# Security Attribute Evaluation Method: A Cost-Benefit Approach

Shawn A. Butler
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213
412-268-8101
shawnb@cs.cmu.edu

## ABSTRACT

Conducting cost-benefit analyses of architectural attributes such as security has always been difficult, because the benefits are difficult to assess. Specialists usually make security decisions, but program managers are left wondering whether their investment in security is well spent. This paper summarizes the results of using a cost-benefit analysis method called SAEM to compare alternative security designs in a financial and accounting information system. The case study presented in this paper starts with a multi-attribute risk assessment that results in a prioritized list of risks. Security specialists estimate countermeasure benefits and how the organization's risks are reduced. Using SAEM, security design alternatives are compared with the organization's current selection of security technologies to see if a more cost-effective solution is possible. The goal of using SAEM is to help information-system stakeholders decide whether their security investment is consistent with the expected risks.

## 1 INTRODUCTION

Security managers have trouble evaluating alternative security designs and justifying security technology investments because the technology benefits are difficult to estimate. Security technology benefits depend on how often an attack is expected, how much damage is likely to occur and how effective the security technology is in mitigating the damage from an attack. Cost-benefit analyses can bridge the communication gap between security managers and information technology (IT) managers. IT managers want to know that their investment in security has reduced risks to an acceptable level and security managers want to be sure that their designs are the most secure. The advantages of a structured cost-benefit analysis are:

- Security managers make their assumptions explicit and capture decision rationale.
- Sensitivity analysis shows how assumptions affect design decisions.
- Design decisions are re-evaluated consistently when assumptions change.
- IT managers see whether investment is consistent with risk expectations.

IT managers are motivated to minimize security costs but maximize security benefits. Comparing costs among alternative security architectures is significantly easier than comparing benefits, since proven financial analysis tools can more precisely estimate costs. In contrast, benefits are based on uncertain events and imperfect knowledge. Although no one can accurately predict how often an attack will occur and how effectively the security will mitigate the damage, experienced security managers intuitively, and implicitly, estimate the risk and the effectiveness of their risk mitigation strategies. The key to security cost-benefit analyses is to make these intuitions explicit.

In addition to the difficulties in estimating risk, comparing alternative designs is challenging because the strength of the design depends on a relaxed adherence to security engineering design principles. For example, security managers usually adopt a *precautionary* approach to security engineering. Security architectures that have at least one mitigation strategy for each risk are usually preferred to those that leave gaps for rarely expected attacks. Simplistic cost-benefit analyses will have difficulty showing economic justification for rare events.

### 1.1 Security Attribute Evaluation Method

This paper presents the Security Attribute Evaluation Method (SAEM) and the results of using this method to compare alternative security designs for a non-profit organization's financial and accounting system. SAEM is a cost-benefit analysis process for analyzing security design decisions that involves four steps: 1) a security technology benefit assessment, 2) an evaluation of the effect of security technologies in mitigating risks, 3) a coverage assessment and 4) a cost analysis. Steps 3 and 4 can be done in parallel.

SAEM relies on a quantitative risk and benefit assessment, in which an analyst or investigator conducts structured interviews of IT and security managers to elicit the initial data. The organization carefully reviews the results of each step before continuing to the next step. If the results do not appear to represent the managers' concerns or experience, then the managers can revise the original input data and assumptions.

Risk mitigation strategies include procedures and security technologies. As presented in this paper, SAEM only evaluates the technology design choices, not risk mitigating procedures. Though SAEM could be expanded to include procedures as part of the assessment, these procedures tend to be organization

specific and SAEM was initially designed to be used across organizations.

Currently, I am developing SAEM in collaboration with security managers of real systems. In addition to the financial and accounting system, I am also using SAEM to evaluate a government e-commerce system and a Department of Defense web-based logistics system.

## 1.2 Roadmap

Section 2 assesses the current practice of security architecture development. Section 3 lays the foundation for the Security Attribute Evaluation Method with the multi-attribute risk assessment. The section presents a brief description of the multi-attribute risk assessment method, which was used to prioritize the organization's risks based on the security manager's estimate of attack frequencies and potential outcomes. Section 4 describes SAEM and presents an example of how SAEM could be used to select the most cost-effective security technology from among three alternatives. Section 5 shows the results of conducting sensitivity analysis on the initial results and Section 6 reports on the feedback from case study clients. Finally, Section 7 discusses future work.

## 1.3 Considerations

SAEM depends upon a risk assessment and an initial assessment of security technology effectiveness. These assessments depend upon several assumptions, such as:

- That the organization has established security policies and procedures that are sufficient and robust for the business operation;
- That security products have been correctly installed, configured and maintained;
- That attacks result in predictable outcomes and variances.

These assessments will vary among security managers based on their experience, expertise, and information system environment [7]. SAEM provides a framework to analyze how their assumptions affect design decisions and a mechanism to determine how sensitive these decisions are to the assumptions. Although one may argue with the estimations presented in this paper, they do represent an experienced security manager's best estimations.

