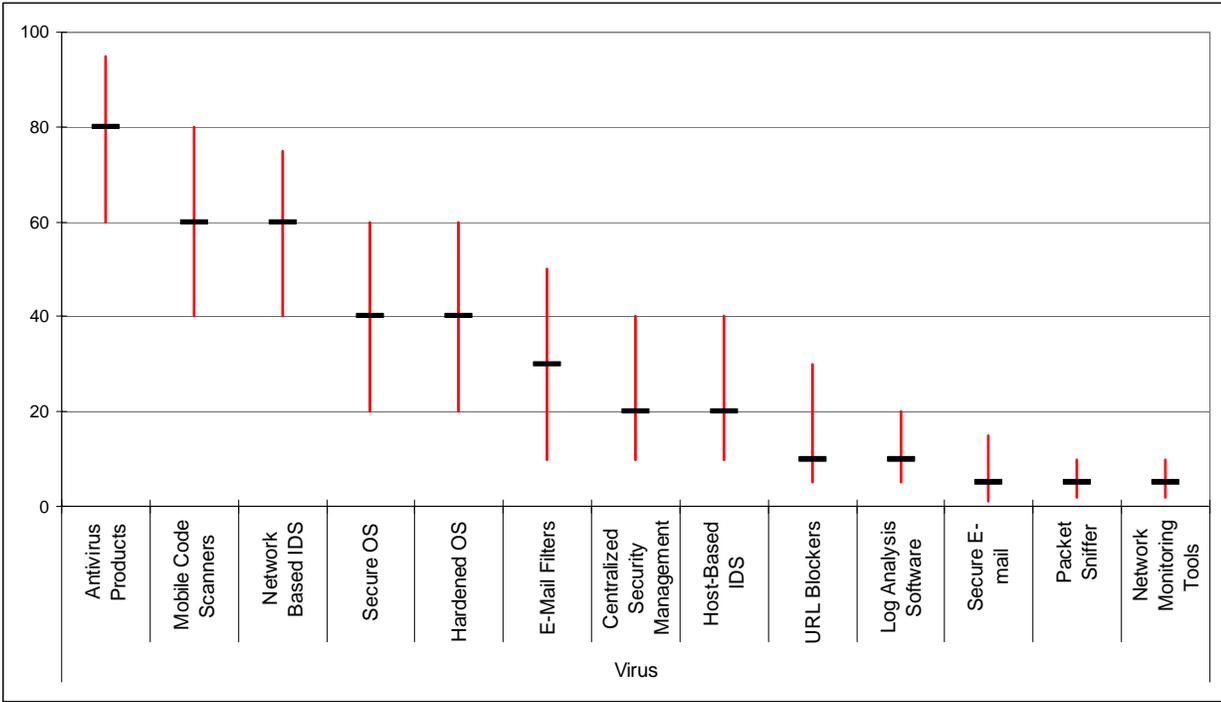
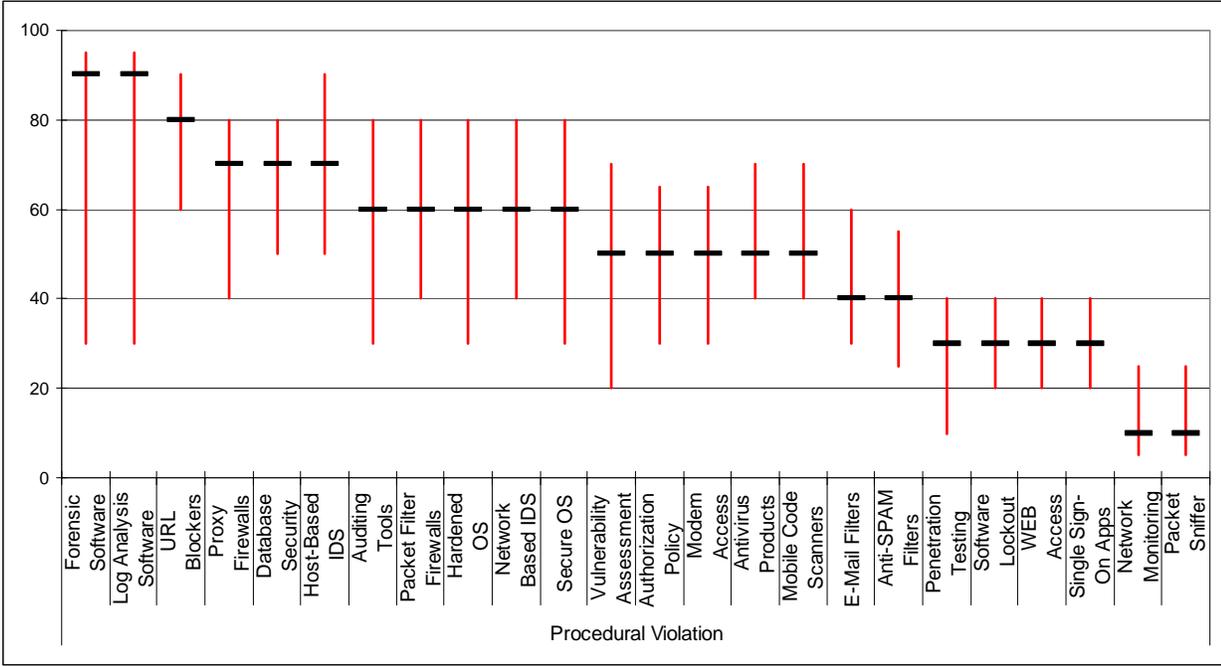


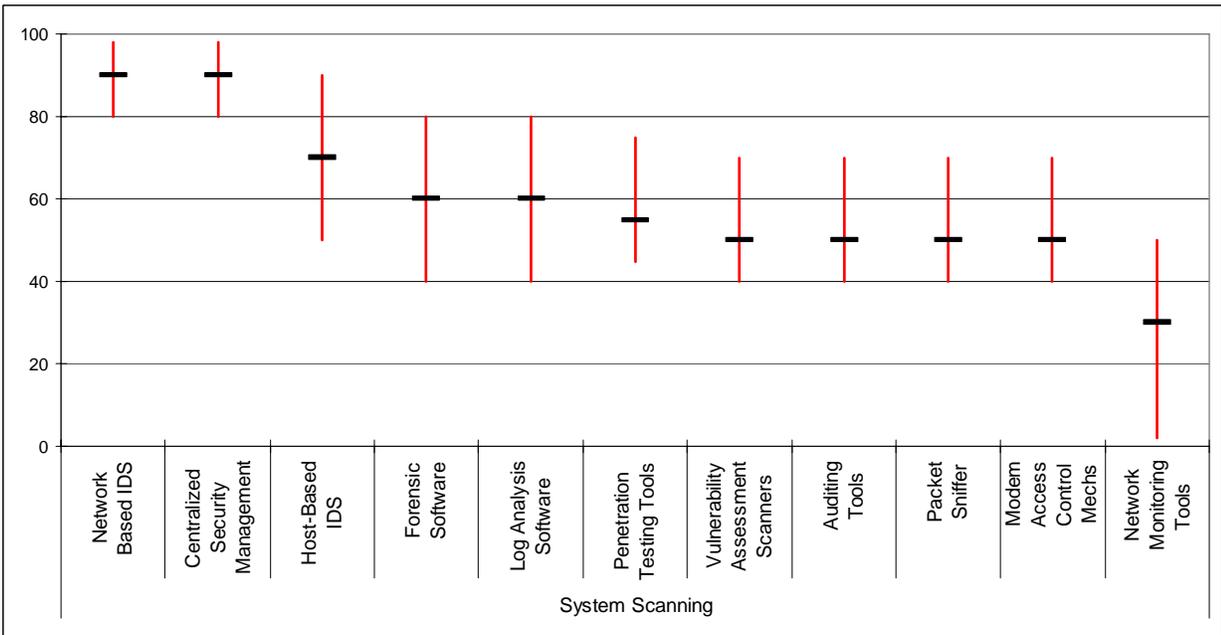
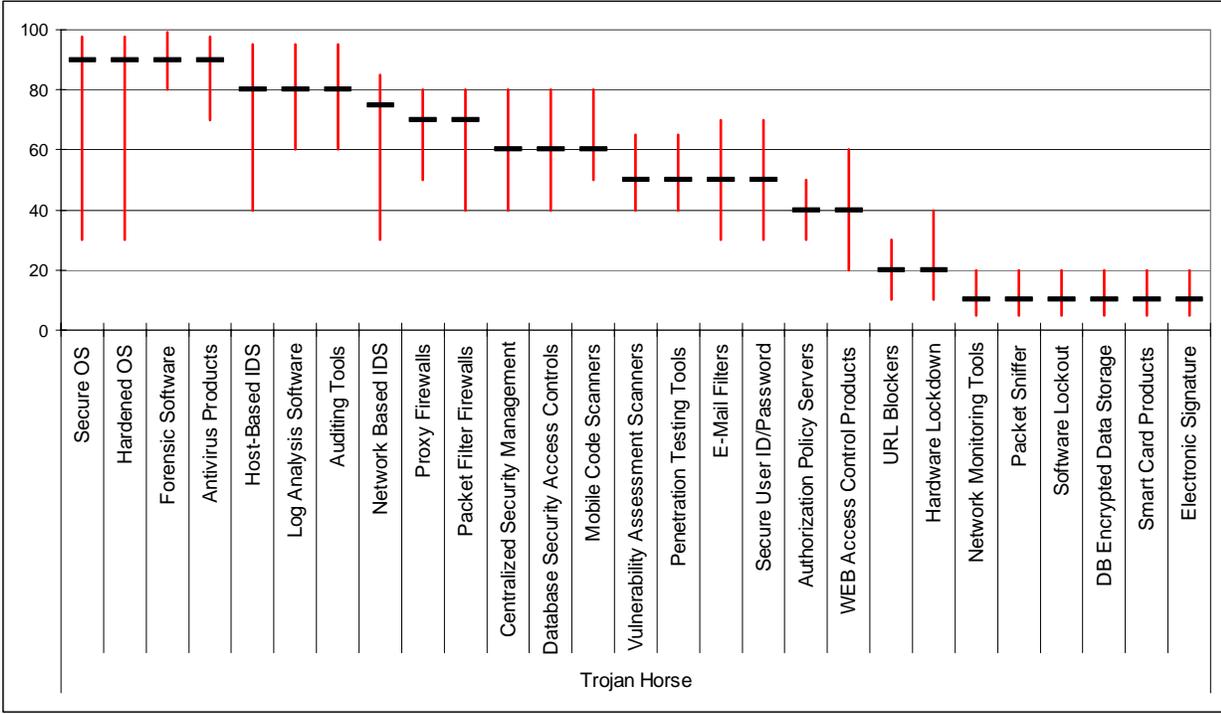


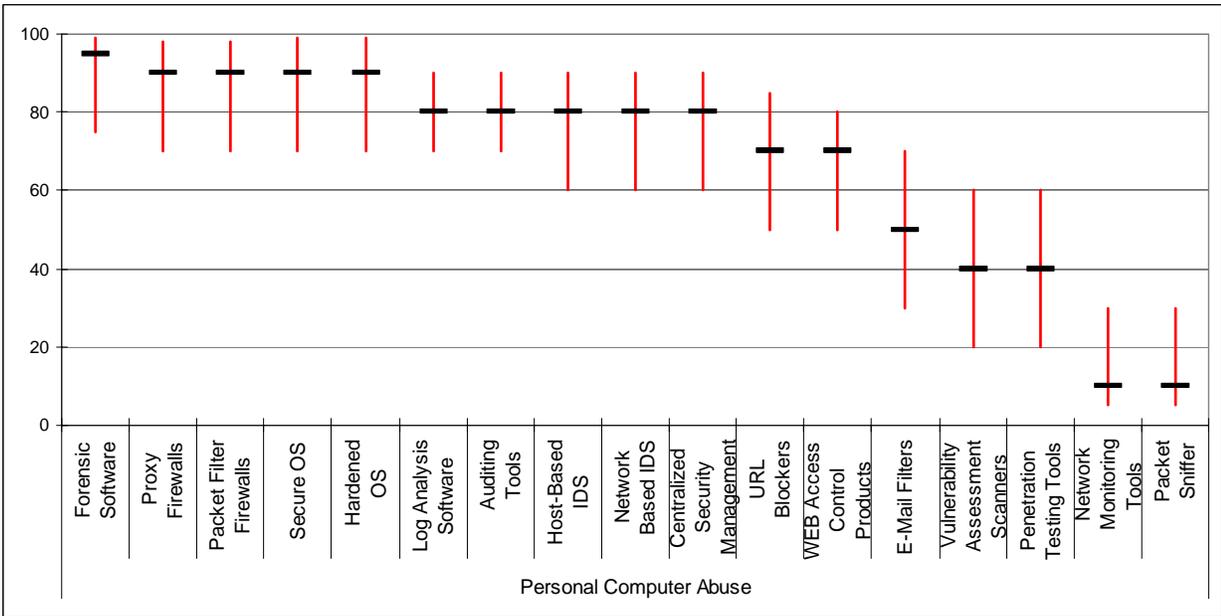
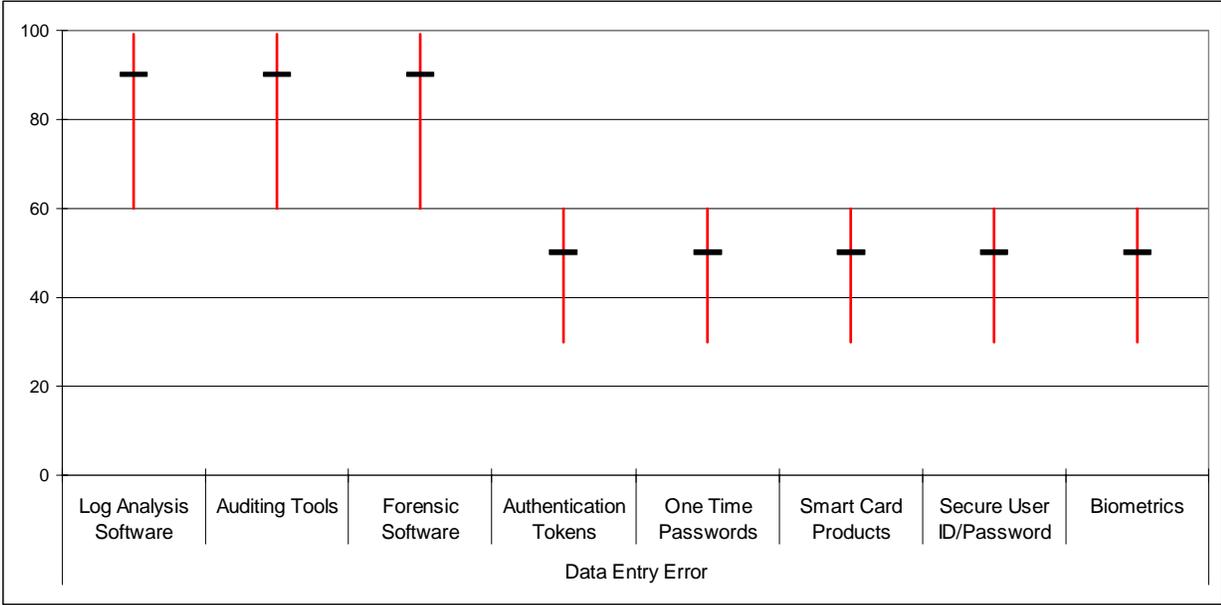


