

A Framework for Reasoning about Assurance

April 23, 1998

Document Number: ATR 97043

Developed for:

National Security Agency
9800 Savage Road
Fort George G. Meade, MD 20755

Developed by:

Arca Systems, Inc.
8229 Boone Blvd., Suite 750
Vienna, VA 22182

A Framework for Reasoning about Assurance

Jeffrey R. Williams
Arca Systems, Inc.
williams@arca.com

George F. Jelen
G-J Consulting
gjelen@erols.com

April 23, 1998

Copyright © 1998 Arca Systems, Inc. All Rights Reserved

**This research was partially supported by the National Security
Agency under Contract Number MDA904-97-C-0223**

Executive Summary

Informed decisions about security depend upon a complex set of factors related to both assurance and risk. In this paper we argue for a new definition of assurance, specifically describing its relationship to measurements of risk or security. The following definition expands the traditional definition of assurance to include a broad range of evidence, while narrowing the scope to a specific type of confidence.

Assurance is a measure of confidence in the accuracy of a risk or security measurement.

This paper offers a way to build upon risk and security measurement methodologies and to employ them in such a way as to yield a rough measure of assurance. We do not advocate a particular method for measuring risk or security, but assume that such methodologies and tools are available.

Consider the decision maker who must decide whether to accept a risk or make an investment to mitigate the problem. He may use a risk assessment tool to help quantify the risk, but he may have very little confidence that the numbers are accurate. In essence, he is not sure whether the risk is acceptable or not.

Assurance is a major factor in security decisions.

The decision maker has two options for handling such situations. If his confidence in the risk measurement is high, he can attempt to reduce risk by adding security mechanisms. If his confidence is low, however, adding another mechanism may not help. In fact, adding a new mechanism may even increase uncertainty in the risk measurement. In this case, the decision maker needs to improve assurance by obtaining better information about the severity of the risk.

If the decision maker then decides to reduce this uncertainty, we offer a structure for assurance arguments as a logical way to communicate the information used in making security decisions. An assurance argument starts with claims about risks and then packages all the evidence and supporting arguments into a logical hierarchical structure. The goal is that these arguments will be capable of reuse in a wide variety of applications, easing the burden of security evaluations.

Assurance arguments are a powerful tool to reduce the uncertainty in risk or security assessments.

Although this paper does not provide a means by which one can determine assurance need in the sense of some quantitative or even qualitative statement, it does provide a way of deciding whether or not the assurance one has is sufficient, and this, we claim, is quite good enough.

Table of Contents

1. INTRODUCTION	1
1.1 Purpose of the Paper	1
1.2 Overview	1
1.3 Organization of the Paper	2
1.4 Indebtedness to Previous Research.....	3
2. DISTINGUISHING ASSURANCE FROM SECURITY AND RISK	4
2.1 Measurement Dimension and Assurance Dimension.....	4
2.2 Measurement Alternatives	5
2.3 Risk Background	6
3. UNDERSTANDING ASSURANCE	8
3.1 Estimating Risk.....	8
3.2 Assurance	8
3.3 Risk and Assurance	9
3.4 Investigating the Uncertainty.....	10
3.5 Normalizing Assurance	11
3.6 Relationship between Confidence Coefficient and Uncertainty.....	12
3.7 Consequence Uncertainty	13
4. RISK PLANE DISCUSSION	15
4.1 Case Studies	16
4.2 Factors Involved in Decisions	18
5. DECIDING A COURSE OF ACTION	20
5.1 Adding Mechanisms	20
5.2 Adding Evidence	21
6. INGREDIENTS OF ASSURANCE ARGUMENTS	22
6.1 Claims.....	22
6.2 Evidence	26
6.3 Reasoning	27
6.4 Assumption Zone.....	27
7. STRUCTURING AN ASSURANCE ARGUMENT	29
7.1 Choosing a Top Level Claim.....	29
7.2 The Subject Tree	29
7.3 Evidence Claims	30
7.4 The Predicate Tree	30
7.5 Structuring Principles	33
8. COMPOSING AND DECOMPOSING RISK AND ASSURANCE	35
8.1 Composing Risk	35
8.2 Combining the Risks From Different Events.....	36
8.3 Combining Assurance From Different Events.....	36
9. RELEVANCE AND UTILITY OF THE FRAMEWORK	38
9.1 Network Rating Methodology	38
9.2 System Security Engineering Capability Maturity Model (SSE - CMM).....	39
9.3 Common Criteria	40
9.4 Summary.....	40