# 2 BACKGROUND

Software engineers consider security an attribute of the system architecture, but security managers often refer to a system's collection of security policies, procedures and technologies as a *security architecture* [5]. Security architecture development follows three phases (See Figure 1). Development begins with a risk assessment, which analyzes the expected threats and outcomes to produce a prioritized list of threats. Analysis and selection of risk mitigation strategies or countermeasures follows the risk assessment. System design constraints and security requirements (in the form of policies, regulations, and requirements) guide the selection of countermeasures. Finally, security managers integrate the selected countermeasures into the system and configure the security technologies to enforce organizational security policies. The security architecture development process is revisited periodically to ensure that the architecture stays current.

In practice, the process is complicated because risk assessment data is based on the security manager's subjective and qualitative estimate of each threat and the anticipated risk reduction from the countermeasures. In addition, security managers lack the tools to systematically evaluate alternatives and justify security expenditures.

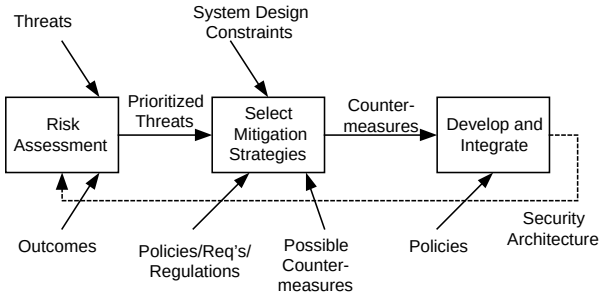


**Figure 1.** SecurityArchitecture Development Process

## 2.1 State of the Practice

Ideally, a risk assessment should be based on statistical threat data which captures how often attacks occur and their outcomes. The security architecture should be the most effective given cost and system design constraints. In the author's 20 years of experience in system development, risk assessments, if done at all, rarely contain sufficient information from which to conduct meaningful cost-benefit analyses. Security managers lack the tools to transform qualitative and subjective assessments into credible results.

Security managers could use simplistic equations, such as *Risk = Cost * Threat*, if threat data were available and all costs could be described in economic terms. Unfortunately, for many reasons, statistical threat data is rarely available. In addition, countermeasures should impact either outcomes or threat frequencies, but the effectiveness, or *benefit* of security technologies, is difficult to estimate without considering an organization's information system environment.

Without firm statistical data, security managers rely on experience, judgment, and the best available knowledge to determine which attacks are most likely and what damage will probably occur if the attack is successful. In addition, during discussions with system security managers in the case studies, it was clear that they had some sense, though largely just intuitive, of the effectiveness of the security technologies used in their systems.

## 2.2 Investment Priorities

Although software engineers and IT managers leave security decisions to security specialists, they often wonder whether the investment in security is well spent. Most importantly, managers want to be sure that the allocation of security resources reduce risk to an acceptable level. The difficult questions are not whether basic security mechanisms, such as access control and perimeter firewalls, should be part of the security design, rather which type of mechanism is most cost-effective given the expected risks. For example, there are three different types of firewalls: packet filter, circuit gateways, and application gateways. In addition to packet filtering, constraint gateways perform intrusion detection,

encryption, URL filtering, etc. Constraint gateways are significantly more expensive than packet filtering firewalls. Is the additional functionality cost effective?

## 2.3 The Case Study

This paper illustrates SAEM using a financial and accounting system. This system is part of a larger network. I interviewed the security manager and information system department managers for the risk assessment data. The security manager provided a previous risk assessment, but it lacked specific frequency and outcome data. In the SAEM risk analysis, the managers identified twenty-eight risks, each having four possible outcomes. Although carefully structured interviews are designed to elicit the best possible estimations, the managers reviewed the results and clarified or adjusted their answers when necessary.

After the risk assessment, the analyst interviews the security manager for the security benefit assessment information. This benefit assessment included over 40 security technologies. Due to article length considerations, the financial and accounting risk and benefit assessment data presented in this paper represents only a small portion of data collected and analyzed.

# 3 RISK ASSESSMENT

The purpose of the risk assessment is to identify threats and the consequences of successful attacks so that security managers can prioritize security resources. If an organization does not have a risk assessment that contains attack and outcome expectations, then an analyst conducts a multi-attribute risk assessment, which was developed as part of this research to provide a prioritized list of threats.

As a precursor to SAEM, I conducted a multi-attribute risk assessment for the case study. A multi-attribute risk assessment results in a *threat index* (TI) for each risk based on the expected frequency of an attack and probable outcomes from a successful attack. A complete description of a multi-attribute risk assessment process is described in [2]. This section provides a brief summary of the risk assessment's key elements and selected case study results. SAEM uses data from the risk assessment to determine the security technology benefits.