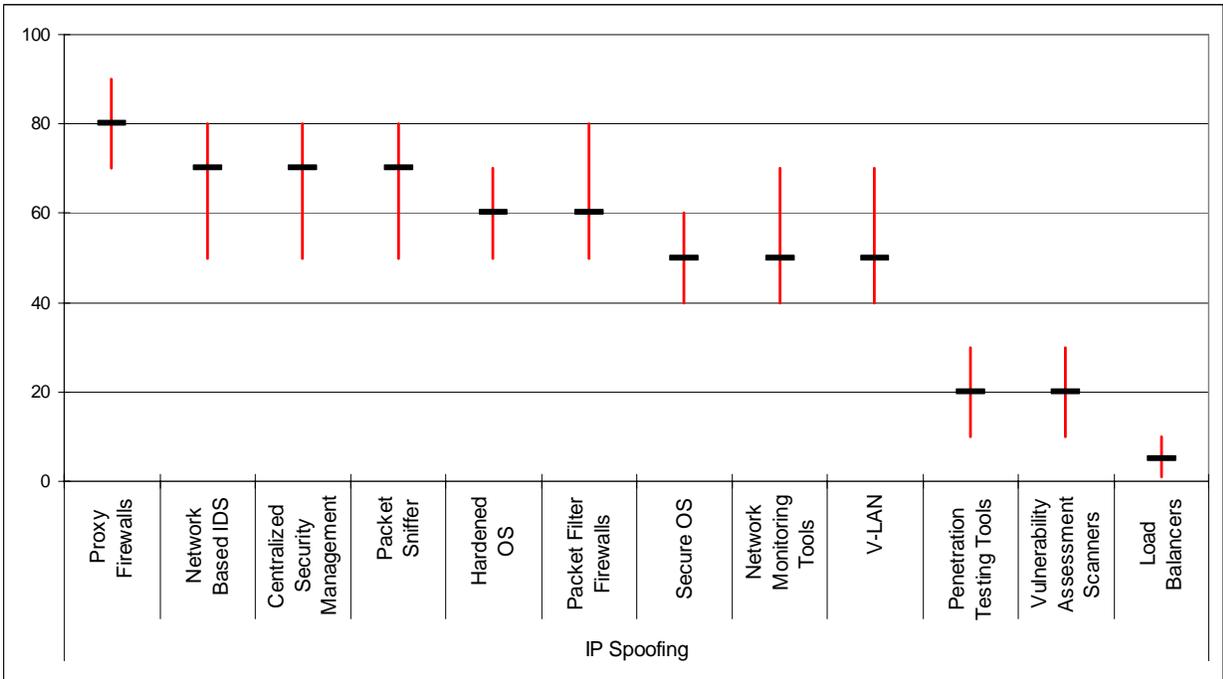
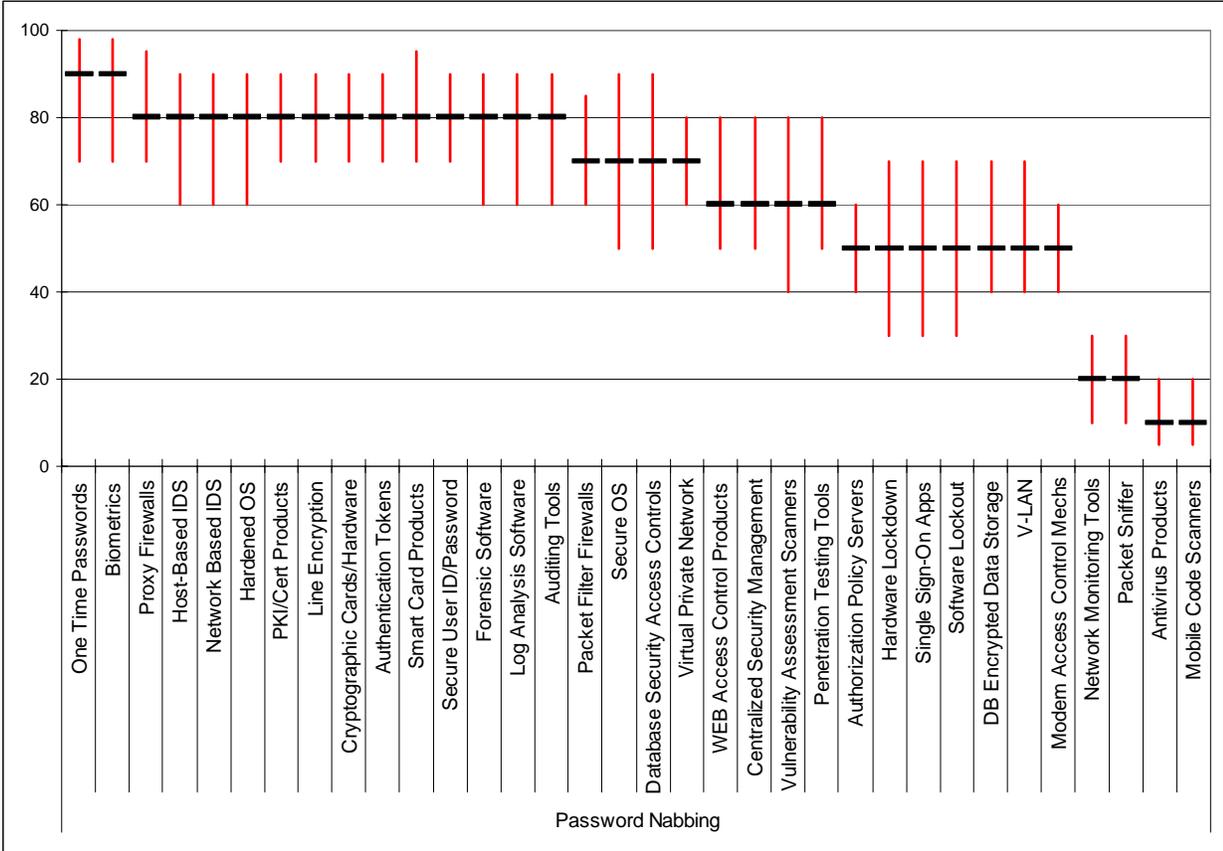


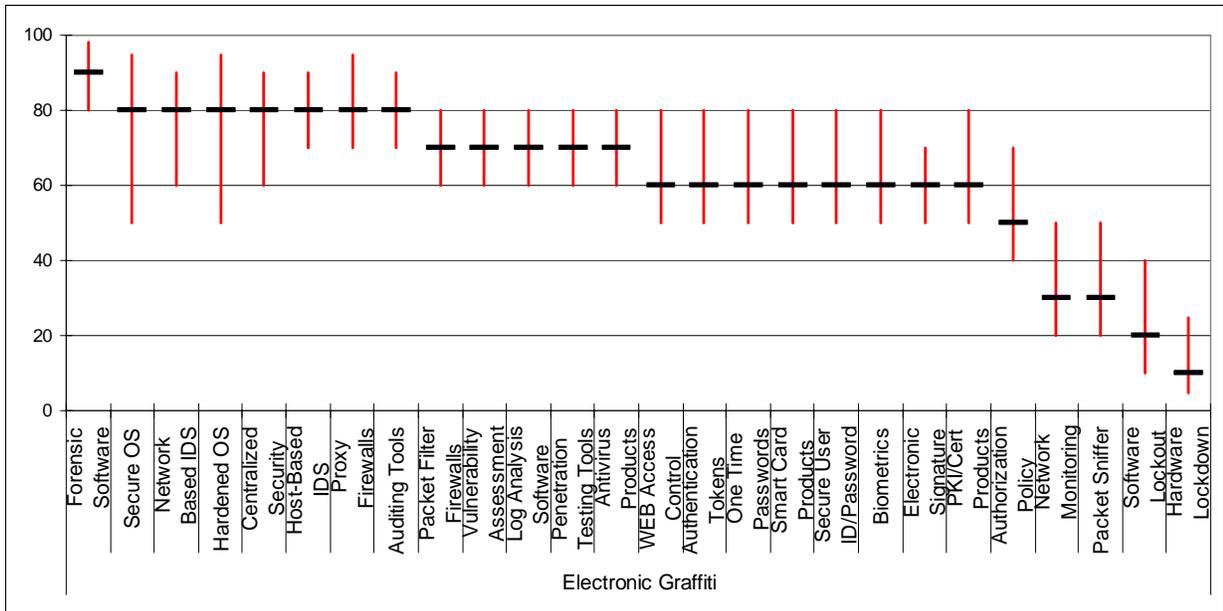
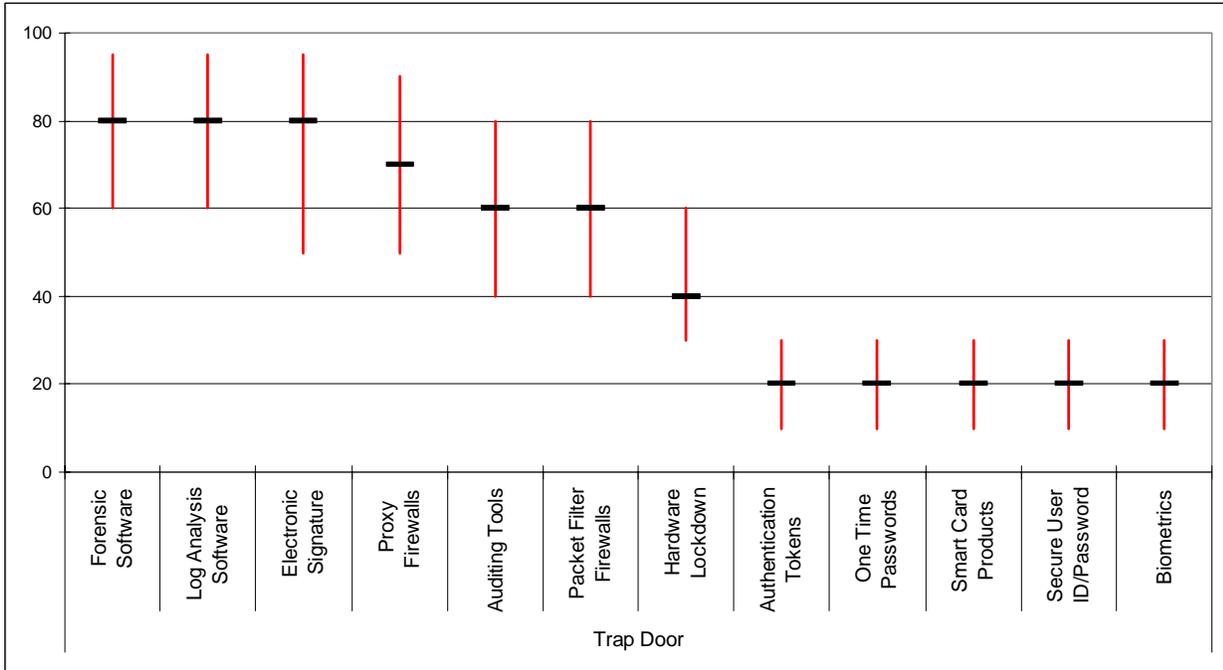


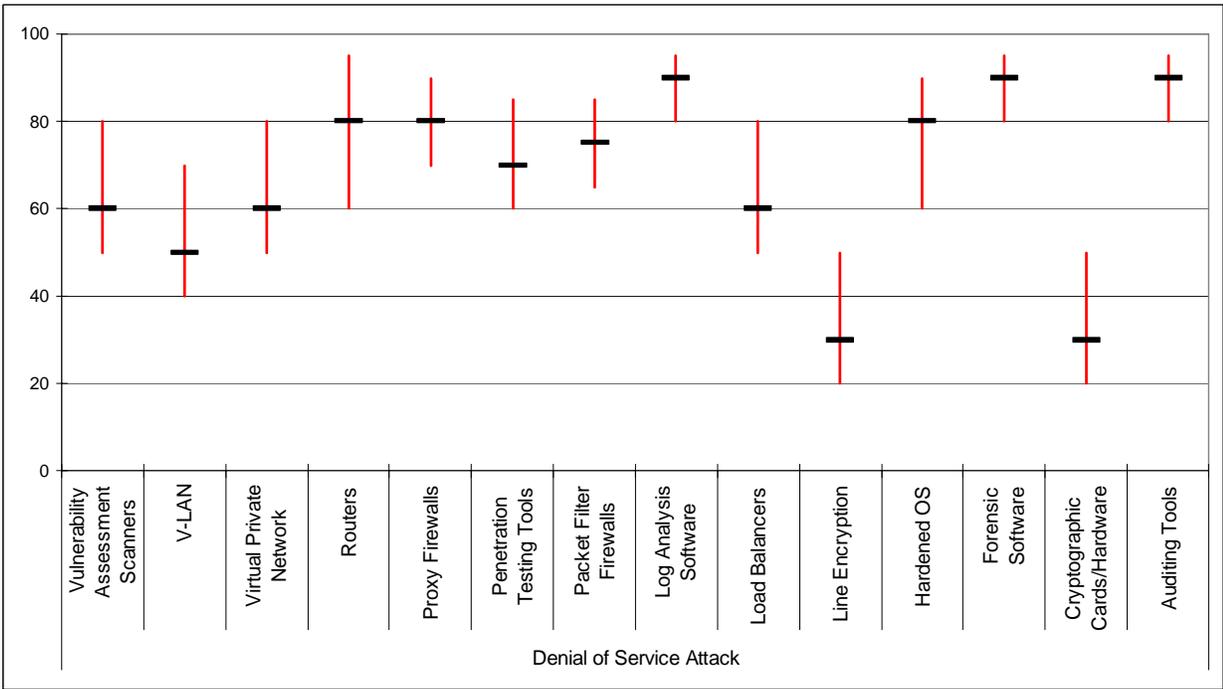
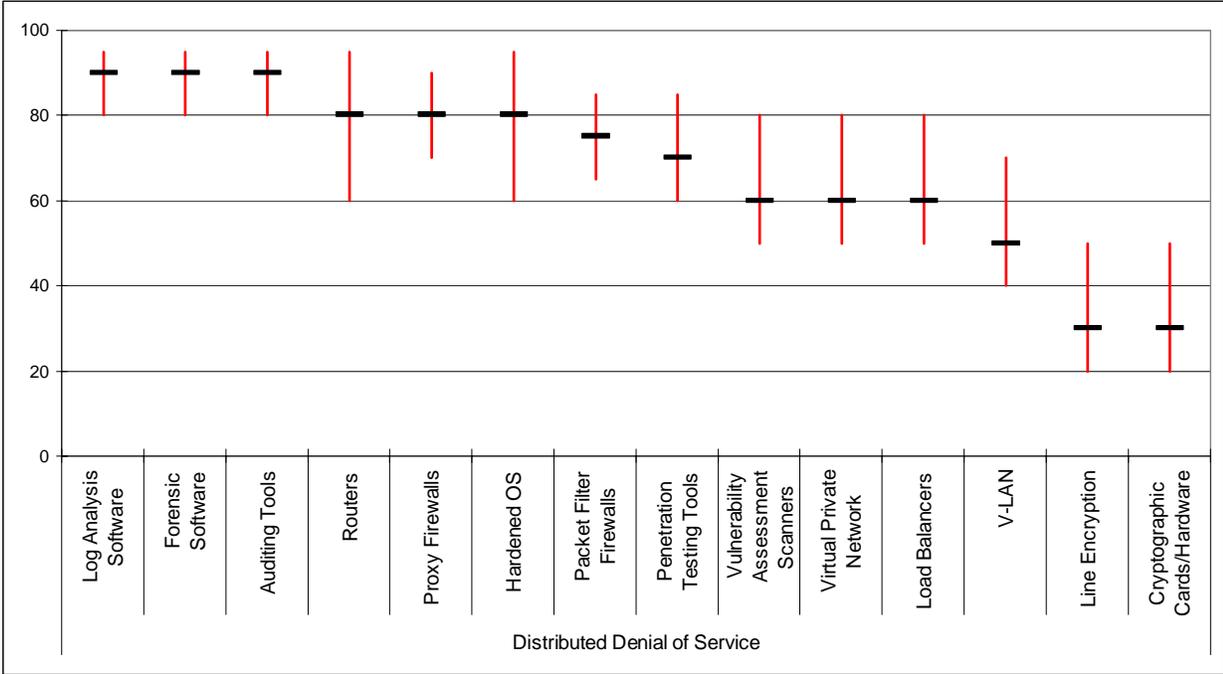


