

A Simple Matrix Approach to Risk Assessment

1. Risk Assessment - General

Risk is uncertainty and security risk is an expression of the uncertainty concerning the security of a system. A security risk assessment is an analysis which attempts to measure the uncertainty in the security of a system. Decision-makers may use a security risk assessment to decide questions affecting the security of a system such as:

- Is the system good enough to deploy?
- What deficiencies should be fixed?
- How do proposed modifications affect risk?
- When should a system be abandoned?

1.1 Methodology - General

Risk assessments can be either qualitative or quantitative, but the inherent uncertainty concerning the results of a security evaluation and the motivation or actions of potential adversaries demands a more qualitative approach to security risk assessment. A qualitative security risk assessment of a client's system according to a general security risk methodology can be used to assess security risk in any information security system.

The methodology measures security risk in terms of:

- the probability an adversary succeeds against the system,
- the corresponding system security consequences of an attack,
- the four primary threats against the system:
 - Organized crime,
 - Sophisticated fraudster with a profit motive,
 - Sophisticated fraudster with a disruptive motive,
 - Non-sophisticated fraudster with a profit motive.

For a client's system we closely examine all the known types of attacks on the system. In order to compare the attacks, we defined several attributes which determine either the probability that an attacker will succeed or the seriousness of the system security consequences. The combination of these two results, the probability of success and the consequences, indicates the risk which the attack presents to the system. For each attribute we define four categories of seriousness ranging from "Very low", "Low", "Moderate", "High." For each attack we then combine these "scores" to arrive at an expression of "Risk"¹ which ranges from {Very low, Very low} to {High, High}. Attacks in the latter category for a given adversary represent significant security risk to the system; attacks in the former category represent the least security risk to the system. There are 16 possible ratings using this method.

¹ Risk is expressed as a pair of ratings, e.g., {Low, Moderate}, indicates that there is a "Low" chance of the attacker succeeding and a "Moderate" consequence to a successful attack.

2. METHODOLOGY

2.1. BASIC CONCEPTS

Risk assessment is an analytic process which provides a person, the decision-maker, information relevant to a particular critical high stakes decision involving uncertainty. What is at risk varies depending on the issue, the decision may be to take action or not, and the analytic process may be more or less scientific or quantitative depending on the subject at hand. In the area of security systems the risk issue is whether the system provides adequate security now and in the future against existing and postulated threats and, if not, how the system might be modified to do so. The decision-maker might be the owner or developer of the system.

2.1.1. Models

There are many different but related models used for risk assessment. The insurance or industrial model is fundamentally a cost-benefit model. The insurance industry derives premiums, and therefore decides company profitability, based on life expectancy or probability of catastrophic event occurrence. The automobile industry uses a similar analysis to decide whether to fix a safety-related problem. Automobile executives might weigh the probability of accident and consequential lawsuit judgments against the cost of fixing the problem in an entire car line. The medical industry uses a model which is fundamentally based on probability of event occurrence and associated consequences; e.g., if one takes a particular drug there is a 1 in 100 chance of having a given side effect, or there is a 1 in 20 chance of surviving a particular operation. Patients, the decision-makers in this case, can use this information to decide what course of action to take.

Both of these approaches benefit from real data gathered over time; clinical trials in the medical field and payouts in the insurance industry. Not only are probabilities of occurrence and expected loss calculable but probability distributions with associated statistics (mean, variance, etc.) can be derived. This leads to a statistically supportable risk analysis in these cases. Because these assessments are statistical, it is rare that a single piece of data significantly and suddenly changes the decision-makers view of risk.

2.1.2. Judgment vs. Measures

Inherent to all risk methodologies is the subjective vs. objective problem. The decision-maker wants as much objective data as possible on which to make risk judgments yet the decision of what is a significant factor and what constitutes acceptable risk is highly subjective. What constitutes acceptable risk for some may be unacceptable for others even when the same factual data and circumstances are used to convey risk. This is exacerbated in the context of security analysis which is itself more art than science. Much of the data on which security judgments are made is postulated and must be projected to the future. One may not know when the subject system is under attack and at risk. New data (attacks) can appear at any time which can thoroughly defeat the system's security and radically alter prior judgments of risk. Additionally, technology advances significantly affect security risk over time. These factors tend to differentiate security risk from other types of risk and call for a slightly less structured approach to security risk assessment.