## 3.1 Multi-Attribute Risk Assessments

Multi-attribute analysis provides a convenient framework for conducting risk assessments. Traditionally used in the Decision Sciences, multi-attribute analysis is used to systematically evaluate decision alternatives when the decision outcomes are uncertain [8]. Multi-attribute methods used in risk assessments result in a *threat index* for each risk based on estimations of threat frequencies and expected outcomes. An outcome can have several consequences. For example, an attack could result in lost revenue, public embarrassment, and regulatory penalties. These consequences are called *attributes* in multi-attribute analysis. Therefore, an outcome is a vector of attributes where the value of the attribute is the level of damage.

The multi-attribute risk assessment consists of four steps. In the first step, the security and IT managers identify the outcome attributes. In the case study, the IT and security managers identified four significant outcome attributes: Lost Productivity, Lost Revenue, Reputation, and Regulatory Penalties[1]. Next, the security manger identifies the frequency and outcome attribute values for each threat. After the security manager identifies the attribute values, the security and IT mangers rank the outcome attributes relative to their concerns. Finally, threat data and attribute ranks are used to generate the threat index.

After identifying the outcome attributes, the security manager estimated the frequency and attribute values for each threat. The security manager estimated that the frequency among attacks ranged from 2-3 times per hour to 2-3 times per year. Together, the security and IT managers provided the expected attribute values and an upper and lower bound attribute value for each outcome. In addition, the analyst assigned probabilities to the expected, upper, and lower bounds ($P_{exp,}$ $P_{high, and}$ $P_{low}$) to estimate how likely an expected, high and low outcome event will occur. The initial values for expected, high, and low were .89, .01, and .1, respectively. Although these values were based on interviews with security experts, security mangers could adjust them if they believed the values should be different.

Table 1 shows the results of six risks out of the 28 risks identified. Notice that attribute values do not have to be in similar units. In

| **Threats** (Estimated Attacks/year) | | **Lost Revenue** ($$) | **Reputation** (0- 6 Scale) | **Lost Productivity** (hours) | **Regulatory Penalties** (0- 6 Scale) |
|---|---|---|---|---|---|
| Scanning (10,220/yr) | Low | 0 | 1 | .25 | 0 |
| | Exp | 0 | 1 | .5 | 0 |
| | High | 1,000 | 4 | 1 | 0 |
| Procedural Violation (4,380/yr) | Low | 0 | 0 | 0 | 0 |
| | Exp | 0 | 1 | 2 | 0 |
| | High | 12,000 | 4 | 40 | 3 |
| Browsing (2,920/yr) | Low | 0 | 0 | 0 | 0 |
| | Exp | 0 | 1 | 0 | 0 |
| | High | 0 | 4 | 8 | 0 |
| Distributed Denial of Service (DDoS) (156/yr) | Low | 0 | 1 | 1 | 0 |
| | Exp | 0 | 2 | 3 | 0 |
| | High | 0 | 3 | 200 | 0 |
| Password Nabbing (365/yr) | Low | 0 | 0 | .5 | 0 |
| | Exp | 0 | 0 | .5 | 0 |
| | High | 0 | 1 | 1 | 0 |
| Personal Abuse (110/yr) | Low | 0 | 0 | 0 | 0 |
| | Exp | 0 | 0 | .17 | 0 |
| | High | 0 | 1 | 4 | 0 |

**Table 1 Threat Frequencies and Outcome Values**

[1] Regulatory Penalties are defined as the increased oversight and additional operational procedures needed to satisfy external auditors or regulators.

this risk assessment a 7-point Likert scale [4] is used for reputation and regulatory penalties. Lost revenue is reported in dollars lost, and lost productivity is estimated in person hours lost. The first three risks are the top system risks derived from the multi-attribute risk assessment. The data from these risks are the security manager's actual estimates of the frequency and outcomes of successful attacks.

In step three of the risk assessment the *Swing Weight Method* [3] determines the relative weights of the outcome attributes. First, the managers ordered each outcome attribute to reflect their relative concerns among attributes. For example, in this case study, the managers were more concerned about lost productivity than reputation, more concerned about reputation than regulatory penalties and more concerned about regulatory penalties than lost revenue. After the initial ordering, the managers ranked each outcome attribute on a scale of 1 to 100. The attribute that is of most concern to the managers received 100 points. The managers ranked the others relative to the first. Finally, the ranks were normalized to values between 0 and 1, which were used to produce the attribute weights. Table 2 shows the ranks and weights for this case study.

| Attributes | Ranks | Weights (W) |
|---|---|---|
| Lost Productivity | 100 | .42 |
| Reputation | 80 | .33 |
| Regulatory Penalties | 40 | .17 |
| Lost Revenue | 20 | .08 |