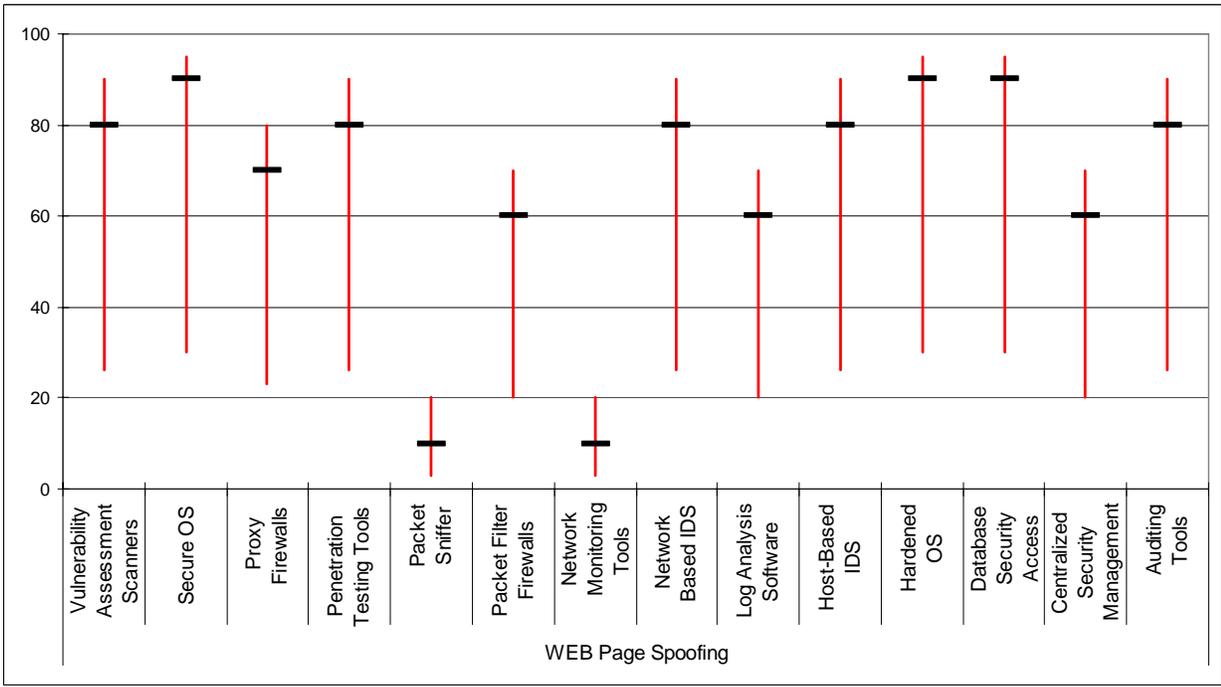
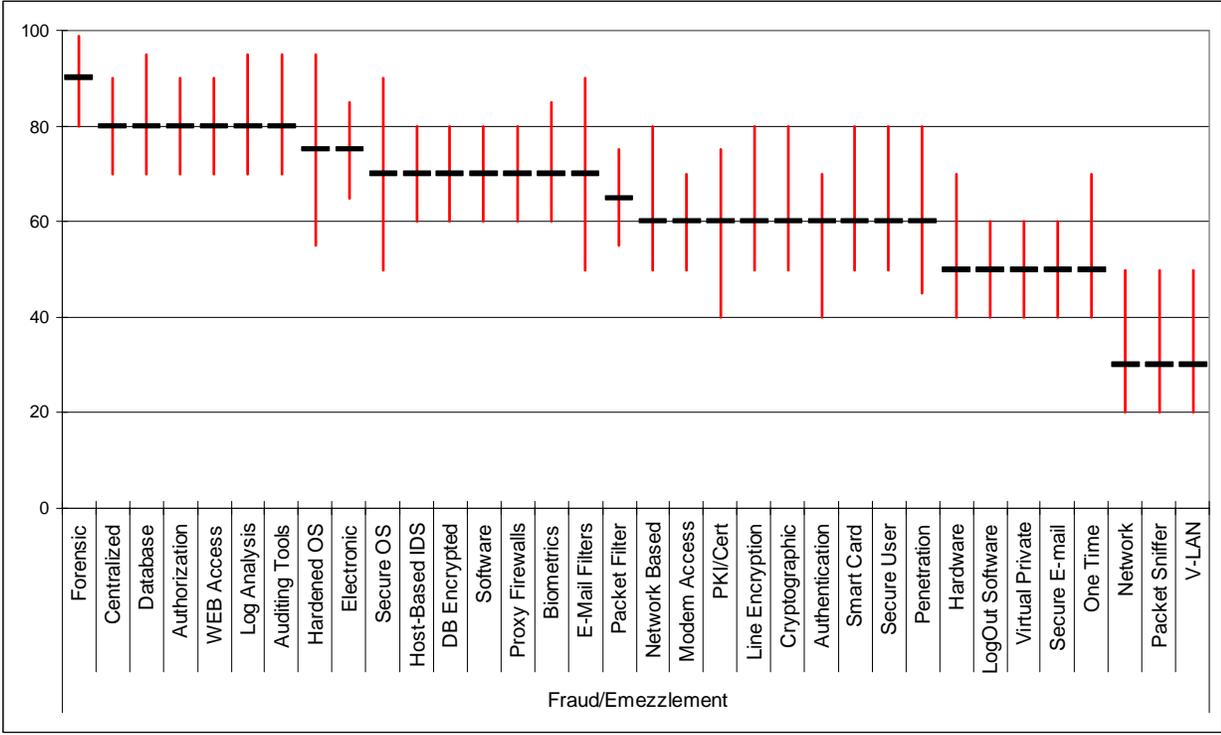


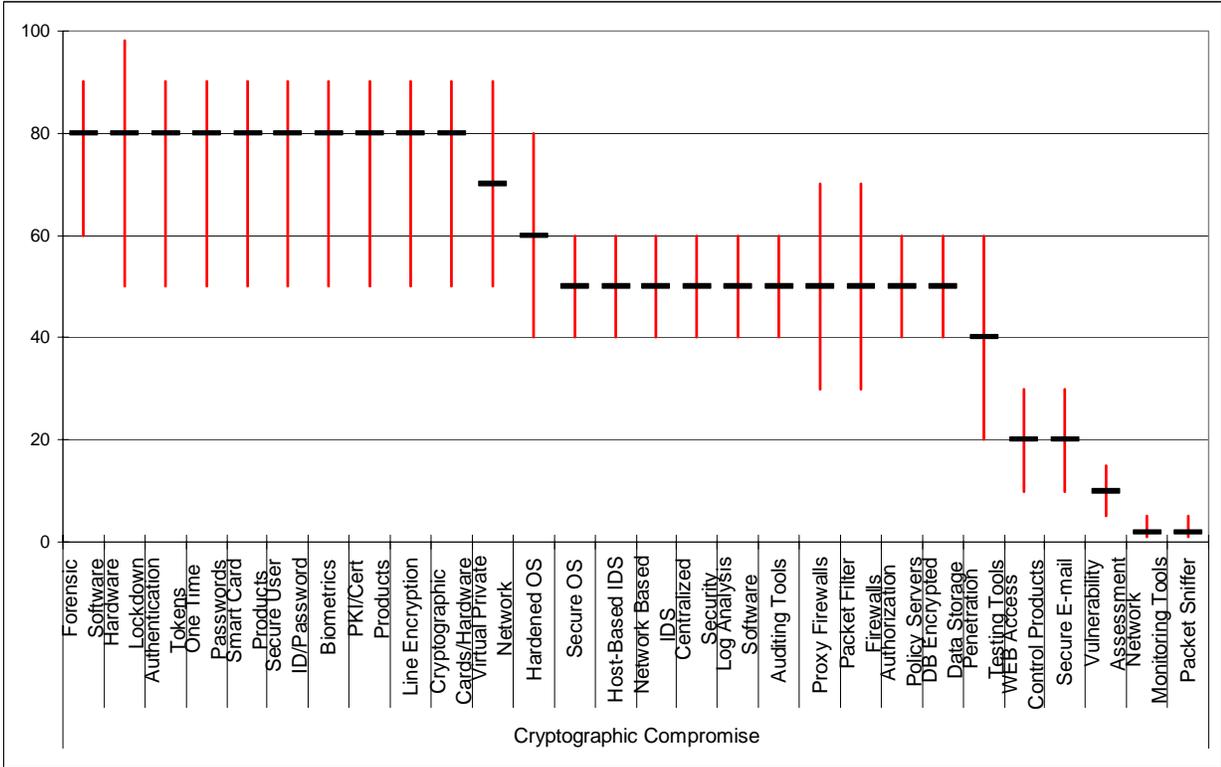












SATISFACTION SURVEY
(Government Case Study)

The Director and Deputy Director completed this survey at the end of the SAEM process. The comments at the end of each section are their unedited comments; however, the Deputy Director did not rate the risk assessment process.

I. Risk Assessment: During the Risk Assessment, participants identified the organization’s threats and estimated the frequency and outcome of attacks. The Risk Assessment resulted in a prioritization of the organization’s threats.

(G*: Director of Mission Assurance, G: Deputy Director)

| | | | | | | | |
|---|-----------------------------|----|-----------------|-----------|-----------|-------------------------|---|
| How difficult was it ... | <i>Not at all difficult</i> | | <i>Somewhat</i> | | | <i>Very difficult</i> | |
| ... to identify and initially rank the threats? | 0 | 1 | 2 | G* | 4 | 5 | 6 |
| ... to estimate the frequency of attacks? | 0 | 1 | G* | 3 | 4 | 5 | 6 |
| ... to estimate the outcomes? | 0 | 1 | G* | 3 | 4 | 5 | 6 |
| <hr/> | | | | | | | |
| | <i>None at all</i> | | <i>Somewhat</i> | | | <i>A great deal</i> | |
| How much insight did you gain about the organization’s threats from the Risk Assessment? | 0 | 1 | 2 | 3 | 4 | G* | 6 |
| How much insight did you gain about the organization’s outcomes from the Risk Assessment? | 0 | 1 | 2 | 3 | G* | 5 | 6 |
| How much did the risk assessment change your perception of the organization’s risks? | 0 | 1 | 2 | 3 | G* | 5 | 6 |
| How much easier would it be to explain the organization’s risk priorities using the SAEM Risk Assessment than previous assessments? | 0 | 1 | 2 | G* | 4 | 5 | 6 |
| <hr/> | | | | | | | |
| How strongly would you approve or disapprove of submitting the risk assessment ranking to your CIO for use in making decisions about risk management? | <i>Strongly disapprove</i> | | <i>Neutral</i> | | | <i>Strongly approve</i> | |
| | -3 | -2 | -1 | 0 | G* | 2 | 3 |
| <hr/> | | | | | | | |
| How satisfied are you with the threat rankings? | <i>Very dissatisfied</i> | | <i>Neutral</i> | | | <i>Very satisfied</i> | |
| | -3 | -2 | -1 | 0 | G* | 2 | 3 |

What did you like about the Risk Assessment?

G*: Made me think on a bigger plane or level

G: It enabled you to take a broad look at potential risk and prioritize them.

What would you like to see improved in the Risk Assessment process or results?

G*:Automated for use by the manager on a frequent basis

G: Several categories were broad and overlapped with many different mitigation strategies. Also I would incorporate a method to support quantifying risks with a financial assessment.

II. Benefit Analysis: In the benefit analysis phase, participants estimated the effectiveness of security technologies against the threats resulting in a prioritization of security technologies.

| | | | | | | | |
|---|-----------------------------|-----------|-----------|-------------|-----------|----------|-------------------------|
| How difficult was it.... | <i>not at all difficult</i> | | | | | | <i>very difficult</i> |
| to identify and initially rank the security technologies? | 0 | 1 | 2 | G,G* | 4 | 5 | 6 |
| to estimate the effectiveness of the technologies? | 0 | 1 | 2 | G,G* | 4 | 5 | 6 |
| How much insight did you gain about the value that security technologies provide? | <i>none at all</i> | | | | | | <i>very much</i> |
| | 0 | 1 | G | G* | 4 | 5 | 6 |
| How much did the benefit analysis change your perception of the organization's security technologies? | 0 | 1 | G | 3 | G* | 5 | 6 |
| How much easier would it be to explain why a particular security technology should be purchased? | 0 | 1 | 2 | G | G* | 5 | 6 |
| How strongly would you approve or disapprove of submitting the benefit analysis results to your CIO for use in making decisions about spending financial resources? | <i>strongly disapprove</i> | | | | | | <i>strongly approve</i> |
| | -3 | -2 | -1 | G | G* | 2 | 3 |
| How satisfied are you with the security technology rankings? | <i>very dissatisfied</i> | | | | | | <i>very satisfied</i> |
| | -3 | -2 | -1 | G | G* | 2 | 3 |

What did you like about the Benefit Analysis?

G*:Again conscious thought versus buying or using because of media hype

G: Illustrated what technologies could be applied to mitigating a particular risk based on the prioritization.

III. Coverage Analysis: The coverage analysis showed the security technology defense-in-depth coverage (protect, detect, and recover) of the top six threats as determined by SAEM.

| | | | | | | | |
|--|----------------------------|-----------|-----------|----------|-----------|-----------|-------------------------|
| How much insight did you gain about the overall defense-in-depth coverage that your organization's current security technologies provide? | <i>none at all</i> | | | | | | <i>very much</i> |
| | 0 | 1 | G | 3 | G* | 5 | 6 |
| How much did the coverage analysis change your perception of the organization's security status? | 0 | G* | G | 3 | 4 | 5 | 6 |
| How much easier would it be to explain why a particular security technology should be purchased if the coverage analysis showed a gap? | 0 | 1 | 2 | G | 4 | G* | 6 |
| | | | | | | | |
| How strongly would you approve or disapprove of submitting the coverage analysis results to your CIO for use in making decisions about spending financial resources? | <i>strongly disapprove</i> | | | | | | <i>strongly approve</i> |
| | -3 | -2 | -1 | G | 1 | G* | 3 |
| | | | | | | | |
| How satisfied are you with the coverage analysis? | <i>very dissatisfied</i> | | | | | | <i>very satisfied</i> |
| | -3 | -2 | -1 | G | 1 | G* | 3 |

What did you like about the Coverage Analysis?

G*: You can easily see your coverage against the important threats.

What would you like to see improved in the Coverage Analysis process or results?

G*: Just automate the whole process.

CHAPTER 8. Analysis of the Method

8.1 Introduction

Since one of the goals of this thesis is to show the feasibility of using multi-attribute analysis techniques in security architecture design decisions, it is important to show that these techniques work across diverse organizations and that the underlying assumptions of the additive model correctly captures the security manager's understanding of the threat environment and effectiveness of risk-mitigation controls. Previous chapters independently analyzed three case studies, but did not explore the underlying assumptions of the threat index function or compare and contrast the case studies. This chapter explores the efficacy of the method and the underlying assumptions of the threat index function.

The three case studies presented in this thesis are from three different types of organizations—commercial, hospital, and government—and while they cannot be considered generally representative of these types of organizations, the security managers' perceptions of threats and the effectiveness of security technologies differed across case studies. However, without additional case study data, it is impossible to determine whether security managers are providing estimates that are consistent within their industries or even consistent with security managers as a whole, but it is clear from reviewing the satisfaction surveys that each organization found SAEM to be a useful and insightful process.

One of the reasons for using multi-attribute analysis in security architecture design decisions is to provide security managers with a general framework that helps them model their organization's security risks and evaluate security technology effectiveness against those risks. Although security managers with different levels of experience and expertise can use SAEM, the quality of the results depend heavily on the security manager's ability to assess the organization's threats and estimate the effectiveness of security technologies, which is often dependent on his or her experience and security expertise. For example, the hospital security manager, who appeared to have the least experience and expertise, identified the lowest number of threats for the organization and the least number of risk-mitigating technologies for those threats. In addition, his estimates of security technology effectiveness were higher than the other case study security managers. Despite his lack of expertise and experience, the security manager stated that the process helped him justify spending \$85,00 in additional detection mechanisms for the organization's security architecture.

In this chapter, each section is an analysis of one of the SAEM steps—risk assessment, benefit analysis, coverage analysis, and security technology tradeoff analysis—and includes the consolidated results of the satisfaction survey. In addition, this chapter presents some cross-case study analysis that shows the differences among case study security managers, their perceptions of their threat environment, and the perceived benefits that various technologies have in mitigating their organizations' threats.

8.2 Risk Assessment Analysis

The risk assessment determines an organization's threat priorities based on the threat index. The threat index³³ is computed using the security manager's estimated frequency of an attack and the most likely consequences of the attack. In the case studies, the consequence values are normalized using a linear function so that the consequence values can be added together. Intuitively, the linear normalization function assumes that as the consequence of an attack increases, so does the security manager's sensitivity or concern about the damage proportionally increase with respect to the damage. Similarly, the intuition behind the threat index function is that, as the frequency of an attack increases, there is a proportional increase in the security manager's sensitivity or concern about the threat. Neither of these assumptions, i.e., the security manager's sensitivity to consequences and frequencies, need be true; however these sensitivities be monotonically increasing.