2.2. SECURITY RISK MODEL

We choose a model for security risk based on the estimated probability that an attacker succeeds, P , and the corresponding estimated system security consequences, C , should the selected attack succeed. We condition our estimated security risk on the capabilities and motivation of various threats or attackers. Different adversaries may attempt different attacks for different reasons and succeed with more or less probability, thereby presenting a different security risk. For each attack (or security weakness), A_i , and defined threat, T_j , we derive an estimated security risk pair $\{P_{A_iT_j}, C_{A_i}\}$ which is intended to convey security

risk of the system for attack A_i executed by an adversary with capabilities defined by threat T_j . As we shall see, the probability that an attacker succeeds with a given attack depends on the threat. However, the security consequences to the system due to an attack are independent of the threat. If the probability of attacker success is significant and the corresponding system consequences are also significant, we say the system has significant security risk defined by that adversary.

Inherent in this security risk model is the assumption that attacks against a system have been identified via a careful security evaluation and that associated probabilities of success have also been identified. (As mentioned above, weaknesses must also be included in the risk analysis but will surely have significantly less effect.) In our risk assessment, we must assume that the potential attacker has knowledge of the system and of attacks known to us. It would be foolish to assume the attacker did not possess information which he might be expected to be able to obtain (e.g., the algorithms used for authentication) or to assume he did not observe a particular weakness discernible from a careful analysis of the system. Although the adversary may have knowledge of attacks unknown to us, we do not factor this into our risk assessment methodology.

By considering all attacks known against the system and estimating the corresponding security risk pairs over a specific type of adversary we can obtain an overall estimate of the system's security risk against specific threats. However, we recognize that a single attack can radically affect security risk. If the probability a particular attacker succeeds with a specific attack is thought to be "high" and the corresponding system security consequences is also thought to be "high", a decision-maker may conclude that the security risk of the system defined by that attack executed by the specified attacker is so high that some appropriate risk-mitigating action to counter the attack is necessary.

2.2.1. Uncertainty

The proposed risk model is based on estimated probability of occurrence that an attacker succeeds, P , and the corresponding estimated security consequences, C , to the system. In this sense the model is closest to the medical risk model however it is much more subjective. The data on which we base our estimate of risk (motivation of attackers, existence of attacks and weaknesses, and corresponding system security consequences) is not numerical and may not be quantifiably supportable. For example, a study which says that a particular attacker will not choose to exploit the system cryptanalytically because there is no shortcut attack against the system may be in error because the evaluation failed to find some particular cheap attack known to the attacker. Although both of the factors on which we base security risk are estimated (and vary with analysts and are subject to error), we choose to ignore variance in these estimates.

2.2.2. Constancy and Uniformity

Security risk is not constant and in fact changes over time for many reasons. Clearly computational technology improves with time making possible some previously impossible or improbably attacks. Attack technology also improves with time so that new attacks can be discovered at any time thereby causing a significant, almost instantaneous, increase in security risk. These factors require that security risk analysis include both identified attacks (specific strategies to defeat security) and less well defined security weaknesses which might lead to attacks. Motivation of attackers can also change over time perhaps because system usage changes. Suppose such changes increase the return on an attack. This increases the probability an attacker will choose to exploit the system and also makes the system consequences due to the attack greater. Both of these factors increase security risk. Similarly the environment in which the system operates may change over time providing more or less opportunity to execute a particular attack.

Security risk is not uniform across all attackers or threats. An insider may have more opportunity to exploit the system using attacks not readily available to an outsider. A non-technologically sophisticated attacker may focus only on simple, quick turn-around, hard-to-prevent but limited return attacks. A sophisticated well resourced adversary may execute more sophisticated long term attacks with low risk of detection but having a greater system return. The organized crime element might be willingly to take more

risk to attack a system and are more likely to employ system insiders. Each of these adversaries would provide a different security risk profile to the same system.

2.2.3. Projection

For the reasons discussed above it is usually necessary to project security risk and update it over time as a function of many factors. Security systems are complex, costly, and usually take considerable time to develop. They have great inertia; it is not easy to react quickly to a significant change in estimated risk. However, security system decision-makers usually find it useful to know when they should be prepared to replace or significantly modify a component of a security system. Projected security risk assessments based on continuing evaluation, risk analyses, and postulated but real threats can assist this endeavor.