**Table 2 Attribute Ranks and Weights**

Although a detailed discussion is beyond the scope of this paper, SAEM uses value functions, $V_j(x_j)$, for each outcome attribute *j*. The value function normalizes each attribute value $x_j$ so that the values can be summed together. A key benefit of using multi-attribute decision techniques is that dissimilar attribute value units, e.g. hours dollars, Likert scales, can be summed together once the attribute values are normalized and weighted according to their attribute weights $W_j$. The method computes a threat index $TI_a$ for each attack *a* based on the security manager's subjective attribute value estimates and frequency assessments using the following function:

$$TI_a = Freq_a * [p_{low} * (\Sigma_{j=attributes} W_j * V_j(x_{j\ low})) + p_{expected} * (\Sigma_{j=attributes} W_j * V_j(x_{j\ expected})) + p_{high} * (\Sigma_{j=attributes} W_j * V_j(x_{j\ high}))]$$

The TI values are dimensionless units that capture the relative importance of each type of risk. Table 3 shows the threat indices for the six risks. Security mangers reviewed the threat indices after providing the risk frequencies and outcome attribute values and confirmed that the relative risk ordering represented their initial assessments. The total TI summed from the risks in Table 3 is 1,507.00.

## 3.2 Risk Assessment Conclusion

Although threat indices are based on best-guess subjective estimates of the threat frequencies and outcomes, the purpose of the multi-attribute risk assessment is to quantify this experience soundly so that assumptions are made explicit and the security managers can see how these assumptions propagate through the assessment. Security managers should revise the risk assessment when they have information that is more reliable. In the next section, the risk assessment results help determine the overall benefit of security technologies.

| $Freq_a * P_{low/\ exp/\ high} * \Sigma (W_j * V(x_{j\ low/exp/\ high}))$ | | | | **TI** |
|---|---|---|---|---|
| Threats | Low $P_{low}$=.1 | Exp $P_{exp}$=.89 | High $P_{exp}$=.01 | |
| Scanning | 85.61 | 765.88 | 34.95 | 886.44 |
| Procedural Violation | 0.00 | 338.39 | 28.59 | 366.98 |
| Browsing | 0.00 | 216.57 | 10.14 | 226.71 |
| DDoS | 1.33 | 23.86 | 0.93 | 26.12 |
| Password Nabbing | 0.03 | 0.28 | 0.31 | .62 |
| Personal Abuse | 0.00 | 0.03 | 0.10 | .13 |
| Total | 86.97 | 1,345.01 | 75.03 | 1,507.00 |

**Table 3 Threat Indices**

# 4 SECURITY ATTRIBUTE EVALUATION MODEL

The purpose of (SAEM) is to provide a structured cost-benefit process to evaluate alternative security designs. This process involves four steps: 1) benefit assessment, 2) threat index evaluation, 3) a coverage assessment, and 4) cost analysis. This section outlines the steps and uses the risk assessment from Section 3 to illustrate the benefit assessment step and shows how security managers can use SAEM to select a security technology for inclusion in the security architecture.

## 4.1 Benefit Assessment

The benefit or effectiveness of a security technology is an assessment of how well the technology mitigates a risk. A technology can mitigate risk in two ways: prevent an attack from occurring or reduce the consequence of a successful attack. Security technologies reduce the consequence of an attack because security managers detect an attack, which gives them an opportunity to either stop an ongoing attack or identify the damage; therefore, security technologies are classified into categories based on their effect on the risk.

### *4.1.1 Security Technology Categories*

Security technologies fail for a variety of reasons so security experts recommend using more than one countermeasure against expected threats. This security engineering principle is known as *defense-in-depth* [1]. Furthermore, the National Institute of Standards and Technology [6] recommends that security managers consider adopting technologies that ensure that protection, detection, and recovery mechanisms are all included. Intuitively, security managers would like to stop a threat from succeeding, but if a threat gains access to the system then it is

essential to quickly detect the intruder and recover from the damage.

Using the defense-in-depth model, security technologies can be classified as *protection*, *detection*, or *recovery* mechanisms. These classifications indicate how a security mechanism mitigates risk. All security technologies fall into at least one category, but some technologies fall into two categories, e.g. auditing.

Inherently, some protection mechanisms detect a threat before stopping it; however, the primary function of the protection mechanism is to prevent the threat from succeeding at all. Since *protection* mechanisms stop a threat from succeeding, they reduce threat frequencies. Detection and recovery mechanisms mitigate threat outcomes from a compromised system. The purpose of a detection mechanism is to identify that an attack is ongoing, or has occurred. Finally, recovery mechanisms detect potential damage and give system administrators the ability to restore system integrity if possible. Table 4 shows a classification of common security mechanisms.