Figures and Tables

Figure 1: Assurance and Measurement Dimensions Are Orthogonal.....	4
Figure 2: Determining Uncertainty Using Best Case-Worst Case Approach.....	10
Figure 3: Combining Consequence and Likelihood Uncertainty.....	15
Figure 4: Example Events Plotted on a Risk Plane.....	16
Figure 5: A Simplistic Approach to Risk Decisions.....	18
Figure 6: Subject Tree.....	30
Figure 7: Evidence Claims.....	31
Figure 8: Generalized Predicate Decomposition.....	33
Table 1: Summary of Different Risk Cases.....	18
Table 2: Assurance Arguments Are Built from These Elements.....	22
Table 3: Examples of Assurance-Relevant Properties.....	25
Table 4: Combining Four Independent Risk Events.....	37

1. INTRODUCTION

In selecting assurance methods, the need for assurance should be weighed against its cost. Assurance can be quite expensive, especially if extensive testing is done. Each method has its strengths and weaknesses in terms of cost and what kind of assurance is actually being delivered. A combination of methods can often provide greater assurance, since no method is foolproof, and can be less costly than extensive testing.

The NIST Handbook [NIST95]

1.1 Purpose of the Paper

Informed decisions about security depend upon a complex set of factors related to both assurance and risk. This paper is intended to cast light on the relationships among these factors to enable informed decisions about security risks. The motivation for this work comes from the wide diversity of opinions about what assurance is, how it might be measured and communicated, how much one needs, and how various types and sources of evidence relate.

Although one might wish to be able to develop an absolute measurement of assurance (as well as of security and risk), this would imply, in the case of assurance, both a unit of measurement and at least a reasonably consistent means of measuring it. In our view, neither of these now exist nor are they likely to. What this paper offers instead is a way to build upon quantitative risk measurement methodologies and to employ them in such a way as to yield a rough measure of assurance that ought to permit one to trade off the relative merits of seeking more evidence, and thus gaining greater assurance, against employing more safeguards, thus reducing risk. Although the method does not tell one how much assurance she has, it does tell her whether or not she has enough.

Today's systems—and the enterprises in which they reside—are so complex that even the most capable risk measurement tools are unlikely to yield risk values that are much better than rough indications of *relative* risk—which, we should quickly add, is often quite good enough in many situations. The problem is that the value of risk, whatever it turns out to be, is likely to be surrounded by a fairly large but unknown amount of uncertainty. This can create a dilemma for the decision maker who must decide whether to invest in further safeguards, which will undoubtedly reduce the overall risk but could be both expensive and unnecessary, or to collect more evidence to reduce the amount of uncertainty surrounding the risk calculation—what this paper calls *assurance*.

1.2 Overview

Assurance is commonly defined as “the degree of confidence that security needs are satisfied.” [WITAT95]. The problem with this definition is that, unless one has a way to specify security needs quantitatively, assurance can not be expressed quantitatively either. Many methodologies, however, exist for expressing risk quantitatively, and by tying our definition to the quantification of risk or security rather than that of security needs, it is possible to approach a quantitative expression of assurance. So, for the purposes of this paper, the following definition will be used:

Assurance is a measure of confidence in the accuracy of a risk or security measurement.