2.3. SECURITY RISK FACTORS

As discussed above, we choose a model for security risk based on the estimated likelihood or probability that an attacker with specified adversarial capabilities succeeds in attacking the system, P_{AiTj} , and the corresponding estimated system security consequences, C_{Ai} , should the selected attack succeed. For each attack (or security weakness), A_i , and defined threat, T_j , we derive an estimated qualitative security risk pair $\{P_{AiTj}, C_{Ai}\}$ which is intended to convey security risk of the system for attack A_i if executed by an adversary with capabilities defined by threat T_j .

2.3.1. Probability Attacker Succeeds

For a given attack, A_i , the probability an attacker with capabilities defined by threat T_j succeeds against the security system, P_{AiTj} , is the product of the probabilities that the attacker attempts a particular attack on the system and the probability that the selected attack succeeds; i.e.,

$$P_{AiTj} = P(\text{Attacker } T_j \text{ tries}) * P(\text{Attack } A_i \text{ succeeds})$$

In many cases we can obtain good information on the attack success probabilities, but we usually have poor information on the probability an attacker chooses to exploit a system. Whether an attacker chooses to attack is out of the system's control but we can identify a number of factors which might influence a decision to attack the system.

2.3.1.1. Factors Affecting Opponent's Decision to Attack

An adversary must be motivated to attack a security system; i.e., a successful attack must yield results consistent with his goal. For this reason it is important to include adversarial goals in the definition of threats against the system.

Not only must an attack exist for the adversary to execute but the probability of success (given the risk of detection and attack return) must be acceptable to the adversary. When evaluating the probability of attack success, it is important to consider various strategies of attack execution.

The exploitation potential of the attack also plays a critical role. This factor might be measured in dollar cost, attack execution time, or required technological sophistication all of which might vary depending on the attacker. Attack execution time is particularly important if exploitation depends on certain parameters remaining constant. In a cryptographic authentication system the execution time for attacks which recover keys is important since the keys change over time. If an attacker fails to possess the resources necessary to execute a particular attack on the key management system quickly enough to make use of the keys that are recovered, he may be forced to select a less productive attack.

What motivates an attacker varies but usually the more that can be gained, the more likely it is that an attacker will choose to attack. Thus the potential return to the attacker and its consistency with his exploitation goals is a critical factor for the attacker to consider in deciding whether to attack a system.

The opportunity the attacker has to execute the selected attack against the system is also important. An adversary with no opportunity to execute the most efficient cost effective attack against the system will, in all likelihood, be forced to choose a different, perhaps less effective, attack which may affect his decision to attack at all. Alternatively, a less-motivated attacker may choose to attack a system because he has easy insider access to the system or its components and the opportunity to execute lucrative attacks not considered by others.

Another factor affecting the adversary's decision to attack the system is the difficulty of detection of the selected attack. Depending on the threat, the probability of detection and risk of punishment may be more (e.g., a sophisticated commercial competitor exploiting the system for profit) or less (e.g., organized crime) important. This factor also has a direct effect on the system consequences.

Thus the attack attributes influencing a motivated attacker's decision to select an existing attack and execute it include:

- i) probability of attack success if chosen;
- ii) exploitation potential in terms of dollars, execution time, technology;
- iii) return on the attack given attacker's goal;
- iv) opportunity to execute the attack;
- v) difficulty of attack detection.

For each of these factors, a high value increases the likelihood that the attacker will select that attack to execute.

2.3.1.2. The Probability That The Attack Succeeds

The second factor in estimating the probability the attacker succeeds is the probability that the selected attack succeeds. Estimating the probability of success of an attack is highly dependent on the attack and can be tricky. Cryptanalytic or mathematical attacks are usually the easiest attacks on which to estimate probability of success. Multi-disciplined system attacks are far more difficult. In all cases various strategies of attack execution must be considered. Often attacks which require continuing physical penetration of a data base or facility can be structured so that only the initial penetration has a single event probability of success. Subsequent executions of the attack may have a significantly higher success rate because the initial attack made subsequent penetration easier by, for example, disabling detection mechanisms and stealing user account data.

Note that this factor also influences the probability the attacker chooses to attack the system in the first place. If all attacks against the system have a "low" probability of success, then, even if the attacker chooses to attack, his overall probability of success is not very high. The risk to the security of the system is then a function of the system consequences of the selected attack and could still be substantial depending on the motivation of the attacker; i.e., some adversaries could be willing to try low probability of success attacks if the return is substantial and consistent with their goals.