| Protection | Detection | Recovery |
|---|---|---|
| Packet Filter Firewall (PF FW) | Host-Based IDS | Log Analysis Apps |
| Application Firewall (AP FW) | Network Monitors | Auditing |
| Circuit Firewall (CIR FW) | Net-Based IDS | Back-up and Recovery Tools |
| Smart Cards | Auditing | Load Balancing |
| Authentication Policy Servers | Key Stroke Replicator | Key Stroke Replicator |
| Virtual Private Networks | | Forensic Software |
| Virtual Private Network (VPN) | | |
| Email Content Inspection | | |
| Vulnerability Assessment | | |
| Anti-virus Software | | |
| Line Encryption | | |
| Hardened OS | | |

**Table 4 Technology Categories**

### *4.1.2 Relevant Security Technologies Benefits*

After classifying security technologies, the next step in assessing the benefits of security technologies is to identify which technologies mitigate each of the threats. For example, the security manager identified several technologies that help mitigate scanning threats, such as application relay firewalls, and vulnerability assessment scanners, virtual private networks, packet filter firewalls, etc. Table 5 shows the six threats for the financial and accounting system and the security technologies that the security manager believed would mitigate the threats.

| Risk | Security Technologies |
|---|---|
| **Scanning** | Application firewall, vulnerability assessment scanners, virtual private network, packet filter firewall, host-IDS, network-IDS, modem access control, hardened OS[2] |
| **Procedural Violation** | All authentication mechanisms, access control policy server, encryption, electronic signatures, auditing, auth policy server, forensic software, host-IDS, log analysis |
| **Browsing** | Host-IDS, database encryption, auth policy server, auditing, log analysis |
| **DDoS** | Hardened OS and network monitoring |
| **Personal Abuse** | Secure OS, auth policy server, forensics, email filters, log analysis, auditing |
| **Password Nabbing** | Line encryption, 1 time password (1xPWD) , smart cards, secure OS, hardened OS |

**Table 5 Relevant Technologies**

### *4.1.3 Benefit Estimates*

Perhaps the most difficult and controversial piece of benefit assessment is to quantify the effectiveness of countermeasures. Although security managers recognize that precise effectiveness metrics are unobtainable, they are able to provide rough estimates. As part of the interview, the analyst asks the security manager to estimate the effectiveness of each technology against each threat. These estimates are based on the security manager's experience of working with many of these technologies, his assessment of the organization's ability to correctly configure and maintain the technology, his expectations about the skill level and motivation of the attackers, and the organization's policies and system design.

Table 6 shows the security manager's benefit estimates for technologies that could be used to mitigate the risks identified during the risk assessment. Each value in Table 6 represents the percentage that a threat is reduced. For example, the security manager estimated that vulnerability assessment scanners would reduce the frequency of successful scanning attacks by 66%[3].

In a few cases, the security manager identified a pair of technologies that in combination provide mitigation. For example, the security manager stated that hardening the OS together with network monitoring software prevented 75% of the distributed denial of service attacks. The network monitoring software could detect an ongoing Distributed Denial of Service (DDoS) attack, and a hardened OS could mitigate the outcomes from a DDoS. In cases where a pair of technologies mitigated a risk, a combined technology is used (See Table 6, last row).

| Threats / Technology | Scanning | Proc Violation | Browsing | DDoS | Password Nabbing | Personal Abuse |
|---|---|---|---|---|---|---|
| Hardened OS | 66% | | | | 50% | |
| AP FW | 75% | | | | | |
| Vuln Asses Scanners | 66% | | | | | |
| Forensic Software | | 100% | | | | 40% |
| Email Filters | | | | | | 100% |
| Host-IDS | 33% | 50% | 50% | | | |
| Network-IDS | 33% | | | | | |
| 1xPWD | | 95% | | | | |
| Biometrics | | 100% | | | 100% | |
| Smart Card | | 25% | | | 95% | |
| Auditing | | 40% | 30% | | | 40% |
| Log Analysis | | 40% | 30% | | | 40% |
| Auth Policy Server | | 25% | 30% | | | 75% |
| DB Encrypt | | | 30% | | | |
| Net Monitor w/ Hard OS | | | | 75% | | |

**Table 6 Effectiveness Estimates**

In developing security technology benefit assessments, security managers selected all the technologies that they believed would mitigate their risk and estimated the effectiveness. These effectiveness values were based on the organization's ability to employ and maintain the technology; therefore, the effectiveness of the technologies varies across organizations. The security manager's estimates may not reflect the actual effectiveness of a security technology, but until better approximations are made, managers use these estimates to allocate security resources.

Two points about the effectiveness estimations are worth mentioning. The first is that security experts may disagree with the security manager's selection of technologies; however, when questioned, the security manager had clear and convincing reasons for selecting each technology. Second, security experts may also disagree about the effectiveness values. However, since better estimates are not available, a rational approach is to conduct sensitivity analysis to understand how sensitive the decisions are to the security manager's assumptions and estimates. Section 5 discusses sensitivity analysis.