Within the security community, the reliance on various criteria and evaluation processes to measure assurance has clouded the rationale for producing and using assurance evidence. In this paper, we argue that assurance is an integral part of the risk and security management processes rather than a special or separate attribute. The above definition clearly links assurance with measurement, making the costs and benefits of producing and evaluating assurance evidence clear and comparable.

1.3 Organization of the Paper

The paper will show how assurance can be measured by structuring the risk inquiry to yield information about uncertainty. We show how risks, along with their uncertainty, can be compared and evaluated. Next, we suggest ways of determining a course of action if the risks are found to be unacceptable. If there is too much uncertainty, we show how an assurance argument can be created to reduce that uncertainty. Finally, we discuss how some current assurance related efforts can be viewed using this framework.

Section 2 distinguishes between the measurement dimension and the assurance dimension. For the purposes of this paper, any measurement technique can be used. However, the examples used in this paper generally assume a risk-based measurement technique. Therefore, a general overview of this approach is also provided here.

Section 3 discusses the idea that any risk assessment process necessarily involves some uncertainty. Fortunately, there are ways to quantify the magnitude of this uncertainty and use it to help make decisions.

Section 4 explores the possibilities for determining whether risks and their attendant uncertainties are acceptable. At a minimum, the person making decisions must be able to compare risks in order to decide which to address first. However, it is also useful to establish an independent threshold of which risks are tolerable.

Section 5 describes various options for deciding a course of action. Certainly, one option is to do something to lower the likelihood component of the risk. However, in cases where there is low assurance, that is, where the uncertainty associated with the risk is high, another option is to try to reduce that uncertainty.

Section 6 explains a way to reduce uncertainty through the construction of an “assurance argument” that captures all the information relevant to assessing a particular risk. The argument begins with a claim that a subject has some particular quality. This claim is supported with evidence and some reasoning which shows that the argument is complete and convincing.

Section 7 shows how risk and assurance can be aggregated to form a complete picture. The approach described here is scaleable and applies whether the subject is an enterprise or a low-level enterprise component.

Section 8 explains that the risk at one level of the assurance argument is not necessarily the sum of the risks at the level immediately below. It also shows how one can determine the total risk and the associated assurance from several independent events. Finally, some thoughts on consequence uncertainty and how it might be reduced are presented.

Section 9 discusses how various assurance methods, criteria, and tools can be compared using the framework. Because each of these approaches yields evidence that can be used to support claims, the particular claims that each supports can be examined to categorize and contrast these different approaches. This section discusses three of these approaches and the types of claims that each supports.

1.4 Indebtedness to Previous Research

Several previous efforts and works greatly influenced this paper. For some time, many different people and groups have recognized the desirability of having some type of framework by which to organize and present an assurance argument. An unpublished paper, also entitled “A Framework for Reasoning about Assurance” [WILL95a], written in 1995 by Jeffrey Williams and Douglas Landoll of Arca Systems for NIST, served as a major source during the writing of this paper. The present paper extends some of the concepts introduced in the earlier paper by more clearly articulating the connection between assurance and risk. Also, “A Framework for Assessing Network Security” [NRM97], targets the development of a “consistent, cost-effective framework within which network security may be assessed and evaluated.” The NRM addresses several topics related to the development of workable security measurements.

2. DISTINGUISHING ASSURANCE FROM SECURITY AND RISK

The view of assurance advocated in this paper depends upon having a technique for measuring risk or security in which uncertainty is an issue. Since there will never be a way to measure either risk or security without introducing uncertainty, the approach to assurance advocated here should be broadly applicable.

In this paper, we focus more on risk measurement than on security measurement primarily because risk tends to be a more easily measurable quantity than is security.¹ Certainly, techniques for measuring or at least estimating the extent of risk are more prevalent than methods for measuring or estimating the amount of security. However, to the extent that security is measured, the techniques described herein for measuring assurance would apply equally well to security-based assurance as they do to risk-based assurance.