2.3.2. Security Consequences of Attack

The second major factor in our risk model is the consequences of the attack from the system owner's point of view. What will be lost if a particular attack against the system is successful? To the owner of a commercial system, lost revenue is clearly the most visible and probably most important system consequence to an attack. But there are other intangible adverse system consequences to a successful attack including; denial-of-use, loss of public confidence, embarrassment, bad publicity, loss of customers,

legal liability, etc. Also of importance is the publicity of the attack which may encourage others to attempt similar attacks against the system or other systems of interest to the owner.

The scope or range of an attack, in terms of numbers of components or geographical areas affected by the attack, directly affects the system consequences.

Aside from the direct and immediate outcome of an attack, certain attributes of an attack contribute directly to the system consequences. Many of the same factors which contribute to the attacker's decision to attack also contribute to the system's security consequences but perhaps from a different perspective. For example, from a system's perspective, probability of detection directly affects expected loss over time. Thus difficult-to-detect attacks not only lower risk for the attacker but their longevity causes higher system consequences. Additionally hard-to-detect attacks may have an accumulating adverse consequence on the system which is difficult to "cleanup." Opportunities for attacks not only affect the probability an attacker succeeds but also indirectly the system consequences. This is particularly relevant in the case of insider attacks. There might be more opportunities for insider attacks at an installation or service center where more people are involved and physical security might be lax. Because of this, these attacks could have substantially higher system security consequences.

The attack attributes which directly influence the overall consequences to the system of a successful attack are:

- i) revenue loss from attack
- ii) scope or range of attack application
- iii) intangible losses from attack
- iv) difficulty of attack detection.

For each of these factors, a high value increases the adverse system consequence and, for a given likelihood of attack success, the resulting system security risk.

2.4. THREATS

We have already established that system security risk is a function of the threat as described by the motivation and capabilities of the attacker. For a client's system we define and describe four distinct types of attackers based upon existing and postulated threats to the system.

2.4.1. Non-Sophisticated Fraudster-Profit (NSF-P)

This adversary intends to attack the system for the purpose of selling a service to a user for profit or economic gain. It may be information (proposal data, client lists, credit card data), money, or a product (a digital copy of a commercial film). He is generally not technologically sophisticated and does not have access to significant computational or signal resources. He may act as an individual or in a group functioning as a service enterprise. He is least likely to be a system insider; he represents the principal threat to many current systems.

2.4.2. Sophisticated Fraudster-Profit (SF-P)

This adversary also intends to provide a service to his customer for profit but is much more sophisticated than NSF-P. He can obtain and use current commercial grade technology to run mathematical attacks (using individual or networked personal computers), to intercept signals, to modify signals, or to insert signals all of which might be necessary as part of his attack. He does not have unlimited resources, however. He may or may not be an insider himself but if not, he probably has the resources to buy some inside help.

2.4.3. Sophisticated Fraudster-Disruption (SF-D)

This adversary has the same capabilities as the SF-P but has different motivation. He is not interested in selling a service for profit but rather is interested in attacking the system for disruptive or denial-of-service purposes only -- the disturbed individual or group of individuals who get their kicks from bringing down a system.

2.4.4. Organized Crime (OC)

This is an organized group of criminals who are very well resourced, technologically sophisticated, intent on exploiting the system for their own illegal purposes. They either are insiders themselves or can buy or obtain any inside help necessary to run the attack of choice. Taking significant risk to execute the selected attack would not be a problem for them.

2.5. MEASURES

We have already mentioned that security risk analysis is largely a subjective exercise supported by a few objective facts. Models which attempt to convert everything to numbers and force-fit them into a sophisticated mathematical model may be artificial and misleading relative to security risk. These models often suffer from errors in risk definition, selection of restrictive mathematical functions, dependencies which are not accounted for, and cumulative errors.

Recall that in our risk model, for a given threat, T_j , we associate with each attack or security weakness an ordered risk pair $\{P_{AiTj}, C_{Ai}\}$ where P_{AiTj} represents the probability an attacker with capabilities consistent with threat T_j succeeds with attack A_i and C_{Ai} represents the system security consequences of the attack. For reasons already discussed the subjective nature of security risk necessitates that we use qualitative terms, such as {high, moderate, low, very low}, to estimate both the probability or likelihood of attacker success and corresponding system security consequences. Thus, for each of the relevant attack or security weakness attributes discussed above and for the attack attributes contributing to the system security consequences, it is necessary for us to define partially ordered qualitative categories.