## 4.2 Threat Index Evaluation

The second step in using SAEM is to evaluate the effect each security technology has in mitigating risk. In this step the benefit assessment is applied to the threat frequencies and outcomes to determine how the overall threat index is affected. This section describes a simplified example drawn from the case study in which the security manager is considering increasing the system security and he must select a security mechanism from among three alternatives. This example does not consider how the organization's existing security technologies would affect the effectiveness of each alternative.

### *4.2.1 The Alternatives*

Consider the situation where the security manager wants to improve the overall system security and is trying to decide whether to invest in a log analysis package, a host-based intrusion detection system, or hardening the operating system. Assume for now that each alternative is approximately equal in cost. Table 7 shows each option and which threats are affected by each technology. In addition, several security technologies are already installed, such as an authentication policy server, network monitors, auditing and line encryption. Using the defense-in-depth model, Figure 2 depicts these existing security technologies and their placement in the model. Which technology is chosen will depend on the incremental benefit each technology provides, the existing security architecture, and the cost to implement and maintain each one.

| Hardened OS | Host-based IDS | Log Analysis Software |
|---|---|---|
| Scanning | Scanning | Procedural Violation |
| Password Nabbing | Procedural Violation | Browsing |
| DDoS | Browsing | Personal Abuse |

**Table 7 Security Technologies and Risks They Mitigate**

### *4.2.2 Estimating the Overall Effectiveness*

The second step of SAEM is to determine the risk reduction impact. Since each security technology reduces the risk from several threats, comparing technologies requires an overall assessment on the threat index.

The hardened OS is a protection mechanism. It reduces the frequency of each attack by the amount estimated by the security manager. Similarly, host-based intrusion detection systems (detection mechanisms) and log analysis packages (recovery mechanisms) reduce the outcome values by the amount estimated by the security manager. Table 8 shows the new threat indices and the percentage change between the old threat indices and the new ones.

Recall from Section 3 that the total threat index was 1507. The Log Analysis software has the least risk reduction with a new total threat index of 1,309, a 14.1% improvement, whereas hardening the operating system or purchasing a host-based intrusion detection package would yield a significant improvement – 39.69% or 38.67% respectively. The benefits of either the hardened operating system or host-based intrusion detection system appear to outweigh the benefit of the log analysis software, and their benefits are approximately equal. Now the

security manager must decide between the hardened operating system and the host-based intrusion detection system.

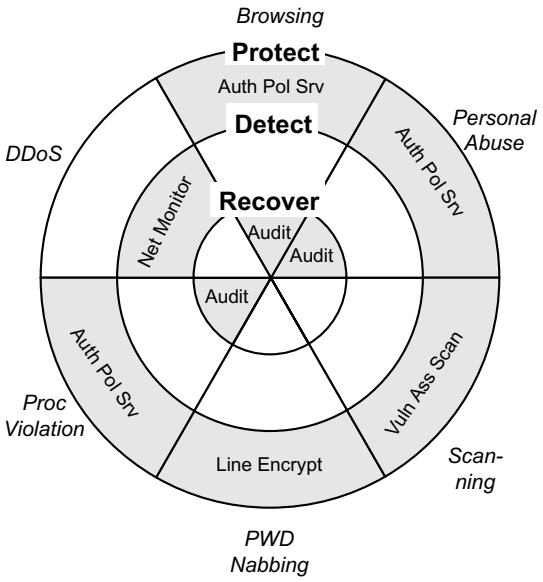


Figure 2. Existing Coverage

| Threat | Hardened Operating System | Host-based IDS | Log Analysis Software |
|---|---|---|---|
| Scanning | 301.39 | 593.92 | 886.44 |
| Procedural Violation | 366.98 | 183.49 | 220.19 |
| Browsing | 226.71 | 113.35 | 158.69 |
| DDoS | 6.53 | 26.12 | 26.12 |
| Password Nabbing | .31 | .62 | .62 |
| Personal Abuse | .13 | .13 | .13 |
| Total | 919.28 | 934.87 | 1,309.38 |
| Change | 39.69% | 38.67% | 14.10% |

**Table 8 New Threat Indices**

## 4.3 Security Architecture Coverage

The decision to select one technology over another might be based on engineering design principles rather than strictly an effectiveness evaluation. In addition, to the defense-in-depth principle, another general security engineering principle applied in security architecture designs is breadth-of-coverage. Breadth-of-coverage requires that there should be at least one mitigation strategy for each risk. A quick inspection of Figure 2 shows that there is at least one security mechanism for each of the threats. If there had been a threat for which there wasn't a mitigation strategy, the security manager might choose a security technology that covers this open threat regardless of the effectiveness values of other technologies.

However, although there is at least one security technology for each threat, in several cases there is *only* one. Recall from Table 4 in Section 4, that a hardened operating system was categorized as a protection mechanism and host-based intrusion detection systems were categorized as detection mechanisms.