2.1 Measurement Dimension and Assurance Dimension

A critical feature of the view of assurance presented here is that it is orthogonal to the measurement of risk and security. Keeping these dimensions separate helps to reduce the confusion associated with the terms. Currently, there is a tendency to associate high assurance ratings with high security and low risk.

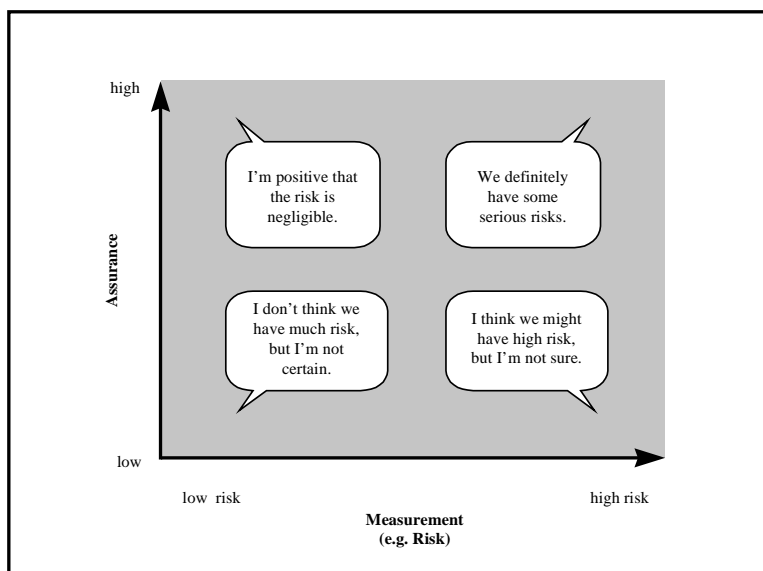


Figure 1: Assurance and Measurement Dimensions Are Orthogonal

Unfortunately, confusing the two dimensions makes it very easy to miss the interesting cases where the two dimensions lead to conflicting conclusions. For example, consider a network component with a large number of security mechanisms, but no information to show whether

¹ In this paper, in those cases when we wish to distinguish between these two, we will use the terms "risk-based assurance" and "security-based assurance."

they are correctly configured or not. Here, combining the measurement of the mechanism with the confidence in that measurement will produce an ambiguous result. By separating the two, it is easy to recognize that it would probably be more useful to gather more information than to add another security mechanism.

Figure 1 graphically presents the two separate dimensions and indicates where some common statements about security might fall. Although the figure depicts a risk-based measurement, that of a security-based measurement would work equally well.

By distinguishing the measurement dimension from the assurance dimension, the approach described in this paper is intended to enable the use of all relevant information in making a security decision. In particular, the framework allows all types of security and risk evidence to be embraced within a single assurance argument. For example, when attempting to determine the security of a local-area network, an assurance argument can use TPEP evaluation results of the components, the vulnerability analysis results for the system, and penetration testing results to support specific claims. This argument informs the risk or security measurement process, allowing an estimate of assurance to be determined.

2.2 Measurement Alternatives

The approach described above can be applied regardless of the specific measurement technique selected. There are many ways to measure the risk or security of something. One popular method is to take a close look at the security features or protections, to see how good they are. This approach has been adopted by various product and system evaluation and assessment programs. This approach is most often used to assess products and systems independently from any particular operational environment. Given the difficulty of foreseeing the impact of security problems when the ultimate application is unknown, this analysis usually starts with the lowest levels of implementation and works upwards.

Another widely used method is to closely examine the risks, to see how bad they are. Risk-based approaches place greater emphasis on the expected impact of a problem. The basic idea behind the risk approach is to look for combinations of threat, vulnerability, and impact that create unacceptable exposure to harm. Generally, this approach is used in operational settings in which the impact of successful attacks can be assessed. Therefore, risk approaches generally start at the highest levels of the enterprise and work downward, looking for exposures that pose unacceptable risks to the enterprise.