2.5.1. Probability of Attack Success

We define attack success as having achieved the intended goal of the attacker. We caution that these likelihood estimates should apply to complete attack strategies not single event tries. We arbitrarily define four categories which we believe are consistent with the range of threats under consideration. Thus a higher likelihood of success is favorable to the attacker and unfavorable to the system.

Very Low (V), less likely than 1 in 1000

Low (L), between 1 in 1000 and 1 in 10

Moderate (M), between 1 in 10 and 1 in 2

High (H), more likely than not

2.5.2. Attack Exploitation Potential

We characterize categories of attack exploitation potential in terms of sophistication required, equipment cost, and execution time consistent with the threats we have previously defined. Note that the attack exploitation categories listed below are absolute and not threat dependent. However, some well resourced attackers may find certain attacks with Very Low absolute exploitation potential easy to execute. We have chosen to display a relative score, i.e. such attacks would be judged to have a High exploitation potential score for that attacker in our assessment. A high value for attack exploitation potential is favorable to the attacker and unfavorable to the system.

Very Low (V), requires state-of-the-art mathematical, engineering and signal processing knowledge, access to networked high performance workstations and complete laboratories, substantial financial resources up to \$10M for exploitation, thorough understanding of system details, ability to reverse engineer components and software, ability to intercept, modify, create, and override selected signals in near real time, ability to spoof system audit detection mechanisms, executable sometime within the cryptoperiod of changeable information.

Low (L), requires substantial commercial computational and signals collection technology, access to networked PCs, university laboratory, limited financial resources up to \$100K, ability to replace or reprogram components, good understanding of the system or component details, ability to intercept, replay, and modify signals, executable well within the cryptoperiod of changeable information.

Moderate (M), requires some computational and signals collection technology available in electronic stores, ability to replace or reprogram components, very limited financial resources up to \$10K, limited understanding of the system or component details, ability to intercept and replay but not modify signals, executable over nearly all the cryptoperiod of changeable information.

High (H), requires no signal processing or computer resources, almost no financial resources, no understanding of system or component details, executable almost instantaneously.

2.5.3. Return on Attack

The return on an attack is very much a function of the attacker's goals. For example, if an adversary intends to exploit the system to send short messages undetectably for illegal purposes, then an attack which allows this would have a high return for that attacker. Below we have characterized classes of attack returns against the mainline authentication system in terms of target, geographical applicability, and time. We note that attack return is correlated to the system security consequences. Again a high value is in favorable to the attacker and unfavorable to the system.

Very Low (V), very limited use of system (e.g., registration only) or denial-of-use of system for a single component in limited area for a short period of time.

Low (L), limited use of system functions for single component over limited period of time or denial-of-use of multiple components in limited area over limited period of time.

Moderate (M), full use of system for single component over a limited period of time or limited use of multiple components over long period of time or denial-of-use for multiple components in multiple areas for limited time.

High (H), full use of system for single component over long period of time or full use for multiple components over limited period of time, or denial-of-use of the system for a geographical area, or unlimited use of the system for other purposes.

2.5.4. Opportunity for Attack Execution

The opportunity to execute an attack depends on the attack and the system operation. It is presumed that insiders can execute the most serious attacks, require the least amount of time, and, in general, would have the greatest opportunity to do so. Opportunity for attack execution is correlated with the system security consequences. Again a high value is favorable to an attacker and unfavorable to the system's security.

Very Low (V), attack can be run against a single component operating during a short time period or in a specific area.

Low (L), attack can be run against a limited number of operational components during a short period of time or in limited areas, or against a specific activity at a specific facility during a very short period of time.

Moderate (M), attack can be run against a large number of components in multiple areas over a limited period of time or can be run against a specific facility or installation at a specific time.

High (H), attack can be run against any component in any area.

2.5.5. Difficulty of Attack Detection

The probability of attack detection depends greatly on the type, quality, and usage of system mechanisms intended to detect anomalous behavior and on the characteristics of the attack itself. Some attacks can be segmented and run over a period of time or in different locations to avoid system detection mechanisms. In other cases the attacker could spoof audit mechanisms or “condition” the system to confuse detection mechanisms based on profiling. Recall that this factor affects both the likelihood the attacker selects the attack as well as the system consequences from the attack. Below we characterize four classes of attack detection based on expected client’s system behavior and the mainline attacks against the system. A high value is favorable to the attacker and unfavorable to the system.