Although the benefit of selecting a hardened operating system is almost the same as selecting the host-based intrusion detection system, their role in the defense of the system is considerably different. Figures 3 and 4 show where the host-based intrusion detection system and hardened operating system the fit into the layered defense strategy. The hardened OS provides additional protection, but the application relay firewall and vulnerability scanners already provide some protection. In contrast, the host-based intrusion detection system adds to the overall security because it provides additional security in the middle layer where detection mechanisms were missing. Now, the security manager might prefer the host-based IDS because it provides greater defense-in-depth.

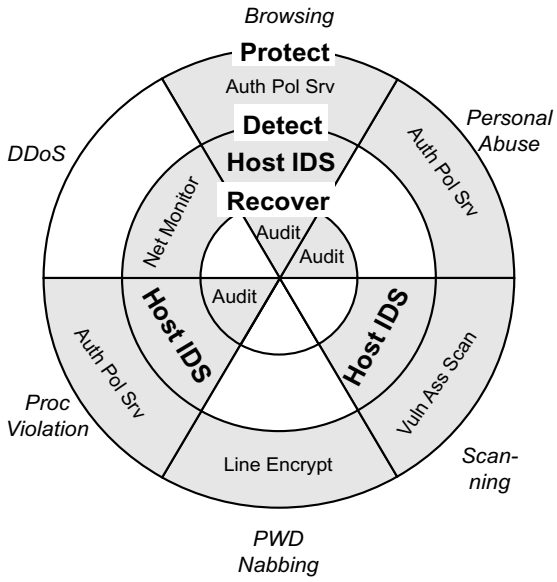


Figure 3. Host IDS Coverage

## 4.4 Cost

The security manager can select a security technology also on the basis of cost, such as purchase, training, maintenance and installation costs. Technology costs are highly dependent on system architectures and detailed designs so it is nearly impossible to determine costs in isolation from a particular design. For example, a system architecture that has several Internet portals would need a firewall for each portal.

Since determining security technology costs can require a significant amount of time, SAEM directs the security manager's attention to the technologies that will provide the most benefit first. The security manager doesn't waste time on security technologies that provide little value. In addition, the defense-in-depth model highlights potential gaps in system defenses. For example, the risk assessment rated scanning as the most significant threat.

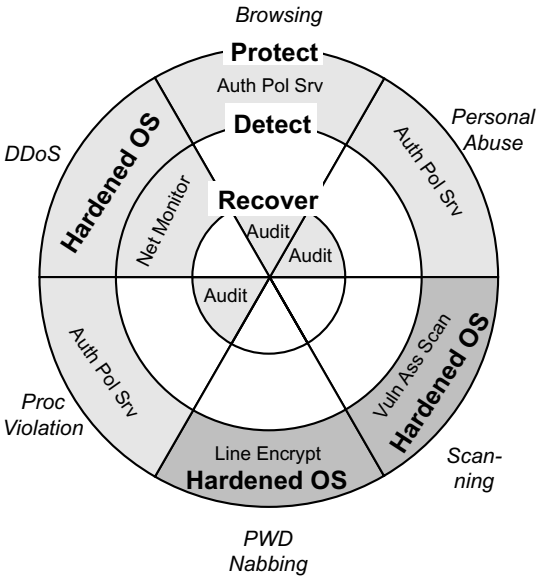


Figure 4. Hardened OS Coverage

Figure 3 showed that the only defense was a vulnerability assessment scanner. The defense-in-depth model indicates that additional security resources may be warranted.

In the case where two technologies appear to provide similar benefits and one does not have a design advantage over the other, then other factors, such as purchase, implementation, or maintenance costs will likely determine the final decision. More detailed information, such as specific capability, operational costs, and licensing fees, must be obtained before an exact system cost can be determined.

# 5 SENSITIVITY ANALYSIS

The purpose of sensitivity analysis is to determine how sensitive the analysis is to the security manager's range of uncertainty about key variables or assumptions. The security manager could make errors in estimating the benefits. In addition, the method presented in this paper assumes that benefits apply equally across the outcomes. Sensitivity analysis allows the security manager to explore the estimates and assumptions to understand how they affect the selection. The security manager can also explore different assumptions about how benefits apply across different outcomes.

## 5.1 Benefit Assessment Sensitivity Analysis

Benefit estimates can have different types of errors. First the security manager could be uniformly optimistic or pessimistic about his benefit estimates. Second, the security manager could error in his estimation of a particular security technology and threat. Unsurprisingly, uniformly optimistic or pessimistic errors do not usually affect the final decision, because the estimates and risk assessments produce relative orderings, not absolute values.

Additional interviews, which determined upper and lower bounds for the benefit assessments could help managers see the affect of the second type of error. For example, the security manager stated that he thought hardening the operating system would stop approximately 66% of the scanning attacks. An additional interview might also determine that the security manager believed that at least 50% of the scanning attacks would be stopped, but it would not be realistic that more than 75% would be stopped. These values can be substituted for the expected values to see if the revised benefits assessments would result in a different security selection.