This section explores alternatives to the threat index function and the linear normalization function, and it compares the risk assessment information provided by each case study security manager. It evaluates ranking threats by frequency as an alternative to the threat index function and explores two different normalization functions that could be used to indicate the security manager's sensitivity to frequency and consequences when ranking threats. These normalization functions reduce the influence that highly-frequently occurring threats have on the organization's threat ranking. Next, four alternative value functions are evaluated to determine whether SAEM's threat index function could more closely predict the government security manager's final ranking using these alternative functions. Since the goal of SAEM is to provide insight into the security manager's security technology selection process, the results of using these alternative normalization functions can help explain how the security manager's sensitivity to frequencies and consequences influence his or her threat rankings.

Exploration of alternative normalization functions resulted in different threat rankings. An important point to consider when looking at the changes in threat rankings is how much of the variation in rankings is due to the variation caused by running different simulations. ASESS conducted ten simulations of 1,000 iterations of the data from each case study. The analyst compared the results from each simulation and determined that the rankings were consistent among all simulations, although the threat indexes varied slightly among the simulations.

8.2.1 Analysis of Assumptions

8.2.1.1 Threat Prioritization Based on Expected Frequencies

At first glance, the frequency of an attack appears to have a significant impact on the final SAEM threat rankings in each of the case studies. In the commercial case study, the security manager estimated that *Virus* attacks occurred with far greater frequency than any other attack and SAEM ranked viruses as the most significant threat for the organization. In the hospital and government case studies, SAEM ranked the three most frequently occurring attacks as the three most significant threats to the organizations. Although the case study security managers did not always rank threats according to frequency, it is possible that the frequency of an attack may more closely predict the security manager's final threat ranking than the SAEM risk

³³ $TI_a = Freq_a * (\sum_{j=attributes} w_j * v_j(x_{aj}))$

assessment ranking. If the security manager’s threat ranking is highly correlated with frequency, then the risk-assessment process could be simplified and still produce similar results.

In fact, frequency rankings do not predict security managers’ final threat rankings as the closely as SAEM rankings do. The tables at the end of this chapter show a comparison of the final SAEM and security manager rankings with threats ranked according to expected frequency. Table 8 - 1 shows the correlation across the rankings, i.e., the SAEM final ranking, the Expected Frequency ranking, and the security manager’s final ranking. For each case study, the SAEM Rank is more correlated with the security manager’s final ranking than is the Expected Frequency ranking. For example, the SAEM final rank had a .96 correlation with the commercial case study security manager’s final ranking, whereas the Expected Frequency Rank had a .85 correlation with the her final ranking. Therefore, the threat index function provides a more robust prediction of the security manager’s final threat ranking than does the expected frequency alone. This demonstrates that the consequences of an attack influence the threat rankings of the case study security managers.

Table 8 - 1 SAEM and Expected Frequency Correlations to Security Manager Final Rankings

| | Commercial | | Hospital | | Government | |
|-------------------------------|------------|-------------------------|-----------|-------------------------|------------|-------------------------|
| | SAEM Rank | Expected Frequency Rank | SAEM Rank | Expected Frequency Rank | SAEM Rank | Expected Frequency Rank |
| Expected Frequency Rank | 0.76 | 1.0 | 0.97 | 1.0 | 0.88 | 1.0 |
| Final Security Manager's Rank | 0.96 | 0.85 | 0.85 | 0.81 | 0.57 | 0.33 |

8.2.1.2 Alternative Frequency Transformation Functions

Since ranking threats by estimated frequency resulted in less correlation to the government case study manager’s final threat ranking, it is possible that the manager does not consider the frequency of attacks to be as important as the outcomes of attacks in assessing threats. Perhaps, the threat index function did not adequately reflect the importance of frequency in his determination of threat priorities, i.e., the threat index function overemphasizes frequency in the government case study. The results in Table 8 - 1 indicate the consequences of an attack influence the security manager’s threat rankings. Since security managers did not always rank highly-frequently occurring threats as significant to the organization, it is possible that the frequency variable of the threat index function does not adequately capture the security manager’s assessment of the effect that frequency has in determining threat rankings.

The relative threat index function, as presented in Chapter 3, assumes that there is a direct linear relationship with frequency, because the threat index function multiplies the summed outcomes of an attack by the frequency. However, at some point an increase in attack frequency may not result in proportional increases in a threat’s relative significance. At some point, the security manager’s perception of the threat is that “it happens a lot” and any increase in the frequency is still perceived simply as “a lot.” For example, the security managers

estimated that several of their organizational threats occurred hourly, resulting in over 200,000 incidents per year. If the frequency increased by an additional 50,000 incidents per year -- a 25% increase -- a security manager would not likely change his or her ranking of the significance of the threat relative to other threats. In contrast, a weekly occurring threat, the frequency of which suddenly increased 25% might change the significance of the threat relative to other threats.

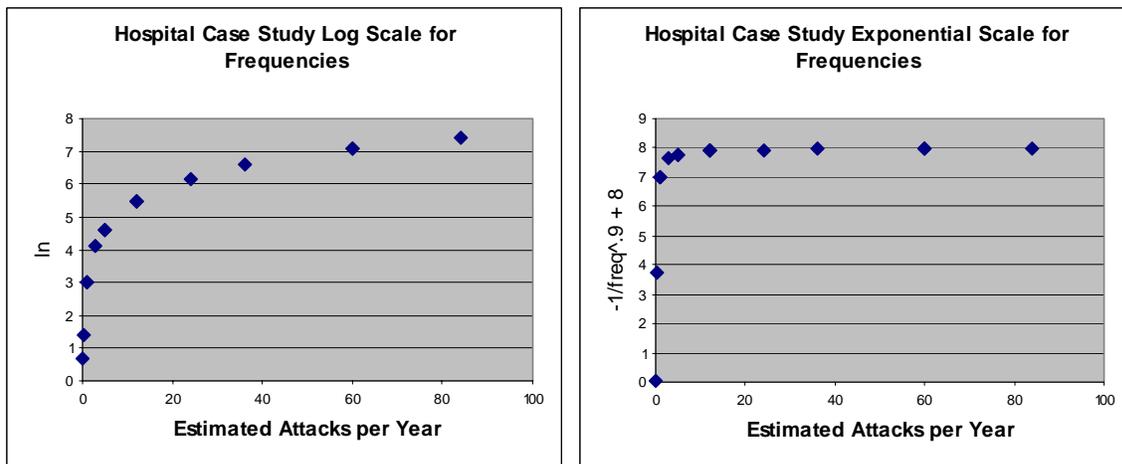
Convex frequency transformation functions can reduce the impact of highly-frequently occurring threats in SAEM's risk assessment threat ranking. Ideally, the analyst could elicit specific thresholds that would establish the specific shape of a frequency function, but this section explores two convex frequency functions to determine whether elicitation would have resulted in a closer prediction of the government security manager's final threat ranking. The idea is to explore the following functions in reducing the effect of frequency in computing the threat index:

$$\ln(freq_a) + b$$

$$-1/freq_a^{.9} + b$$

where $freq_a$ is the estimated frequency of a threat (a) and b is a constant that shifts the function to ensure that the function does not result in negative values. For example, Figure 8 - 2 shows the natural logarithmic and the exponential frequency normalization functions for the hospital case study that reduced the effect of the frequency term in the threat index function.

Figure 8 - 1 Examples of Frequency Transformation Functions



The logarithmic frequency function has a more gradual slope than the exponential frequency function. The slopes of these functions reduce the significance of highly-frequent attacks relative to less frequently occurring attacks; the exponential slope reduces the effect more quickly than does the logarithmic function. Table 8 - 2 compares the correlation results of the case study threat rankings using the natural logarithm (ln) and exponential (exp) frequency normalization functions with the security managers' final threat rankings.

Table 8 - 2 Frequency Function Results Correlated with Security Manager Final Rankings

| | Commercial | | | Hospital | | | Government | | |
|------------|------------|---------|----------|-----------|---------|----------|------------|---------|----------|
| | SAEM Rank | In Rank | exp Rank | SAEM Rank | In Rank | exp Rank | SAEM Rank | In Rank | exp Rank |
| <i>In</i> | 0.81 | | | 0.61 | | | 0.98 | | |
| <i>exp</i> | 0.81 | 0.78 | | 0.14 | 0.81 | | -0.09 | 0.03 | |
| SM Final | 0.96 | 0.74 | 0.69 | 0.85 | 0.57 | 0.22 | 0.57 | 0.61 | 0.42 |

Although the frequency transformation functions' correlation with the security manager's final threat ranking did not improve in the commercial case study or the hospital case study, the natural logarithm frequency transformation did result in a slightly higher correlation with the government case study's security manager's final threat ranking than the with the SAEM risk assessment's correlation with the security manager's threat ranking. The improved correlation between the natural logarithm function and the security manager's final threat ranking might indicate that the analyst should conduct additional interviews to determine the security manager's assessment of frequency in ranking threats. However, recall that in the government case study the Director of Mission Assurance indicated that the manager may not have ranked the threats according to their significance to the organization independent of existing risk mitigation efforts. Given the remarks of the Mission Assurance Director and the strong correlations seen in the hospital and commercial case studies, the SAEM risk assessment threat index function appears to provide the most robust prediction of the security manager's final threat prioritization.

8.2.1.3 Value Functions

Recall from Chapter 3 that the purpose of the value function ($v_j(x_{ij})$) is to transform the outcome attributes to a 0-1 scale so that the outcome values can be compared. In previous chapters, the analysis of the case studies used a linear value function.³⁴ The underlying assumption of the linear value function is that the security manager's concern about the outcome values increases linearly as the damages from an attack increase. As with the frequency of an attack, security managers may have thresholds at which they become more concerned as the damages increase; therefore a linear value function may not adequately capture the security manager's sensitivity to attack outcomes.

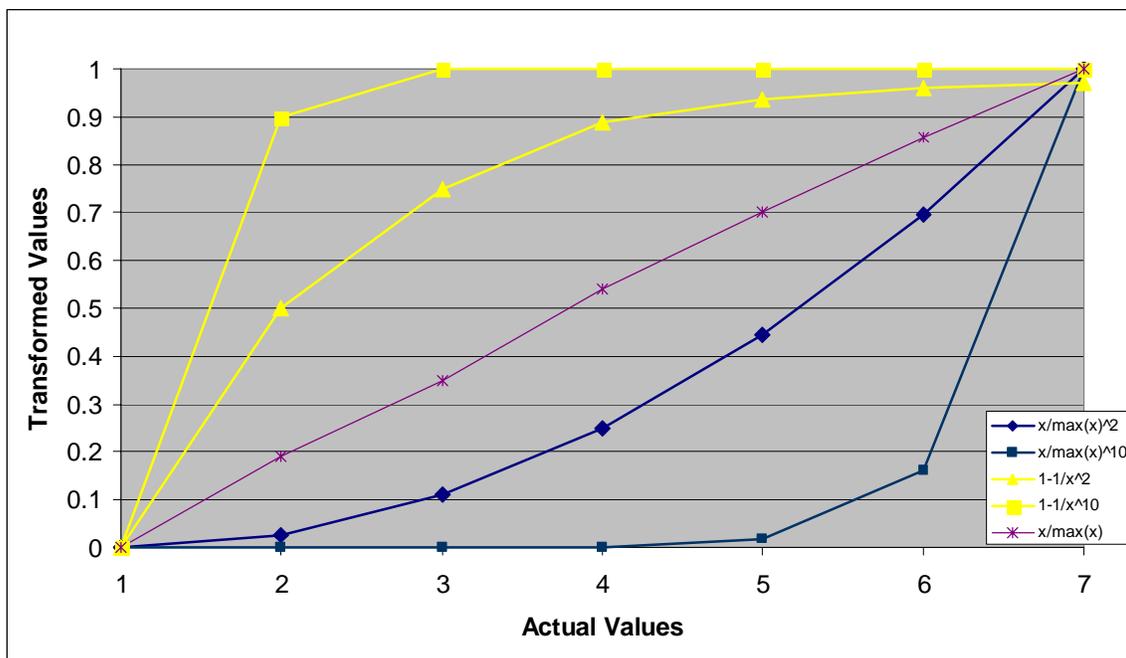
Two possible scenarios are worth exploring with respect to the security manager's sensitivity to attack outcomes. First, the security manager may be relatively insensitive to attacks that result in low consequence values, but concern may rapidly increase as consequences reach a particular threshold. For example, many of the most significant threats, such as Personal Computer Abuse and Viruses, resulted in relatively low productivity losses or had minimal impact to the organization's public image. The security manager could be equally minimally sensitive to attacks that result in none or minimal damage to public image, (a rating of 1 or 2 on the Likert scale), but become rapidly more concerned when a Virus attack results in at least moderate damage to public image (4 on the Likert scale), and become greatly concerned about

³⁴ x/x_i^* , where x_i^* is the max value of the attributes

attacks that result in moderately or moderately severe damage to public image (> 5 on the Likert scale). Therefore, security managers could rank attacks that result in less severe consequences disproportionately lower than attacks that result in more severe consequences.

In contrast, the security manager might be very concerned earlier, when the consequence of an attack appears to have little impact and be equally concerned for any attack that results in moderate damage or greater. For example, again using the 7-point Likert Scale, the security manager may not significantly distinguish between attacks that result in moderate damage or greater (a rating of 4 or greater on the Likert scale). Each of these scenarios can be represented by non-linear value functions. Intuitively, concave value functions represent the first scenario where the security manager does not really distinguish among consequences that result in relatively low amounts of damage, but grows rapidly more sensitive to attacks that result in relatively higher amounts of damage. Convex value functions represent the second scenario where the security manager is relatively concerned about attacks that result in low amounts of damage and at some point nearly equally highly concerned about attacks that result in more than a moderate amount of damage. Figure 8 - 2 represent two concave and convex value functions whose curves model possible changes in consequence sensitivity.

Figure 8 - 2 Concave, Linear, and Convex Transformation Functions



Although the analyst did not elicit specific value function thresholds during the case study interviews, the analyst can use these concave and convex value functions to help understand the security manager's sensitivity to outcome attribute values. Value functions may differ for each attribute. Additional interviews could determine the security manager's specific value function for each attribute, but the analyst could conduct an initial analysis to see whether the identification of a non-linear value function would provide additional insight about the security manager's threat rankings.

Table 8 - 3 shows the correlation results of using concave and convex value functions.³⁵ In the commercial and the hospital case studies, the correlation of threat rankings produced using these value functions with the security manager’s final threat ranking is slightly lower than the correlation of the threat ranking produced using the linear value function. In the government case study, only the threat ranking using the $x/\max(x)^{10}$ value function resulted in a higher correlation, but not significantly higher, with the security manager’s final threat ranking. The correlations of all the threat rankings are shown in Table 8 - 3 to indicate how much the rankings varied using different value functions.

Table 8 - 3 Correlation of Threat Rankings for Different Value Functions

| | | linear | 1-1/x ² | 1-1/x ¹⁰ | x/max(x) ² | x/max(x) ¹⁰ |
|------------|------------------------|--------|--------------------|---------------------|-----------------------|------------------------|
| Commercial | 1-1/x ² | 0.93 | 1.00 | | | |
| | 1-1/x ¹⁰ | 0.93 | 0.98 | 1.00 | | |
| | x/max(x) ² | 0.94 | 0.86 | 0.88 | 1.00 | |
| | x/max(x) ¹⁰ | 0.78 | 0.84 | 0.84 | 0.78 | 1.00 |
| | SM Final | 0.96 | 0.91 | 0.91 | 0.89 | 0.78 |
| Hospital | 1-1/x ² | 0.85 | 1.00 | | | |
| | 1-1/x ¹⁰ | 0.85 | 1.00 | 1.00 | | |
| | x/max(x) ² | 0.86 | 0.94 | 0.93 | 1.00 | |
| | x/max(x) ¹⁰ | 0.54 | 0.77 | 0.75 | 0.85 | 1.00 |
| | SM Final | 0.85 | 0.73 | 0.75 | 0.70 | 0.49 |
| Government | 1-1/x ² | 1.00 | 1.00 | | | |
| | 1-1/x ¹⁰ | 0.99 | 0.99 | 1.00 | | |
| | x/max(x) ² | 1.00 | 1.00 | 0.99 | 1.00 | |
| | x/max(x) ¹⁰ | 0.90 | 0.90 | 0.87 | 0.90 | 1.00 |
| | SM Final | 0.57 | 0.57 | 0.53 | 0.56 | 0.66 |

8.2.2 Case Study Variability

Although there were similarities among the security managers’ threat rankings, overall, the security managers differed in their final threat rankings. All case study security managers ranked *Compromise* as one of their top 5 concerns. In addition, the security managers all ranked *Alteration* and *Signal Interception* within the top ten organizational threats and ranked *Cryptographic Compromise* near the bottom.