Very Low (V), attack is automatically detectable by system on subsequent component or system action and positive action is taken to counter the attack (e.g., immediate key change, two person control, encryption of data base, etc.).

Low (L), attack is detectable by the system with on-line mechanisms during selected activities.

Moderate (M), attack is detectable by the system on random audit, profiling, or during infrequently performed actions.

High (H), attack is detectable by user or by the system via monthly audits but not by automatic on-line system mechanisms.

2.5.6. Intangible Losses from Attack

There are many intangible losses that could result from an attack including: loss of public confidence, embarrassment, bad publicity, loss of customers, legal liability, system disruption, denial of service to legitimate users. We define the levels to be:

Very Low (V), effects are hardly noticeable, little or no public awareness or consequences.

Low (L), noticeable effects but minor negative consequences to system or service provider.

Moderate (M), severe disruption for brief periods of time, noticeable loss of customers, or other effects requiring significant company resources to counteract.

High (H), major and frequent disruption of service, severe loss of customers, direct effect on company image.

2.5.7. Scope and Range of Attack Application

We define attack scope and range to mean the extent to which the damage applies when the attack is successful. It is based on a combination of the number of components affected, the duration in time that the attack might apply, and the geographical range over which the attack applies per success.

Very Low (V), damage or loss is very limited to a few components per occurrence and only for a short period of time, e.g. a few minutes or less.

Low (L), damage or loss extends to multiple components, but only in a limited geographical area and/or for a short period of time.

Moderate (M), damage or loss affects multiple components for a moderate time (e.g. hours) and over a moderate geographical area.

High (H), damage or loss effects many components in a wide geographical area and over an extended period of time, e.g. days.

2.5.8. Revenue Loss from Attack

We shall consider revenue loss over time, but shall restrict the measures to qualitative statements since we have no method to measure actual expected loss at this time. We define the levels to be:

Very Low (V), little or no loss to the system as a whole. Losses confined to small amounts or to a few customers per billing period from any one attacker.

Low (L), losses less than any other type of fraud (e.g. fraudulent credit cards used to open an account).

Moderate (M), losses equal to or greater than any other type of fraud.

High (H), losses greater than all other types of fraud combined.

2.6. ASSESSING SECURITY RISK

For each attack A_i and threat T_j we can assess each attack attribute (probability of success, exploitation potential, etc.) using the qualitatively defined categories {High, Moderate, Low, and Very low}. We then combine these qualitative factors to produce an estimated attack risk pair $\{P_{AiTj}, C_{Ai}\}$ where P_{AiTj} and C_{Ai} each takes on one of the values {High, Moderate, Low, and Very low}. The task of combining the qualitative attribute estimates into an attack security risk pair estimate is subjective and not without error; it is basically a judgment. Ideally a set of weights would be derived and a weighted average used to determine the probability of attack success and system consequence components. We choose a more simplistic qualitative approach: the more “highs” and “moderates” exhibited by the attributes of a particular attack, the more likely it is that the components P_{AiTj} or C_{Ai} will be given a “high” value. Similarly the more “very lows” and “lows” exhibited by the attributes of a particular attack, the more likely one of the components P_{AiTj} or C_{Ai} will be given a “low” value. The entire exercise is complicated by dependency of several attack factors including difficulty of detection and opportunity for attack execution.

Having generated attack risk pairs for each attack and threat we can sort these values on threat and plot them in a matrix to provide an estimate of the security risk based on threat. The vertical axis is the probability attacker succeeds, P_{AiTj} , the horizontal axis is the system security consequences, C_{Ai} , of the attack. This is useful to a decision-maker who might be intent on developing or modifying a system to counter a specific threat.

System security risk increases along the diagonal of the matrix running from {Very low, Very low} to {High, High}. Clearly the more attacks that appear in the {High, High} category the more risk the system exhibits against the particular threat. However, existence of attacks in the {High, Moderate}, {Moderate, High} and {Moderate, Moderate} categories should not be dismissed for many reasons. Slight errors in judgment could move these pairs to a more serious category and over time these pairs are the ones which are more likely to move to the {High, High} category.