## 5.2 Uniform Distribution of Benefits

Recall from Section 4 that protection mechanisms reduced the frequency of attacks and that detection and recovery mechanisms reduced the outcomes. In Section 4, the benefit of a security technology was determined by reducing the frequency of the attack; therefore, highly skilled attackers and moderately skilled attackers were stopped equally well. In reality, security experts will claim that highly skilled attackers can penetrate most defenses and that these technologies are usually more effective against the unskilled and moderately skilled attackers. Furthermore, it might be reasonable to assume that the highly skilled attacker is more likely to seek greater rewards for his efforts so that the unexpectedly high outcomes are a result of the highly skilled attackers that can bypass the security mechanism. If the benefits of hardening the operating system and host-based intrusion detection are applied to only the low and expected results and one assumes that highly skilled attacks result in the more severe outcomes then the benefits of hardening the operating system and host-based intrusion drop to 38.56% and 37.06% respectively. Both technologies remain relatively unchanged. Again, this is unsurprising because the values are relative, not absolute.

# 6 FEEDBACK FROM CLIENTS

Initial feedback from this and other case studies has been very positive. Security managers recognize that their estimates are not precise predictions about security technology performance, but they appreciate the structured approach and the ability to see how their assumptions and estimates affect their decisions. Surprisingly, the security managers were less concerned about how the estimates were derived than that their estimates were consistent.

IT managers are also very positive about the cost-benefit approach. These managers can see how security managers make decisions and understand how these decisions fit together with previous security designs.

# 7 FUTURE WORK

The benefit assessments presented in this paper represent one security manager's experience and expertise. Clearly, his assessment may not reflect general opinion in some cases. Also, since the method depends on the security manager having expertise, new or inexperienced security managers will find it difficult to use SAEM. This method shows that it is worth investing in developing better benefit estimates. These estimates could be used to support security managers in making estimations when they lack expertise.

In addition to better security benefit estimates, this method determines the security technology benefit assuming no other security technology is present. Obviously, the full benefit may not be realized if other, similarly mitigating technologies are present. SAEM needs to show incremental benefit estimations in the context of other technologies. For example, the security manager estimated a 66% reduction in scanning attacks from hardening the OS, and 75% reduction from an application relay firewall. If the application relay is already installed, then the benefit from the OS

will unlikely be 66%. Additional work needs to be done to determine the incremental benefit of technologies when others exist.

Finally, the security managers in each study have asked for an automated tool to support sensitivity analysis. Twenty-eight risks with more than forty security technologies make it difficult for the security manager to ensure that the estimations are consistent. With so many variables, sensitivity analysis is tedious and security managers would prefer to conduct their own sensitivity analysis. Automated support could help ensure consistency. In addition, as new risks and security technologies appear, the security managers see that they could use the tool to quickly enter the new information and see the effects of these changes in their system. A prototype tool is currently under development.

## 8 SUMMARY

SAEM provides a structured cost-benefit analysis technique that helps security practitioners evaluate the selection of security technologies. SAEM depends on a risk assessment that results in a prioritized list of risks based on attack frequencies and outcomes. Security managers provide best-guess estimations of the security technology effectiveness, which is used to determine the potential risk mitigation benefit. Sensitivity analysis is conducted to determine how sensitive the security manager's estimates are to assumptions. The purpose of SAEM is to help security managers select the most cost effective security technology.

## 9 ACKNOWLEDGEMENTS

This research was supported by the National Science Foundation under Grant CCR-0086003 and by the Software Engineering Institute, under the Networked Systems Survivability Program

## 10 REFERENCES

[1] Anderson, Ross. Security Engineering: *A Guide to Building Dependable Distributed Systems*. Wiley Computer Publishing. 2001. Chapter 22.

[2] Butler, Shawn and Fischbeck, Paul. *Multi-Attribute Risk Assessment.* Technical Report CMU-CS-01-169, December 2001.

[3] Clemons, Robert T. and Reilly, Terence. *Making Hard Decisions*. Duxbury, 2001.

[4] Fenton, Norman and Pfleeger, Shari. *Software Metrics: A Rigorous and Practical Approach*, PWS Publishing Company, 1997 (Second Edition).

[5] King, Christopher M. , Dalton, Curtis E., Osmanoglu, Ertem T. *Security Architecture: Design, Deployment and Operations*. Osborne/McGraw Hill - RSA Press, 2001.

[6] National Institute of Standards and Technology Special Publications 800-30:*Risk Management Guide* (DRAFT), June 2001.

[7] National Institute of Standards and Technology Special Publication 800-14: Generally Accepted Principles and Practices for Securing Information Technology Systems. 1996.

[8] Yoon, K. Paul and Hwang, Ching-Lai *Multiple Attribute Decision Making: An Introduction*, Sage Publications, 1995.