At a high level of abstraction, all of these techniques rely on similar information to produce a measure of the same concept, and differ only in perspective or emphasis. The more sophisticated approaches recognize that the measurement must occur at multiple levels of abstraction. In the examples in this paper, we have generally assumed a risk-based technique for assessing security. However, this should not be read to imply that a risk-based technique is required. The definition of assurance presented here can be applied to formal evaluation, certification, accreditation, or to any other security assessment techniques.

2.3 Risk Background

As defined in this paper, the amount of assurance ultimately depends upon the confidence in the accuracy of statements or claims made about the risk associated with some enterprise. Therefore, in order to understand the concept of assurance, one first needs to understand risk. Since the concept of risk embodies threat, vulnerability and consequence, the paper discusses these first.

In this paper we do not advocate a particular method for measuring risk, but simply assume that methodologies and tools exist for such a purpose. We recognize that these methods are not perfect; indeed, if they were, there would be no reason to explore uncertainty. This framework is intended to be general enough that it could be used with any risk measurement method as a means of determining assurance.

Threat has been defined as “any circumstance or event with the potential to cause harm to a system in the form of destruction, disclosure, modification of data, and/or denial of service” [ISU96]. Expressed in mathematical form, *threat* can be defined as the probability that there exists an adversary or circumstance that, given an exploitable vulnerability, possesses the capability, and in the case of the adversary, the motivation and opportunity, to exploit that vulnerability.

Vulnerability is defined as a weakness in the physical layout, organization, procedures, personnel, management, administration, hardware or software that could be exploited to cause harm to an ADP system or to the enterprise in which it resides. Mathematically, vulnerability is the probability, given a threat as defined above, that the threat will succeed.

Events are threat-vulnerability pairs that lead to unwanted consequences. They include successful malicious attacks on the part of adversaries, natural occurrences such as hurricanes or floods, and unintentional errors.

Likelihood is the probability that an event will occur. It is a function of both threat and vulnerability. It is the name we attach to the joint probability that not only does a vulnerability exist that can lead to harm, but also that the specific threat that will exploit that vulnerability is present. Applying the rule for calculating joint probabilities,² likelihood can be expressed as:

$$\textit{Likelihood} = \textit{Threat} \bullet \textit{Vulnerability}$$

Consequence expresses the impact, either harm or loss, associated with an exploited vulnerability. The impact could be economic, political, military, social, psychological, or any combination of them. The most common, and usually the easiest, method of expressing consequence is in monetary form, even if the impact is other than economic.

The concept of risk combines that of consequence with that of likelihood. *Risk* is a measure of the expected negative effect of a particular unwanted event. It can be expressed as a product of the likelihood of the event and the consequence or impact should that event occur. Likelihood,

² For a discussion of joint probabilities, and how threat and vulnerabilities combine mathematically, see Appendix A.

then, is the probability that the particular unwanted event will occur, and, as described above, is a function of vulnerability and threat.

$$Risk = Likelihood \bullet Consequence$$

One common method of expressing risk is as expected loss. Here, consequence is expressed in monetary form, such as the value in dollars that could potentially be lost, and likelihood is expressed as the probability that this consequence or monetary loss will actually occur.

3. UNDERSTANDING ASSURANCE

In the paper, “What Color is Your Assurance” [WICH95], assurance is said to be “something said or done to inspire confidence.” But one can inspire confidence in a *statement* about a thing or about the *thing* itself. The notion of assurance deals with the confidence we have in the statement, not the thing itself. In this security-assurance realm then, there are really two questions:

1. How secure am I?
2. How confident am I that my assessment of security is accurate?

The first question deals with the notion of security, while the second deals with the notion of assurance. Assuming that a risk assessment is used to answer the security question, assurance addresses the confidence that one has in the risk value produced by that assessment.