2.6.1. Thresholds

In general the more attacks which appear in the upper right quadrant, the higher the security risk for a given threat. However, it would be a mistake for a decision-maker to focus only on this area. It is well known that a security system exhibiting a single exploitable attack with consequences consistent with

the attacker's goals might be at substantial risk should the attacker choose to run the attack. In the security business it is dangerous to counter all but a few attacks; it only takes one attack for the adversary to be successful. Some attackers will choose to execute very low probable attacks if the cost is cheap and the return is sufficiently high. It is also possible, although less likely, that attackers will choose to execute attacks having a high probability of success but very little return. Each attack risk pair must be carefully examined by the decision maker, particularly those in which either component is in the High or Moderate category.

2.6.2. Cumulative Risk

A decision-maker looking at a security risk analysis having a lot of entries, none of which are in the {High, High} category, may conclude that the system exhibits acceptable security risk. However, this would be dangerous if the system at hand is one against which there are a substantial number of known attacks. This should not give the decision-maker much confidence in the security of the system. For the reasons already discussed there may be ample opportunities for any of these attacks to percolate up the risk matrix and surely over time we would expect this to happen due to advances in technology. For such a system the security risk could change dramatically and a security risk assessment exhibiting these characteristics should be very carefully examined.

2.6.3. Projecting Risk

As we have already discussed, security changes over time. Given a static system new attacks will appear and existing attacks will get cheaper; what requires sophistication now will require little sophistication in the future. Given a system whose design remains fixed but which changes in application or environment (this is not uncommon), more opportunities for attack may exist, new threats may appear, attack return may increase, probability of attack detection could change, etc. The weaknesses which have been included in our methodology may give rise to attacks over time. Because of attack or computational technology advancements or system changes, new weaknesses will appear. For some of these factors accurate projections can be made to determine how quickly attacks may become viable. This is particularly true for cryptanalytic attacks which depend on computational capability. However, this is not true for system attacks which might depend more on system access for example. In general, unless some countermeasures are taken in the present, over time new attack entries will appear and the entries present in today's security risk assessment will tend to move toward the {High, High} category of security risk.