Despite these few similarities, case study security managers differed greatly in their perceptions of their organizations’ threats. This is, of course, appropriate, because they have different missions resources, responsibilities, and legitimate concerns. Therefore, additional research is necessary to establish baseline industry-specific threat data so that security managers can compare their assessments against their industry estimates—adjusting their own estimates as necessary. However, the diversity of the case studies shows the feasibility of using multi-attribute analysis risk assessments with diverse organizations.

³⁵ In this analysis, one type of value function was used for all outcome attributes.

8.2.3 Survey Assessment

During the risk assessment interviews, the commercial and hospital participants struggled with the identification and estimation of threats and attack frequencies. Neither of these organizations had completed a risk assessment prior to these interviews, so the risk assessment process was particularly difficult for them. Although the government case study security manager had not completed a formal risk assessment prior to the interview, the incident response center had been collecting data about the organization's threats so the manager could rely on empirical data for many of the estimates.

Overall, all participants found the process insightful and very helpful in explaining their organization's risk priorities. After completing all the phases of SAEM, each security manager said that he or she would have liked to have developed a different set of threats for the risk assessment. They thought that some of the threats were too general or overlapping with each other. Recall that the analyst provided an initial list of threats so the security manager could refine or tailor the list to the organization's specific risks; however, none of the security managers substantially modified the analyst's initial threat list, despite encouragement to do so. Each of the security managers realized how he or she should have modified the list at the end of the SAEM process. As organizations repeat the SAEM risk assessment process, they will continue to develop their organization's threats, but less experienced security managers would greatly benefit from an industry specific established set of threats.

Table 8 - 4 Consolidated Satisfaction Survey

| | Commercial | Hospital | Government | Average |
|---|---|-----------------|-------------------|----------------|
| How difficult was it ... | Using 7 pt scale with 0 = not at all difficult and 6 = very difficult | | | |
| ... to identify and initially rank the threats? | 5 | 5 | 3 | 4.3 |
| ... to estimate the frequency of attacks? | 4 | 5 | 2 | 3.6 |
| ... to estimate the outcomes? | 5 | 3 | 2 | 3.3 |
| | Using 7 pt scale with 0 = none at all and 6 = a great deal | | | |
| How much insight did you gain about the organization's threats from the Risk Assessment? | 4 | 5 | 5 | 4.7 |
| How much insight did you gain about the organization's outcomes from the Risk Assessment? | 5 | 4 | 4 | 4.3 |
| How much did the risk assessment change your perception of the organization's risks | 4 | 4 | 4 | 4 |
| How much easier would it be to explain the organization's risk priorities using the SAEM Risk Assessment than previous assessments | 5 | 4 | 3 | 4 |
| | Using 7 pt scale -3 =strongly disapprove and 3 = strongly approve | | | |
| How strongly would you approve or disapprove of submitting the risk assessment ranking to your CIO for use in making decisions about risk management? | 1 | 1 | 1 | 1 |
| | Using 7 pt scale -3 very dissatisfied and 3 = very satisfied | | | |
| How satisfied are you with the threat rankings? | 2 | 2 | 1 | 1.7 |

8.2.4 Conclusions

The risk-assessment process results show the feasibility of using multi-attribute analysis techniques to help security manager's rank their threats. The commercial and hospital case study security managers modified their initial threat rankings based on the risk assessment

process, which shows that the process influenced their final threat priorities. In addition, SAEM's final rankings most closely correlated with the commercial and hospital case-study managers' final threat rankings, showing that the SAEM rankings closely represent the security manager's experience and knowledge about their threat environment.

Often one of the results of using decision analysis techniques is greater insight about the uncertainty, objectives, and tradeoffs of the problem. It appears that the security managers gained greater insight into the organization's threat priorities, since all case study security participants stated that the process was very insightful. Whether this greater insight produced better threat rankings is still a research problem, but using multi-attribute analysis techniques helped structure the risk analysis and identified the security manager's assumptions about frequency and outcomes.

For the case studies presented in this thesis, the threat index function that uses linear value functions and multiplies the summed outcome values by the estimated frequencies appears to most closely represent the security managers' understanding of attack consequences and attack frequencies. Future risk assessments may show that frequencies and outcomes influence other security managers differently when prioritizing their organizations' risks, but the threat index function presented in this thesis provides a robust mechanism for modeling these influences.

The variability among the different case studies also shows the robustness of the threat index function in providing insight for various levels of experience and expertise. It also points out the need for baseline empirical threat data so that security manager assessments can be evaluated. Although the task of identifying threats and estimating frequency and outcome values was difficult for the less experienced security managers, all security managers gained insight into their organizations' risks. An established threat database would be very helpful to security managers who had not carefully identified their organizations' threats prior to completing the SAEM risk assessment process.

8.3 Benefit Analysis

8.3.1 Benefit Ranges

During the benefit analysis phase, each case study security manager estimated the effectiveness of security technologies for each of their risk assessment threats. The analyst asked the security manager to provide his or her best estimate of the effectiveness of each security technology given the organization's threat and operational security environment. As in the risk assessment, each case study security manager appeared to have different levels of experience and knowledge about the security technologies. For example, the government security manager was very knowledgeable about all of the security technologies and how they could mitigate the organization's risks, but the hospital's security manager asked many questions about how some of the technologies specifically functioned. The commercial case study security manager's knowledge about security technologies appeared to fall closer to the government security manager's expertise, but was not quite as advanced.

Unfortunately, security managers cannot consult an authoritative source for the effectiveness of security technologies because 1) no such authoritative source exists, and 2) the actual effectiveness of a security technology is dependent not only on the sophistication of the attack, but also on the organization's ability to establish effective security policies, correctly install and configure the security technology, and appropriately maintain the technology once it is

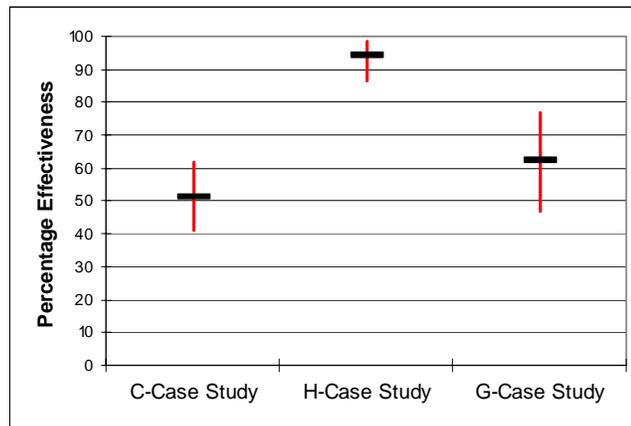
implemented into the security architecture. In addition, some organizations are more likely to be targeted by highly skilled hackers than are others, so the percentage of skilled attacks versus unskilled will affect the effectiveness rating of a security technology. Table 8 - 5 shows some statistics about each case study security manager's effectiveness estimates for security technologies. For example, the commercial case study manager made 196 effectiveness estimates during the benefit analysis phase. On average, this security manager estimated that the most likely effective rate of a security technology was 51.39%, with a standard deviation of 25.18%. The government case study security manager made 530 security technology effectiveness estimates, but the hospital case study security manager made only 69 security technology estimates with an average effectiveness rate of 94.42%, significantly higher than the other case study security managers.

Table 8 - 5 Case Study Effectiveness Estimate Descriptive Statistics

| | Commercial | | | Hospital | | | Government | | |
|--------------------|-------------|---------------|------------|-------------|---------------|------------|-------------|---------------|------------|
| | <i>High</i> | <i>Likely</i> | <i>Low</i> | <i>High</i> | <i>Likely</i> | <i>Low</i> | <i>High</i> | <i>Likely</i> | <i>Low</i> |
| Mean | 61.78 | 51.39 | 40.94 | 98.38 | 94.42 | 86.45 | 76.72 | 62.12 | 46.88 |
| Standard Error | 1.76 | 1.81 | 1.80 | 0.28 | 0.50 | 1.11 | 0.91 | 0.92 | 0.85 |
| Median | 60 | 50 | 40 | 99 | 95 | 90 | 80 | 70 | 50 |
| Mode | 80 | 40 | 30 | 100 | 95 | 90 | 80 | 80 | 50 |
| Standard Deviation | 24.60 | 25.29 | 25.18 | 2.31 | 4.13 | 9.24 | 20.96 | 21.28 | 19.51 |
| Minimum | 15 | 10 | 5 | 90 | 80 | 50 | 5 | 2 | 1 |
| Maximum | 100 | 98 | 95 | 100 | 99 | 95 | 99 | 95 | 80 |
| Count | 196 | | | 69 | | | 530 | | |

Figure 8 - 3 Effectiveness Ranges

The government case study security manager's estimates had the greatest difference between average high (76.72%) and average low (46.88%) estimates of security technology effectiveness -- a difference of 29.44%. In contrast, the hospital case study's security manager had the least difference. His average high estimate was 98.38% and his average low was 86.45%, a difference of 11.93%. Figure 8 - 3 compares the average range of effectiveness estimates for each case study.



Although there is no way to confirm whether the estimates provided by the hospital security manager reflect the actual effectiveness of the technologies, many security professionals may find his estimates optimistic. Additional research might be able to determine whether the manager's estimates were consistent with other, more experienced, hospital security managers.

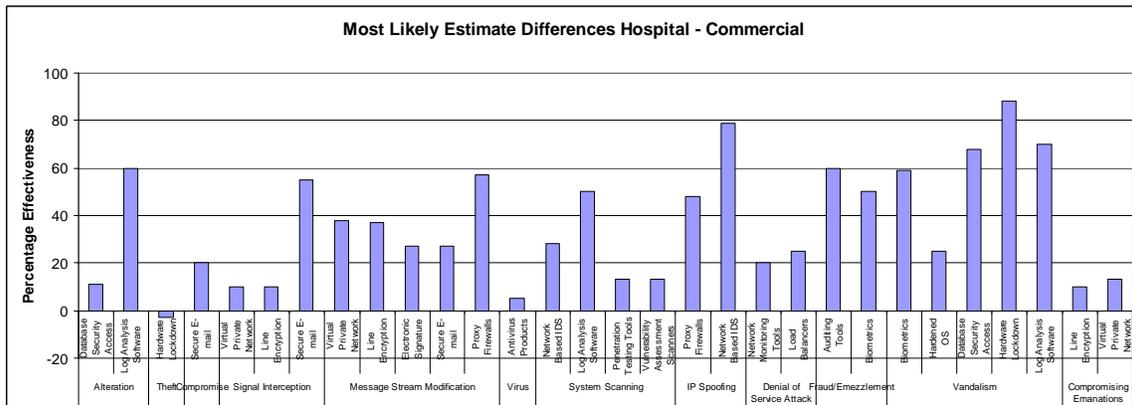
8.3.2 Comparison of Technology Effectiveness

8.3.2.1 Security Technology Effectiveness Estimates Among Case Studies

The descriptive statistics show that the hospital security manager's average effectiveness estimates were significantly higher than those of the other security managers, and that the commercial security manager's effectiveness estimates were lower, on average, than the government security manager's estimates. However, the security managers did not select the same security technologies for each threat so comparing estimated effectiveness averages may not actually reflect an accurate picture of their differences.

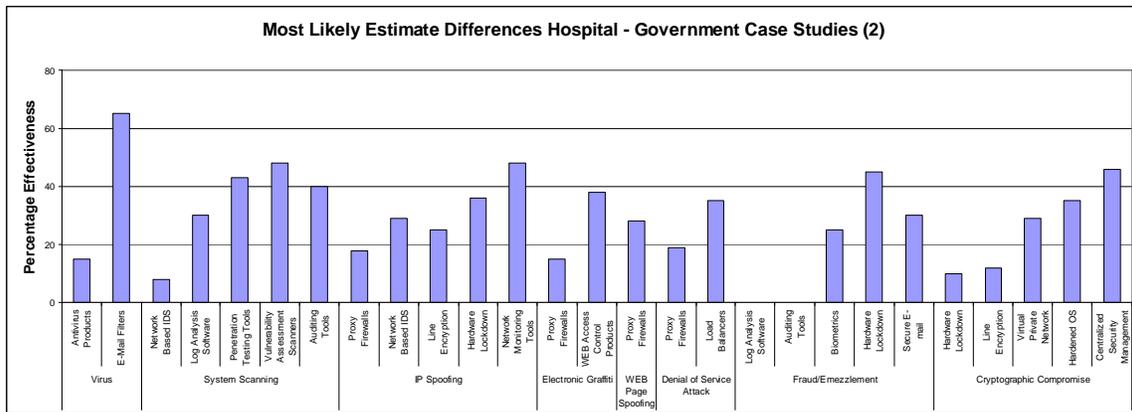
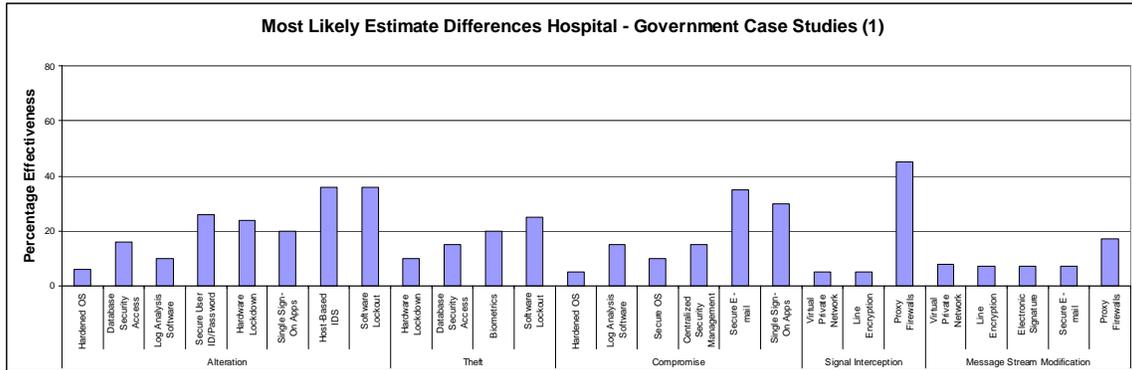
However, when assessing the same threats, the hospital security manager consistently estimated the effectiveness of a security technology higher than the other security managers. Furthermore, the government case study security manager estimated the effectiveness of a given security technology higher than the commercial case study security manager did in 70% of his estimates. Figure 8 - 4 shows a histogram of the differences between the most likely effectiveness estimates for all threats for which both the hospital and commercial security manager identified a similar security technology as risk mitigating. In all cases except one, Hardware Lockdown for *Theft*, the hospital security manager estimated the technology as being more effective than the commercial security manager did for the given threat.

Figure 8 - 4 Differences between Hospital and Commercial Security Manager Effectiveness Estimates



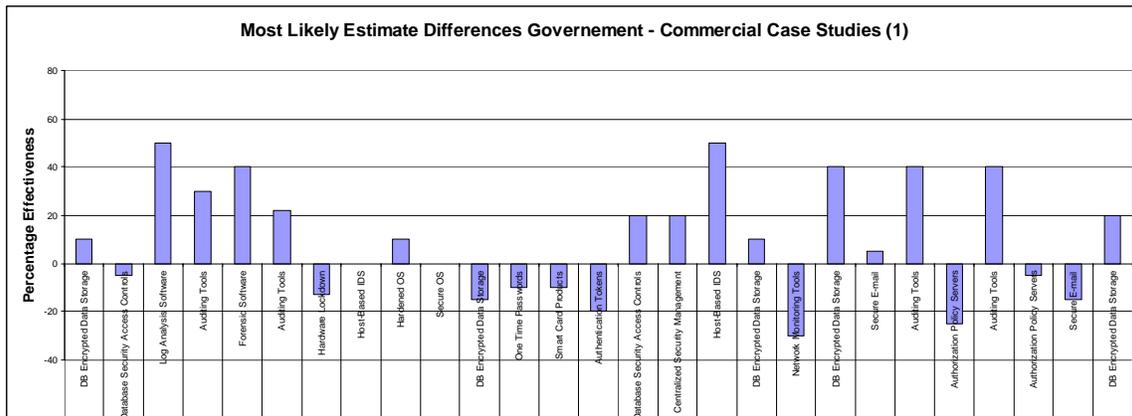
The hospital security manager identified some security technologies as risk mitigating for threats that the commercial security manager did not, so the two managers overlapped for only 30 total threat/security technologies of the hospital security manager's 69 security technology estimates. Figure 8 - 5 shows a histogram of the differences between the most likely effectiveness estimates for all the threats for which the hospital security manager and the commercial security manager identified similar security technologies as risk mitigating. Again, the hospital security manager nearly always estimated that the technology was more effective than the government security manager estimated.

Figure 8 - 5 Differences between Hospital and Government Security Manger Effectiveness Estimates



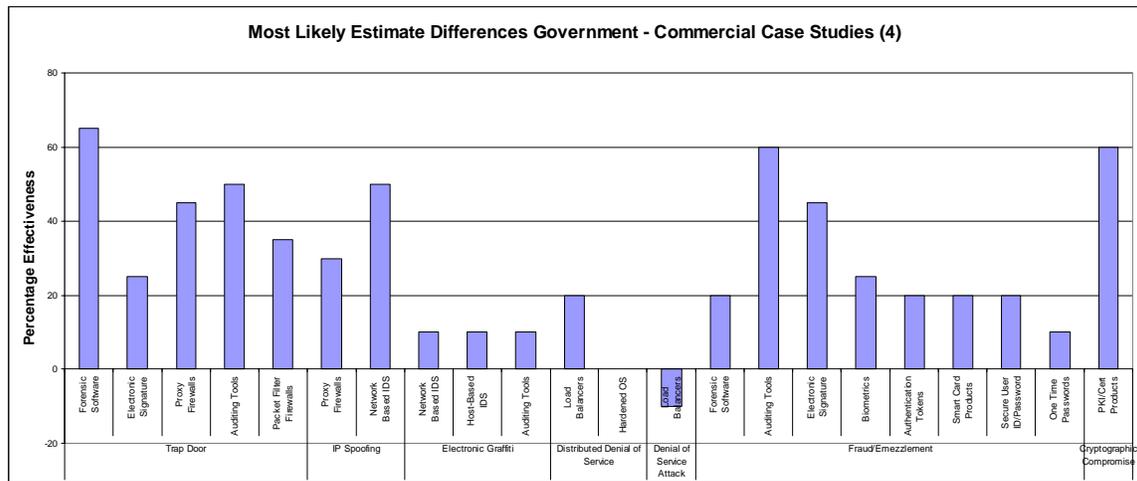
Although, on average, the commercial case study security manager estimated the effectiveness of security technologies lower than did the government case study manager, she did not always estimate the effectiveness lower for a given threat. Figure 8 – 6 shows the differences between the government security manager and the commercial security manager for each technology that both security managers identified as risk mitigating for the same threat.

Figure 8 - 6 Differences between Government and Commercial Case Study Security Manager Effectiveness Estimates



The average difference between the commercial security manager's estimates and the government security manager's estimates was 28%, but 10% of the time, the difference between their estimates was greater than 50 percentage points, with the government security manager registering the higher estimate. For example, the government security manager estimated that Forensic software is 95% effective, but the commercial case study manager estimated that it was only 15% effective for Logic Bombs, a difference of 80%.

Figure 8 - 7 Differences between Government and Commercial Case Study Security Manager Effectiveness Estimates

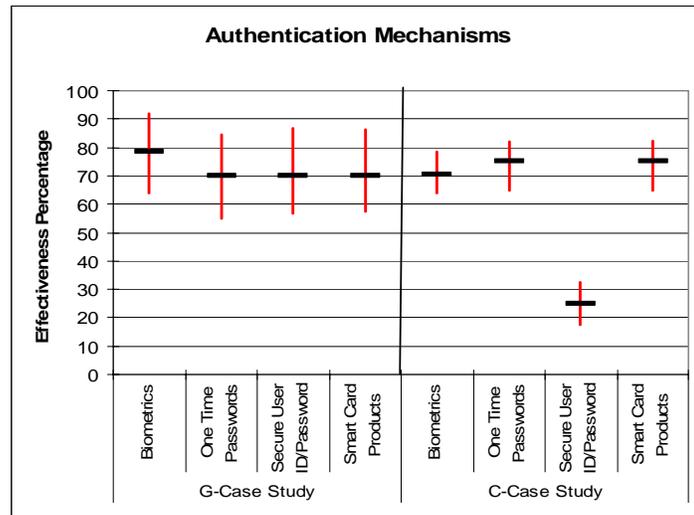


8.3.2.2 Authentication Mechanisms

During the interviews the government case study security manager and the commercial case study security manager appeared to rate all authentication mechanisms as about equally effective. Figure 8 - 8 shows the average effectiveness of the four authentication mechanisms used during the benefit analysis phase: biometrics, one time passwords, smart cards, and User ID/passwords. (The hospital security manager only selected User Id/passwords as an effective authentication mechanism.)

Figure 8 - 8 Authentication Mechanism Comparisons

Not only are most of the technologies rated approximately the same—somewhere between 70-80% effective—but the government case study security manager didn't perceive any difference between the effectiveness of one time passwords, smart cards, and User Id/passwords. The commercial case study security manager rated User Id/passwords as only 25% effective on average because the organization had so many incidents of password sharing that she felt that User Id/passwords were relatively ineffective compared to the alternative authentication mechanisms.



If security managers do not perceive a real benefit from switching from a user id/password based authentication system to a stronger form of authentication, then it is unlikely that they will select these alternatives because they are generally more costly and difficult to maintain. Interestingly, the commercial and hospital case study security managers included smart cards in their security tradeoff analysis phase as technologies that they might consider for their organizations although they didn't consider them any more effective than what they already had in place.

8.3.3 Survey Assessment

Overall, the security managers found it difficult to identify and rank security technologies. The hospital security manager found it most difficult, perhaps because he didn't appear to be familiar with many of the security technologies presented during this phase of SAEM. The government participants did not find it as difficult to estimate the effectiveness of the security technologies, but all of the security managers found the process tedious. All of the managers made mistakes during the process, such as overlooking obvious security technologies or making estimates that were not consistent with similar technologies. Obvious mistakes were usually quickly spotted, but a less tedious process of eliciting effectiveness estimates would help reduce errors.

The commercial and government security managers found this phase somewhat insightful, but the hospital security manager found it the most insightful. During the interviews, the analyst described some of the security technologies, so the process was more helpful to him than to the other managers. Generally, the security managers felt that the process would be helpful in justifying the purchase of a security technology, despite their assessment that the process would not be terribly helpful to present to their CIO's. Of the three case studies, only the hospital security manager was actively considering the purchase of additional security technologies. This process may be more valuable to security managers who are actually trying

to make decisions about which security technologies to purchase, rather than as a tool for the general assessment of security technologies for an organization.

While there the satisfaction survey shows moderate insight and value in the benefit analysis phase, the real value in the results from this step is that it provides information necessary for the coverage analysis step. The coverage analysis populates the defense-in-depth model with the risk-mitigating security technologies that the security manager identified during the benefit analysis. In addition, when coverage analysis identifies weaknesses in the security architecture, the prioritized security technology results from this step show which technologies are the best candidates for reducing the risk.

Table 8 - 6 Benefit Analysis Consolidated Survey

(Government Director's and Deputy Director's ratings are averaged)

| | Commercial | Hospital | Government | Average |
|---|---|----------|------------|---------|
| How difficult was it ... | Using 7 pt scale with 0 = not at all difficult and 6 = very difficult | | | |
| to identify and initially rank the security technologies? | 4 | 6 | 3 | 4.3 |
| to estimate the effectiveness of the technologies? | 5 | 5 | 3 | 4.3 |
| | Using 7 pt scale with 0 = none at all and 6 = a great deal | | | |
| How much insight did you gain about the value that security technologies provide? | 2 | 5 | 2 | 3 |
| How much did the benefit analysis change your perception of the organization's security technologies? | 3 | 4 | 3 | 3.3 |
| How much easier would it be to explain the organization's risk priorities using the SAEM Risk Assessment than previous assessments | 4 | 4 | 3.5 | 3.8 |
| | Using 7 pt scale -3 =strongly disapprove and 3 = strongly approve | | | |
| How strongly would you approve or disapprove of submitting the benefit analysis results to your CIO for use in making decisions about spending financial resources? | 1 | 1 | .5 | .67 |
| | Using 7 pt scale -3 very dissatisfied and 3 = very satisfied | | | |
| How satisfied are you with the security technology rankings | 2 | 2 | .5 | 1.5 |

8.3.4 Conclusions

As with the risk assessment, the benefit analysis indicated a need for more industry-specific or canonical information about the effectiveness of security technologies. The hospital security manager's estimates appear to be too high when compared to the other security managers' estimates, but his estimates may be valid, especially if his organization doesn't experience the highly skilled attacks that the government and commercial organizations experience. The hospital security manager would greatly benefit from a database that identified risk mitigation technologies and their effectiveness for the hospital threat environment.

The benefit analysis process successfully provided the case study participants with some insights about the value of security technologies, but the participants were not as pleased with the results as they were with those of the risk assessment process. Although the benefit analysis ranked all the security technologies relative to all of the threats, the process might have been more meaningful if it had been tailored to specific choices that the security manager needed to make. In each case study, SAEM ranked 40+ technologies, which may seem a bit too artificial unless the security manager is trying to decide among a few security technologies. One of the advantages of evaluating all of the security technologies is that a security manager can see if a particularly beneficial technology had been overlooked. The analyst in most interviews found it necessary to remind the security manager to evaluate the technologies effectiveness as if the security architecture were not in place. If the benefit analysis process were modified to focus more on relevant decisions, the number of effectiveness estimates could be reduced, making the process less tedious and the results more meaningful and insightful.

8.4 Coverage Analysis

The purpose of the coverage analysis phase was to show security managers how their existing security technologies provide defense-in-depth and breadth-of-coverage against various threats. The coverage analysis showed the commercial and hospital case-study managers that there were weaknesses in their security architectures. They used the analysis as justification to add detection mechanisms to the architectures. Although the government security manager did not discover any surprising gaps in his security architectures, the Director and Deputy Director were satisfied with the coverage analysis. Overall, each case study security manager reported that he or she gained insight from the process. Each said the coverage analysis would make it easier (a rating of 5 points out of 6) to explain why a particular security technology should be purchased from the coverage analysis, but only the government case study Director of Mission Assurance said that the coverage analysis did not change her perception of the organization's security status. This is not surprising since the organization had nearly all of the security technologies presented during the benefit analysis phase. Overall, it appears from the satisfaction surveys and my observation of the managers' enthusiasm during the coverage analysis presentations, that each one felt that the coverage analysis was valuable and contributed to their decision making process. Table 8 - 7 shows the consolidated results from the case study surveys.

Table 8 - 7 Consolidated Coverage Analysis Surveys

| | Commercial | Hospital | Government | Average |
|--|---|-----------------|-------------------|----------------|
| How difficult was it ... | Using 7 pt scale with 0 = not at all difficult and 6 = very difficult | | | |
| How much insight did you gain about the overall defense-in-depth coverage that your organization's current security technologies provide? | 5 | 4 | 4 | 4.3 |
| How much did the coverage analysis change your perception of the organization's security status? | 4 | 4 | | 3.6 |
| How much easier would it be to explain why a particular security technology should be purchased if the coverage analysis showed a gap | 5 | 5 | 5 | 5 |
| | Using 7 pt scale -3 =strongly disapprove and 3 = strongly approve | | | |
| How strongly would you approve or disapprove of submitting the coverage analysis results to your CIO for use in making decisions about spending financial resources? | 1 | 1 | 1 | 1 |
| | Using 7 pt scale -3 very dissatisfied and 3 = very satisfied | | | |
| How satisfied are you with the coverage analysis? | 2 | 1 | 2 | 1.7 |

8.5 Security Technology Tradeoff Analysis

The security technology tradeoff analysis did not prove to be insightful to the participants, in part because of the hypothetical nature of the analysis and because the analysis compared alternatives as if existing architectures were not already providing some level of security. In addition, the commercial and hospital security manager ranked cost as an important attribute in their decision framework, but the cost of a security technology is not easily determined. Vendors often are not explicit about their prices, and prices can vary by thousands of dollars among vendors for a similar type of technology. For example, a quick review of trade journals determined that the published price of a packet filter firewall ranged from \$500 to \$25,000. During the tradeoff analysis, the analyst asked the security manager to rank each alternative relative to the cost attribute, but it was obvious that the security managers could not rank the alternatives with much confidence.

During the SAEM risk assessment and benefit analysis phases, the analyst asked the security manager to provide initial rankings; then these initial rankings were compared with the results of the process. Much of the security manager's insight is developed when the initial rankings are compared with the final rankings. In the technology tradeoff phase, the analyst used the same comparison process, but the commercial and hospital security managers had chosen security technology alternatives that had comparable functionality in the organization's existing security architecture, so the tradeoff analysis ranking results were negatively correlated with the security managers' technology rankings.

For example, the commercial and hospital security managers wanted to evaluate smart cards as one of their alternatives, but both information system security architectures already included User id/passwords. In contrast, both case study organizations were weak in detection mechanisms and the security managers ranked detection mechanisms as their first choice. Most security professionals would argue that an authentication mechanism is essential to an organization's security architecture, and detection mechanisms are prudent, but not essential. Therefore, if the organization's security architecture lacked an authentication mechanism, the security managers probably would have ranked their alternatives such that the authentication mechanisms were preferred to the detection mechanisms.

Based on the results of the case studies, the security tradeoff analysis process should be modified so that added value of a technology is assessed, rather than its intrinsic value. In addition, future research should focus on specific instances where the security manager is actually comparing security technologies for implementation so that relevant cost information is available. Overall, the modified process would be more meaningful to the security manager and perhaps more insightful.