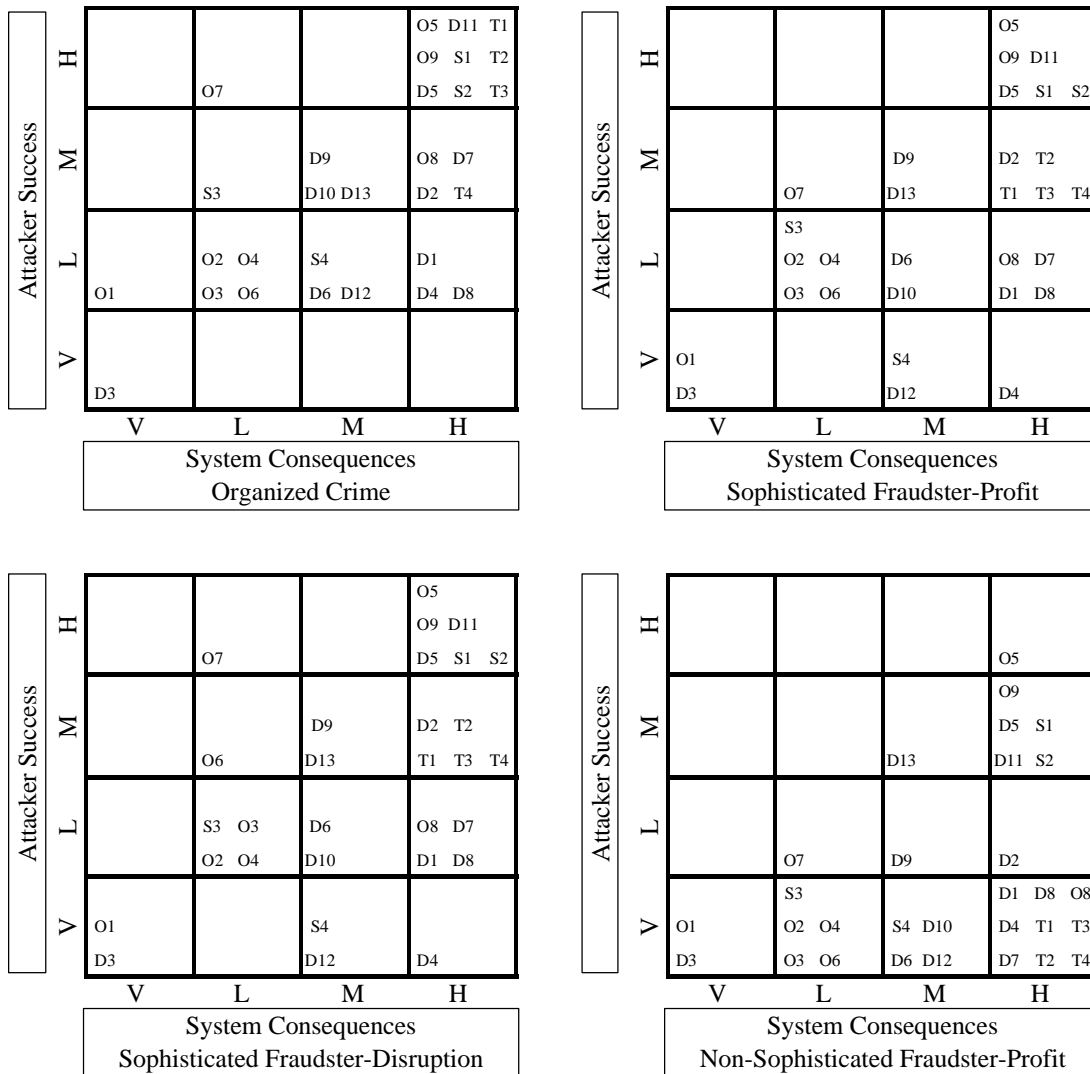


Figure 1: Hypothetical Risk Analysis Diagram

2.7. MANAGING SECURITY RISK

Because security risk changes with time, it is in the best interest of the system to address all attacks which currently are estimated to present a serious security risk against the intended threat. The task of managing security risk is an ongoing endeavor that reassesses system security risk over time as a function of ongoing security evaluation and continuing examination of threat. It is a task which is completed only when the system is turned off.

2.7.1. Risk Reduction Alternatives

At any one time a decision-maker has several alternatives with which to address security risk. In reaction to a particular attack or set of attacks which presents significant security risk, a decision-maker can choose to:

- i) eliminate the risk by fixing the vulnerability via a design change;

- ii) mitigate the risk via some procedural fix which may or may not be followed;
- iii) accept the risk temporarily and immediately plan for a replacement system;
- iv) accept the risk and do nothing.

Assuming the decision-maker wants to mitigate risk now, the goal is to identify cheap fixes which have a significant positive effect on reducing security risk. The decision to fix the vulnerability or adopt a risk mitigating procedural action is usually based on a cost benefit analysis. Procedural fixes are usually cheap although in some cases they can be so user-unfriendly as to adversely affect system or component utility. Although the optional nature of these fixes usually limits their benefit to security, they do have the ability to transfer responsibility to the user and out of the system developer's hands. Hardware or software fixes are often expensive but usually do a good job at eliminating the security vulnerability. This is difficult after a system is designed and partially developed. Just as a procedure that mitigates risk today may do nothing in the future, a quick temporary design modification may be ineffective in the future.

2.7.2. Cost Benefit Analysis

If the choice of the decision-maker is to negate a serious attack and have high confidence that he is successful, a specific security design modification is usually necessary. The cost of such modifications depends greatly on the state of the system (design, development, deployment, new, mature, etc.) and the impact the redesign may have on the rest of the system. It is important that alternative designs be considered and that each undergoes a detailed security evaluation; i.e., an analysis of the modification's security quality and the extent to which the proposed modification counters the specific attack and does not introduce new vulnerabilities. The designer/developer then conducts a detailed cost-benefit analysis of each proposed modification before selecting a specific strategy. Cost is usually measured in terms of dollars, schedule, technology, performance, impact on product already in the field, user impact, system impact, etc. Benefit might be measured in terms of quality of the solution, completeness of the solution, longevity of the solution, backward compatibility with legacy systems, etc.

It is sometimes possible that several security problems can be addressed with a single solution but this is rare. The goal is to select the least costly modifications from the alternatives available which have the maximum benefit in security risk reduction against the threat of interest in the present and in the future. Because most developers are constrained by budget or other factors, it is usually necessary to prioritize risks even after performing a security risk analysis. A tradeoff analysis is done which attempts to maximize the security gain for a given fixed dollar expenditure; i.e. the goal is to find the set of modifications which can be done within budget and other constraints such that the collection of these modifications has a desired combined maximum security risk reduction effect. In our methodology we do not prioritize within categories but against a particular threat we might focus on those modifications which lower risk by moving all attacks out of the {High,High} category. We emphasize however, that a single residual attack which presents significant security risk can be dangerous since the attacker only needs one vulnerability to successfully exploit a security system.