Final SAEM Rankings and Security Manager Rankings Compared with Expected Frequency Rankings

Table 8 - 8 Commercial Case Study

| Threat | SAEM Rank | Expected Frequency Rank | Final SM's Rank |
|-----------------------------|-----------|-------------------------|-----------------|
| Virus | 1 | 1 | 1 |
| Compromise | 2 | 4 | 3 |
| Password Guessing | 3 | 2 | 2 |
| Alteration | 4 | 6 | 5 |
| Denial of Service Attack | 5 | 8 | 7 |
| Internal Vandalism | 6 | 7 | 6 |
| System Scanning | 7 | 3 | 4 |
| Contamination | 8 | 14 | 10 |
| Distributed DoS | 9 | 21 | 16 |
| Electronic Graffiti | 10 | 19 | 9 |
| WEB Page Spoofing | 11 | 18 | 11 |
| Signal Interception | 12 | 5 | 8 |
| Theft | 13 | 22 | 15 |
| Vandalism | 14 | 15 | 12 |
| Compromising Emanations | 15 | 9 | 13 |
| Browsing | 16 | 12 | 14 |
| Procedural Violation | 17 | 11 | 18 |
| Data Entry Error | 18 | 16 | 17 |
| Trojan Horse | 19 | 10 | 19 |
| Fraud/Embezzlement | 20 | 25 | 23 |
| Personal Abuse | 21 | 13 | 21 |
| Password Nabbing | 22 | 17 | 20 |
| IP Spoofing | 23 | 20 | 22 |
| Cryptographic Compromise | 24 | 24 | 25 |
| Message Stream Modification | 25 | 27 | 27 |
| Trap Door | 26 | 23 | 24 |
| Logic Bomb | 27 | 26 | 26 |

Table 8 - 9 Hospital Case Study

| Threat | SAEM Rank | Expected Frequency Rank | Final SM's Rank |
|-----------------------------|-----------|-------------------------|-----------------|
| Virus | 1 | 3 | 1 |
| Alteration | 2 | 2 | 2 |
| Compromising Emanations | 3 | 1 | 4 |
| Compromise | 4 | 5 | 3 |
| System Scanning | 5 | 4 | 5 |
| IP Spoofing | 6 | 8 | 6 |
| Theft | 7 | 7 | 12 |
| Signal Interception | 8 | 6 | 7 |
| Vandalism | 9 | 9 | 8 |
| Denial of Service Attack | 10 | 10 | 10 |
| Cryptographic Compromise | 11 | 11 | 13 |
| Electronic Graffiti | 12 | 12 | 15 |
| WEB Page Spoofing | 13 | 13 | 14 |
| Fraud/Embezzlement | 14 | 14 | 9 |
| Message Stream Modification | 15 | 15 | 11 |

Table 8 - 10 Government Case Study

| Threat | SAEM Rank | Expected Frequency Rank | Final SM's Rank |
|-------------------------------|-----------|-------------------------|-----------------|
| Personal Abuse | 1 | 1 | 15 |
| Browsing | 2 | 2 | 11 |
| Procedural Violation | 3 | 3 | 10 |
| Virus | 4 | 9 | 12 |
| Compromise | 5 | 4 | 4 |
| Contamination | 6 | 5 | 2 |
| Alteration | 7 | 6 | 8 |
| Theft | 8 | 7 | 1 |
| Trojan Horse | 9 | 8 | 13 |
| Password Guessing | 10 | 10 | 5 |
| System Scanning | 11 | 11 | 16 |
| Data Entry Error | 12 | 14 | 14 |
| Fraud/Embezzlement | 13 | 17 | 3 |
| Trap Door | 14 | 16 | 18 |
| Password Nabbing | 15 | 15 | 17 |
| Signal Interception | 16 | 22 | 6 |
| Logic Bomb | 17 | 24 | 9 |
| Electronic Graffiti | 18 | 12 | 20 |
| Message Stream Modification | 19 | 23 | 7 |
| IP Spoofing | 20 | 13 | 19 |
| Cryptographic Compromise | 21 | 19 | 24 |
| WEB Page Spoofing | 22 | 18 | 23 |
| Distributed Denial of Service | 23 | 20 | 21 |
| Denial of Service Attack | 24 | 21 | 22 |

CHAPTER 9. Future Work and Observations

9.1 Introduction

This chapter discusses my observations concerning 1) an incremental SAEM that security managers can use to analyze alternative security technologies; 2) the ability of a security manager to use SAEM without the help of an analyst; 3) the importance of starting with a well defined set of threats; 4) the advantages of using questionnaires during the elicitation process; and 5) SAEM's value to security managers of various levels of experience and knowledge. In addition, this chapter discusses possible future work that would serve to improve the method and provide security managers with additional insight about their estimates.

9.2 Observations

9.2.1 *Observation 1: Point of Reference*

In each case study, the analyst asked the security managers to estimate the frequency and consequence values of a threat, independent of the managers' current security architecture. This was meant to allow the security manager the ability to evaluate their organizations' security architecture or the value of a security technology, to determine the most effective technologies for the security architecture. However, this exercise required security managers to imagine the system without existing countermeasures, which they found very difficult. The security manager's point of reference will affect the type of information that is elicited and the types of questions that the SAEM results will help answer.

All of the case study security managers found it difficult to make the threat and frequency estimates. As previously mentioned in Chapter 7, the government case study Director of Mission Assurance stated that the security manager's rankings were most likely based on current threats and countermeasures. In addition, the analyst needed to remind the other case study security managers several times during the elicitation interviews that the threat assessments should be made independent of the security architecture. Although there are advantages to this approach, some minor modification in the SAEM risk assessment and benefit analysis process could reduce the number of estimates required and allow security managers to simply evaluate a set of security technologies for the security architecture rather than evaluate the security architecture as a whole. This new approach may have been more satisfying to the government security manager.

Assessing the organization's threats and security technologies independent of the security architecture answers a different question than does assessing the organization's threats and evaluating security technology alternatives given the existing security architecture. The former assessment evaluates the current security architecture and prioritizes all possible security technologies. The later assessment—evaluating security technologies given the existing security architecture—does not evaluate the current architecture but, rather, attempts to help the security manager determine the next best set of security technologies for further reducing the organization's security risks. For some security managers, such an incremental SAEM may be

more useful--and less tedious--in helping determine how to allocate, for example, next year's security budget without going through a complete assessment.

If the security manager wanted to use SAEM to help assess a set of potential additional security technologies given the current architecture, then he or she would need to modify the risk assessment and benefit analysis process. First, the security manager would determine and rank the organization's outstanding threats, which may not include threats for which the security manager had achieved an acceptable level of risk control. In addition, the security manager's frequency and outcome estimates would reflect the effectiveness of the security architecture. For example, one of the case study security managers estimated that he would most likely see two or three viruses each week, however, if the security architecture did not have antivirus software then he estimated that the organization would experience two to three virus attacks daily. Thus in the incremental risk assessment process SAEM would rank the *Viruses* much lower.

Finally, in the benefit analysis phase, the security manager would identify and estimate the effectiveness of risk-mitigating security technologies for the threats that were selected in the risk assessment phase. The results of the risk assessment and benefit analysis phases would result in a prioritized, but abbreviated, list of threats and countermeasures based on the security manager's assessments.

The greatest advantage in assessing the threats and security technologies with respect to the security architecture is that the number of estimates required could be far fewer—potentially with fewer errors because the process is less tedious. (Recall that the government security manager made over 500 security technology effectiveness estimates.) In addition, fewer estimates might improve the assessment process because the security manager would have additional time to more thoroughly consider the estimates. Furthermore, the security manager might be able to rely on empirical data to make the estimates because the security manager doesn't have to imagine the system without its security architecture.

Of course, the greatest disadvantage of this incremental SAEM process is that the security manager may overlook a security technology that could greatly benefit the organization. In addition, the security manager could not evaluate the additional risk mitigation that a new technology would provide to the existing threats, but the assumption in this process is that these threats are already at an acceptable level of risk. Conversations with the case study security managers indicated that they would always want to assess their security architectures, but an incremental SAEM process would be very useful subsequent to an initial and complete assessment.

9.2.2 Observation 2: Doing Without the SAEM Analyst

Although the analyst played a key role in the case studies, could security managers successfully use SAEM without the analyst? In the government case study, the Director of Mission Assurance called for automation of SAEM processes, indicating the need for security managers to use SAEM without the analyst. The analyst's key functions during SAEM were to facilitate the elicitation interviews and present and interpret the results. In addition, the analyst answered questions about security technologies and threats during some of the case studies. Although ASESS can quickly produce the risk assessment and technology rankings, security managers would need additional training to interpret the results independent of an analyst.

Since insight into the security manager's decision process is the goal of SAEM, the analyst's critical task was to show the security manager inconsistencies between the initial rankings and the SAEM rankings. The analyst guided the security managers through the refinement process until they established their final rankings. The analyst reviewed the results and highlighted inconsistencies among the rankings. With a few hours of training, each of the security managers could learn to interpret the results and explore the data to determine the source of inconsistent results. However, the danger is that security managers might not be as objective as an outside analyst and the manager could tweak the data to ensure preconceived threat rankings.

In addition to training, security managers would need automated support, such as the ASESS tool, to determine their threat and security technology rankings and analyze their assumptions. The security managers could easily use ASESS to quickly conduct what-if analysis, gaining additional insight into their organizations' prioritizations. Although ASESS is not required in order to prioritize the threats and technologies, the process is extremely tedious and time consuming without it. Automation facilitates what-if analysis, which provides the security managers with important insight. Therefore, security managers could successfully use SAEM without relying on the analyst if they received training and had automated support.

9.2.3 Observation 3: The Initial Set of Threats

In each of the case studies, the analyst began the risk assessment and benefit analysis with an initial set of threats. SAEM does not help the security manager identify threats, but the quality of the results depends on starting with the "right" list of threats for an organization³⁶. All of the case study security managers indicated that they would modify the threat list, and the government security manager stated (on his satisfaction survey) that the initial list needed to be improved. The security managers expressed their dissatisfaction with the threat list despite the analyst's encouragement that the list be modified to represent the organization's threats, but the managers made few modifications to the analyst's initial list.

Ideally, the analyst would prefer that the security manager identify the organization's threats before starting the SAEM risk assessment phase, but none of the case study organizations had completed a risk analysis prior to the SAEM analysis. Instead of articulating their specific lists, each security manager adopted almost all of the threats that the analyst presented in the SAEM risk assessment as a risk to the organization. After completing the SAEM process, the security managers felt that some of the threats were too general or overlapped with other threats. Security managers struggled with estimating frequency and consequence values if the threat was too broadly defined. Therefore, the analyst and the security manager should take time up front to develop a specific list of threats that minimizes overlap with other threats.

Ideally, participants should start with an industry-specific common set of threats—but none exist. Although there are some threats that could be considered common to nearly every organization, the definitions and categories still must be tailored for SAEM. For example, several threat lists include *Information Theft* as a risk, but the risk-mitigation controls are different depending on the source of the attack. Risk mitigation controls, such as firewalls and email content inspection software helps mitigate Theft risks originating externally, but don't

³⁶ Recall from Chapter 2, that the SAEM risk assessment phase supports, but does not supplant, the organization's overall risk analysis process, which is the first step in the organization's risk management process.

help much with the risk of *Theft* from employees, who are internal to the organization. Security managers would select different risk-mitigation controls for *Internal Theft*. Therefore as a minimum, in developing the organization's threat list for SAEM, the security manager should separate internal and external threats if the risk-mitigation controls are likely to be different.

9.2.4 Observation 4: Elicitation Interviews or Questionnaires?

In the hospital and government case studies, the analyst used a combination of elicitation interviews and questionnaires. During the risk assessment, the analyst interviewed the security managers to determine the initial threat prioritizations and attack consequences. After the interview, the analyst developed a questionnaire to capture the frequency and consequence values. During the benefit analysis, the analyst interviewed the security managers to identify the risk mitigating technologies for each of the threats, and again developed a questionnaire to capture the effectiveness estimates. In contrast, in the commercial case study the analyst obtained all of the security manager's estimates through elicitation interviews. Although there are advantages to both elicitation methods, the interview/questionnaire method appeared to provide the most consistent results.

There are several disadvantages in eliciting estimates using only interviews. One disadvantage to interviews is that they are tedious and draining for the participants. The analyst observed that the participants were more engaged in the beginning of the interviews and their estimates were not as deliberate towards the end of the interview. In addition, the analyst detected several inconsistencies during the interview process, whereas there were fewer corrections made to the estimates made using the questionnaires. Perhaps the analyst found fewer inconsistencies on the questionnaires because they allowed security managers to continually compare previous answers, while interview participants appeared more reluctant to review previous estimates. In addition, questionnaires allow the security managers to make their estimates without time constraints, which may make for more deliberate estimates.

Although the interviews are tedious, the greatest advantage to using interviews is that the analyst can often capture the rationale for the estimates, which helps the security manager through the refinement process. When the security manager has little experience or empirical data on which to base estimates, he or she developed plausible scenarios that were used to justify their best estimates. For example, the commercial case study manager made several estimates about *Personal Computer Abuse* based on the percentage of employees she thought were using their computers in violation of company policies for any given hour. Later, during the risk assessment refinement process, when the security manager questioned the *Personal Computer Abuse* estimates the analyst reminded the security manager how she had derived this estimate. Since the security managers make so many estimates, they easily lose track of how they decided on some of the estimates. Therefore, if security managers wanted to use SAEM without a trained analyst, then ASESS should be modified to allow managers to capture rationale for their estimates. Rationale capture could be automated or the security manager could turn it off if it became tedious.

9.2.5 Observation 5: Security Managers Required Level of Security Expertise and Knowledge

Previous chapters in this thesis mentioned that the case study security managers appeared to have significantly different levels of expertise and knowledge about their organizations' threats and security technologies. Since the SAEM process depends on the ability of security managers to make estimates on the best available information, is it realistic to expect SAEM to be of value to security managers regardless of their level of expertise and knowledge? Although all case study security managers found the process insightful, executive-level decision makers who rely on the results may need to assess the level of expertise from which the estimates were made. Moreover, the degree of experience will be important to decisions on whether to use an analyst in the SAEM process and, once it is established, how much to rely on larger sample measurements of managers' assessments.

9.2.5.1 SAEM for the less experienced security manager

Clearly, SAEM can help less experienced climb the security learning curve, but their success with SAEM will likely depend on the presence and input of the analyst. Their experience will be improved once an industry database is compiled to provide baseline measurements against which the less experienced manager can compare his or her results. Therefore, SAEM will not be as valuable to less experienced security managers as it is to more experienced security managers until SAEM establishes a baseline database or additional research provides industry-specific threat and security technology effectiveness data. More importantly, SAEM would not be an appropriate method for less experienced security managers without a trained analyst.

9.2.5.2 SAEM for the moderately experienced security manager

SAEM appeared to have had the greatest value to the moderately experienced commercial case study security manager. This type of security manager understands his or her organizations' threats and most of the security technologies, and can leverage SAEM to analyze their security architectures and communicate the results to their information system executives. In addition, the analysis affirms the managers' known security architecture weaknesses in their information systems, but gives them insight into alternative risk mitigation strategies. For example, the commercial case study security manager identified additional technologies that might help reduce the organization's *Virus* threat as a result of the SAEM process. Moderately experienced security managers would also greatly benefit from a baseline database and additional research about the effectiveness of security technologies.

9.2.5.3 SAEM for the very experienced security manager

SAEM appeared to have the least value to the highly experienced government case-study security manager. This type of manager is very familiar with virtually all of the threats and risk-mitigation strategies. Still SAEM provides the very experienced security manager with an objective tool with which to frame or rationalize their decisions and communicate these decisions to senior information technology executives.

9.3 Future Work

This section describes possible future work that could significantly contribute to the value that SAEM provides to security managers.

9.3.1 *Effectiveness Estimates*

The two greatest areas in need of additional research are: 1) to determine the real rather than estimated effectiveness of risk mitigation technologies; and 2) how organizational factors, such as susceptibility to attack and security policies affect a technology's effectiveness.

The benefit analysis process elicits effectiveness estimates from the security managers, but these estimates are based on their experience and intuition--not empirical evidence. Chapter 8 showed that the hospital effectiveness estimates were significantly higher than the other case study managers' effectiveness estimates, but without hospital industry comparisons one cannot determine whether they are *significantly* inconsistent with other hospital security manager estimates.

Security managers may gain additional insight when their results are compared with the results of other security managers. For example, the commercial case study security manager identified *Viruses* as her organization's number one threat, but only identified a few risk-mitigation technologies. In the commercial case study final report, the analyst was able to identify additional risk mitigating technologies for *Viruses* that the other security managers had identified. Although only the commercial case study manager could determine whether these other technologies were appropriate for her organization, the manager found it very useful to receive this type of feedback. Furthermore, every case study security manager inquired about how their estimates and results compared to others'.

In addition to making industry-specific comparisons, the hospital security manager needed to assess how his hospital compared with other hospitals so that he could have adjusted his own estimates based on the organizational differences among hospitals. For example, large city hospitals may be more susceptible to security compromises than are small suburban hospitals, so the effectiveness of security technologies may be very different between the two types of hospitals. Future research should determine how organizational factors affect the effectiveness of a security technology and establish industry-specific effectiveness baselines.

9.3.2 *Threats and Threat Frequencies and Outcomes*

In addition to developing a baseline of security technology effectiveness metrics, security managers would greatly benefit from an established set of threats--especially threats relevant to their industry. Of the three case studies, only the government case study security manager had collected information about threats and frequency of attacks, but not for all threats. The security research community has made several attempts to establish a taxonomy of vulnerabilities (Lanwehr, Bull et al. 1994; Corporation 2003), but the community lacks uniform definitions of threats³⁷.

³⁷ Vulnerabilities are flaws or defects in software or designs and threats are potential events that take advantage of vulnerabilities.

Establishing a threat taxonomy is only the beginning. Few organizations have attempted to quantify the effects of an attack because the consequences of an attack are sometimes hard to capture. For example, an organization may not be able to measure their public embarrassment damage from a security compromise. Some limited attempts (Larsen 1999; Malphrus 2000; King, Dalton et al. 2001) have been made to estimate the damage that various *Viruses* have inflicted upon organizations, but attempts to quantify damage is rare. Although, the government case study organization had an incident response center but did not attempt to estimate damage from attacks.

Optimal SAEM results greatly depend on the best information available. As future research establishes industry-specific information, security managers will be able to use this information to make informed decisions.

9.3.3 Security Procedures

Security procedures are essential elements of an organization's security architecture, but this thesis included only risk mitigating security technologies. Frequently security managers must establish operational procedures in addition to selecting a security technology to mitigate the organization's risk from a threat. For example, Anderson (Anderson 2001) describes bank security managers who established Automated Teller Machine (ATM) procedures that kept bank employees from having simultaneous access to a customer's ATM card and Personal Identification Number. When the bank violated these procedures, an employee took advantage of the opportunity to steal from customer accounts. Therefore, the security manager cannot completely evaluate the organization's security architecture without including security procedures as potential risk mitigation strategies.

Future research should expand SAEM to include security procedures in the benefit analysis phase. Security managers would need to identify specific security procedures that mitigate threats and estimate the effectiveness of these procedures. Since security procedures tend to differ among organizations, the analyst may not be able to prepare a set of common security procedures, but additional research may determine that a common set of security procedures exist, especially for similar industries. Including security procedures in the benefit analysis phase may significantly expand the security architecture design choices for the security manager; however, security managers may find it difficult to establish cost estimates for security procedures.

9.3.4 More than One Security Technology

The benefit analysis phase evaluated each security technology independent of all the others. In actuality, the effectiveness of a security technology often depends on other technologies. Furthermore, the overall contribution that a security technology makes in mitigating the organization's risk from a threat depends on the existing security architecture. For example, a security manager may have estimated that proxy firewalls are 80% effective in reducing IP spoofing attacks. However, if the security architecture includes a hardened operating system then the effectiveness of the proxy firewall is less because many of the attacks are already mitigated by the hardened operating system.

Security managers could adjust SAEM to include the effect that multiple security technologies have in estimating the effectiveness of a single security technology, but future

research should include this information as part of the effectiveness metrics of security technologies. Until additional research develops more reliable effectiveness metrics, security managers can treat combinations of security technologies as one when using SAEM.

In addition to evaluating the reduced effectiveness that one technology when another technology is present, a security manager might also be aware of conflicts between security technologies. For example, sometimes encryption interferes with firewall effectiveness. These conflicts between security technologies would affect the manager's decision and priorities. The benefit analysis and coverage model should be extended to ensure that security technology conflicts are identified and revealed during the SAEM process.

APPENDIX A

Risk Assessment Data Collection

At our first meeting, you identified several threats. For each threat, I would like you to estimate how often an attack of each type of threat is expected. You should estimate the expected frequency based on what you would expect given that there are no security mechanisms in place to stop or deter an attack. You can provide your estimate in any time units such as hours, days, weeks, months, or years. If there isn't any data to support your estimate, please give your best guess based on your experience as to what occurs or would occur if the current security mechanisms were not in place.

The following statement is an example of how you would think about the frequency:

“I would expect to see ____5____ denial of service attacks each month (year, day, week, etc.).”

In addition to the expected frequency of an attack, please provide an upper and lower bound. The lower bound should fill in the blank for the following type of statement:

“I would be surprised if I saw fewer than _____ denial of service attacks in a given month (week, day, year, #__ years, hour etc.).”

or

I would be surprised if I saw fewer than _____ denial of service attacks every 2 years.

The upper bound should fill in the blank for the following type of statement:

“I would be surprised if I saw more than _____ denial of service attacks in a given year (week, day, month, hour, #__ years, etc.).”

If you decide that there is a threat that you need to add to ensure that all IT risks to your organization are captured, then please add the threat with a brief description, so that I know what you mean by the new threat, and its associated data.

If you have any questions, please call

Attack Frequencies

Theft - The unauthorized taking of information for personal gain

Lower Bound _____ per _____ [hour, day, week, month, year, # ___of years]

Expected _____per _____ [hour, day, week, month, year, # ___of years]

Upper Bound _____per _____ [hour, day, week, month, year, # ___of years]

Contamination - The intermixing of data of different sensitivity levels.

Lower Bound _____ per _____ [hour, day, week, month, year, # ___of years]

Expected _____per _____ [hour, day, week, month, year, # ___of years]

Upper Bound _____per _____ [hour, day, week, month, year, # ___of years]

- 1) **Compromise** - The unintentional release of information to someone not authorized access to the information. This includes information exempt from public disclosure, Privacy Act information, proprietary information, sensitive-unclassified information, and national security information.

Lower Bound _____ per _____ [hour, day, week, month, year, # ___of years]

Expected _____per _____ [hour, day, week, month, year, # ___of years]

Upper Bound _____per _____ [hour, day, week, month, year, # ___of years]

- 2) **Password Guessing** - An automated or manual attempt to obtain user or system privileges by guessing the password

Lower Bound _____ per _____ [hour, day, week, month, year, # ___of years]

Expected _____per _____ [hour, day, week, month, year, # ___of years]

Upper Bound _____per _____ [hour, day, week, month, year, # ___of years]

Threat Outcomes

In addition to the attacks, you also identified outcomes that most concern your organization. Although, confidentiality, data integrity and system availability are important, these ideas do not specifically capture how a security compromise impacts your organization. From our initial session, it appears that there are three major concerns to your organization: 1) Damage to Public Image, 2) Damage to Customer Relations that could arise if sensitive information were compromised, and 3) Lost Revenue.

Any given successful attack could result in one or more of these consequences occurring, or have no consequence at all. For example, an attack may have not damaged public, damage customer relationships, or result in any lost revenues. For each attack (there is a separate sheet for each attack), please indicate your expected outcome for public image, customer relationships, and lost revenue. The expected outcomes reflect what you are likely to see given that the attack is successful. A scale from 1 to 7 is used to assess the severity of the first two outcomes. You can fill in the outcome value that most represents your expectations about the severity of the outcome. I would also like you to provide upper and lower bounds.

Example:

Denial of Service Attack:

Public Image:

The most likely impact to public image would be 3. (most likely)

I would be surprised if the impact were as little as 1. (lower bound)

I would be surprised if the impact were as much as 5. (upper bound)

| | | | | | | |
|------|------|--------------------|----------|----------------------|--------|----------------|
| 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| None | Mild | Moderately Mild | Moderate | Moderately Severe | Severe | Most Severe |

You may find that the definitions are not quite as you would define the attack. Feel free to adjust the definitions to reflect your organization's understanding of the risk. Some of these attacks are broadly defined, if you are not sure about a definition, whether a specific type of attack is included or not, decide whether you want to include it, or develop a new threat. It is important to be consistent through this process and it is also important that your organization's risks are appropriately captured. If you have any questions, please contact me.

Theft

Definition - The unauthorized taking of information for personal gain.

Public Image:

The most likely impact to public image would be _____.

I would be surprised if the impact were as little as _____.

I would be surprised if the impact were as much as _____.

| | | | | | | |
|------|------|--------------------|----------|----------------------|--------|----------------|
| 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| None | Mild | Moderately Mild | Moderate | Moderately Severe | Severe | Most Severe |

Customer Relationships:

The most likely impact to customer relationships would be _____.

I would be somewhat surprised if an attack resulted in no more than a _____ (low) impact in customer relationships.

I would be somewhat surprised if the impact were as much as _____ (high) to our customer relationships.

| | | | | | | |
|------|------|--------------------|----------|----------------------|--------|----------------|
| 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| None | Mild | Moderately Mild | Moderate | Moderately Severe | Severe | Most Severe |

Lost Revenue:

Most likely an attack would result in \$\$ _____ in lost revenue

I would be somewhat surprised if an attack resulted in as little as \$\$ _____ lost revenue

I would be somewhat surprised an attack resulted in as much as \$\$ _____ lost revenue

| | | | | | | |
|------|------|--------------------|----------|----------------------|--------|----------------|
| 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| None | Mild | Moderately Mild | Moderate | Moderately Severe | Severe | Most Severe |

Contamination

Definition - The intermixing of data of different sensitivity levels.

Public Image:

The most likely impact to public image would be_____.

I would be surprised if the impact were as little as _____.

I would be surprised if the impact were as much as_____.

| | | | | | | |
|------|------|--------------------|----------|----------------------|--------|----------------|
| 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| None | Mild | Moderately Mild | Moderate | Moderately Severe | Severe | Most Severe |

Customer Relationships:

The most likely impact to customer relationships would be _____.

I would be somewhat surprised if an attack resulted in no more than a_____ (low) impact in customer relationships.

I would be somewhat surprised if the impact were as much as _____ (high) to our customer relationships.

| | | | | | | |
|------|------|--------------------|----------|----------------------|--------|----------------|
| 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| None | Mild | Moderately Mild | Moderate | Moderately Severe | Severe | Most Severe |

Lost Revenue:

Most likely an attack would result in \$\$ _____ in lost revenue

I would be somewhat surprised if an attack resulted in as little as \$\$ _____ lost revenue

I would be somewhat surprised an attack resulted in as much as \$\$ _____ lost revenue

| | | | | | | |
|------|------|--------------------|----------|----------------------|--------|----------------|
| 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| None | Mild | Moderately Mild | Moderate | Moderately Severe | Severe | Most Severe |

Compromise

Description - The unintentional release of information to someone not authorized access to the information. This includes information exempt from public disclosure, Privacy Act information, proprietary information, sensitive-unclassified information, and national security information.

Public Image:

The most likely impact to public image would be _____.

I would be surprised if the impact were as little as _____.

I would be surprised if the impact were as much as _____.

| | | | | | | |
|------|------|--------------------|----------|----------------------|--------|----------------|
| 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| None | Mild | Moderately Mild | Moderate | Moderately Severe | Severe | Most Severe |

Customer Relationships:

The most likely impact to customer relationships would be _____.

I would be somewhat surprised if an attack resulted in no more than a _____ (low) impact in customer relationships.

I would be somewhat surprised if the impact were as much as _____ (high) to our customer relationships.

| | | | | | | |
|------|------|--------------------|----------|----------------------|--------|----------------|
| 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| None | Mild | Moderately Mild | Moderate | Moderately Severe | Severe | Most Severe |

Lost Revenue:

Most likely an attack would result in \$\$ _____ in lost revenue

I would be somewhat surprised if an attack resulted in as little as \$\$ _____ lost revenue

I would be somewhat surprised an attack resulted in as much as \$\$ _____ lost revenue

| | | | | | | |
|------|------|--------------------|----------|----------------------|--------|----------------|
| 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| None | Mild | Moderately Mild | Moderate | Moderately Severe | Severe | Most Severe |

Password Guessing

Definition - An automated or manual attempt to obtain user or system privileges by guessing the password

Public Image:

The most likely impact to public image would be_____.

I would be surprised if the impact were as little as _____.

I would be surprised if the impact were as much as_____.

| | | | | | | |
|------|------|--------------------|----------|----------------------|--------|----------------|
| 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| None | Mild | Moderately Mild | Moderate | Moderately Severe | Severe | Most Severe |

Customer Relationships:

The most likely impact to customer relationships would be _____.

I would be somewhat surprised if an attack resulted in no more than a_____ (low) impact in customer relationships.

I would be somewhat surprised if the impact were as much as _____ (high) to our customer relationships.

| | | | | | | |
|------|------|--------------------|----------|----------------------|--------|----------------|
| 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| None | Mild | Moderately Mild | Moderate | Moderately Severe | Severe | Most Severe |

Lost Revenue:

Most likely an attack would result in \$\$ _____ in lost revenue

I would be somewhat surprised if an attack resulted in as little as \$\$ _____ lost revenue

I would be somewhat surprised an attack resulted in as much as \$\$ _____ lost revenue

| | | | | | | |
|------|------|--------------------|----------|----------------------|--------|----------------|
| 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| None | Mild | Moderately Mild | Moderate | Moderately Severe | Severe | Most Severe |

REFERENCES

- Alberts, C. and A. Dorofee (2001). OCTAVEsm Threat Profiles. Pittsburgh, Software Engineering Institute.
- Anderson, R. (2001). Security Engineering: A Guide to Building Dependable Distributed Systems, Wiley Computer Publishing.
- Bishop, M. (2003). Computer Security: Art and Science, Addison Wesley.
- Borcherding, K. and T. Eppel (1991). "Comparison of Weighting Judgments in Multiattribute Utility Measurement." Management Science **37**(12): 1603-1619.
- Corporation, T. M. (2003). Common Vulnerabilities and Exposures, The Mitre Corporation.
- Finne, T. (1998). "A Conceptual Framework for Information Security Management." Computers & Security **17**(4): 303-307.
- Finne, T. (1998). "The Three Categories of Decision-Making and Information Security." Computers & Security **17**: 397-405.
- Fischer, G. W. (1995). "Range Sensitivity of Attribute Weights in Multiattribute Value Models." Organizational Behavior and Human Performance **62**(3): 252-266.
- Fraser, B. (1997). Site Security Handbook RFC 2196.
- Gordon, L. A. and M. P. Loeb (2002). "The Economics of Information Security Investment." ACM Transactions on Information System Security **5**(4): 438-457.
- Harris, S. (2002). CISSP Certification Exam Guide, McGraw-Hill/Osborne.
- Jacobson, R. V. (2002). Computer Security Handbook. Fourth Edition, John Wiley & Sons, Inc.
- Kazman, R., J. Asundi, et al. (2000). Quantifying the Costs and Benefits of Architectural Decisions. International Conference on Software Engineering-22, IEEE.
- Keeney, R. L. and H. Raiffa (1999). Decisions with Multiple Objectives. New York, Cambridge University Press.
- King, C. M., C. E. Dalton, et al. (2001). Security Architecture Design: Deployment & Operations, McGraw-Hill RSA Press.
- Knutz, R. and R. V. Vines (2001). CISSP Prep Guide, John Wiley & Sons, Inc.
- Kontio, J. (1996). A Case Study in Applying a Systematic Method for COTS Selection. International Conference on Software Engineering -18.
- Lane, T. (1990). A Design Space and Design Rules for User Interface Software Architecture. Pittsburgh, Carnegie Mellon University.
- Lanwehr, C. E., A. R. Bull, et al. (1994). "A Taxonomy of Computer Program Security Flaws." ACM Computing Surveys **26**(3).
- Larsen, A. K. (1999). "Global Security Survey Virus Attack." Information Week.
- Malphrus, S. R. (2000). The "I Love You" Computer Virus and Financial Services Industry. Subcommittee on Financial Institutions of the Committee on Banking. Washington D.C.
- Meadows, C. (2000). "A Cost-based Framework for Analysis of Denial of Service in Networks." Journal Computer Security **9**(5): 143-164.
- Millen, J. (1992). A Resource Allocation Model for Denial of Service. IEEE Symposium on Security and Privacy, IEEE Computer Society Press.
- NCSC (1985). Department of Defense Trusted Computer System Evaluation, National Computer Security Center, Department of Defense.
- Ramachandran, J. (2002). Designing Security Architecture Solutions, John Wiley and Sons Inc.
- Saaty, T. L. (1990). The Analytic Hierarchy Process. New York, McGraw-Hill.
- SEI (2003). OCTAVEsm Information Security Risk Evaluation, Software Engineering Institute.

- Stillwell, W. G. and D. A. Seaver (1981). "A Comparison of Weight Approximation Techniques in Multi-attribute Utility Decision-making." Organizational Behavior and Human Performance **28**: 62-78.
- Stoneburner, G. (2001). Underlying Technical Models for Information Technology Security. Washington D.C., National Institute of Standards and Technology.
- Stoneburner, G., A. Goguen, et al. (2001). Risk Management Guide for Information Technology Systems. Washington D.C., National Institute of Standards and Technology.
- Stoneburner, G., C. Hayden, et al. (2001). Engineering Principles for Information Technology Security (A Baseline for Achieving Security), National Institute of Standards and Technology.
- Straub, D. and R. Welke (1998). "Coping with Systems Risk: Security Planning Model Management Decision Making." Management Information System Quarterly **23**(4): 441-469.
- Von Winterfeldt, D. and W. Edwards (1986). Decision Analysis and Behavioral Research. New York, Cambridge University Press.
- Yoon, K. P. and C.-L. Hwang (1995). Multiple Attribute Decision Making: An Introduction, Sage Publications.