3.1 Estimating Risk

Risk measurements have a great deal of uncertainty associated with them, which results from the fact that someone must select values to insert in various elements of the risk equation or tool, such as the likelihood of a certain event’s happening within the next year, or the frequency of use of a particular component. The user does not know *exactly* what numbers to select, since there is no way to know. Instead, the user must make an estimate based on her best judgment. Some assessors may tend to be overly pessimistic, resulting in calculated risk values that are too high, while others may tend to be overly optimistic and yield risk values that are too low.

Since the estimates are presumably based on the best information available at the time of the analysis, the numbers supplied may be more or less accurate, depending on the quality and completeness of that information. Therefore, the uncertainty associated with the calculated risk value has a great deal to do with the amount and quality of information available to those supplying the input data. Given a greater amount of high quality information, the risk analysis can produce a result with a narrower uncertainty.

3.2 Assurance

Rather than arbitrarily simplifying the risk picture by accepting a single value for risk, it is possible to manage the uncertainty associated with it in a way that enables a better understanding of the security issues involved. To accomplish this, we have cast the definition of assurance in risk or security terms. For the purposes of this paper, the following definition, presented in Section 1.2, will apply:

**Assurance is a measure of confidence in
the accuracy of a risk or security measurement.**

This is a qualitative definition. Later in the paper, we will give assurance a quantitative definition as well. By expanding the traditional notion of assurance, this definition is broader than many previous definitions. However, this definition is also more specific, since it specifies exactly how

assurance is linked with measurement. In any case, we believe that this definition can be reconciled with previous attempts to define assurance.

The definition expands the traditional notion of assurance in two ways. First, this approach accommodates the effect of “negative” evidence. Negative evidence is defined as information that tends to establish the existence and magnitude of specific risks to the system. All evidence tends to reduce the uncertainty (increase the confidence) in a risk estimate by providing more accurate information to assessors. Unlike positive evidence, however, negative evidence tends to result in higher risk estimates. Examples of such evidence include successful penetrations, criminal records, or OSHA violations.

A second way this definition expands on traditional notions of assurance is that it applies to anything that is considered a security risk, rather than being limited to correctness, confidentiality, or other particular property. This approach means that assurance can be used to help understand security in commercial, Internet, and other environments where security risks may not fall neatly into one of the generally accepted security categories.

3.3 Risk and Assurance

The underlying mathematical model that describes the relationship between risk and assurance is taken directly from probability theory. Let R denote the “real” value of risk, and let R^* denote the estimated value of R resulting from a risk analysis. In order to obtain a measure of the precision of our estimate, R^* , one might attempt to find two positive numbers, δ and ϵ , such that the probability that the true value, R , is included between the limits $R^* \pm \delta$, is equal to $1 - \epsilon$.

$$P(R^* - \delta < R < R^* + \delta) = 1 - \epsilon$$

For a given probability, $1 - \epsilon$, high precision of the estimate would then obviously be associated with small values of δ . The interval $R^* \pm \delta$ is called the “*confidence interval*” or “*uncertainty*” of R , and the probability, $1 - \epsilon$, is denoted as the “*confidence coefficient*” of the interval [CRAM55]. Theoretically, assurance could be expressed as a (δ, ϵ) pair, as a confidence interval given a particular desired value of ϵ , or as a confidence coefficient given a desired value of δ . However, as a practical matter, we have found that expressing it as a confidence interval is the easiest and the most intuitive.

For example, one could say that there is a 95% probability that the overall risk associated with a particular enterprise, expressed as expected loss, is between \$4 M and \$7 M, or in other words, $1 - \epsilon = 0.95$, $R^* = \$5.5$ M, and $\delta = \$1.5$ M.

Although this model combines the uncertainty from threat, vulnerability, and consequence, it is possible to consider them separately. For example, it may prove useful to distinguish between the uncertainty associated with likelihood and that associated with consequence. This is discussed later in the paper.

3.4 Investigating the Uncertainty

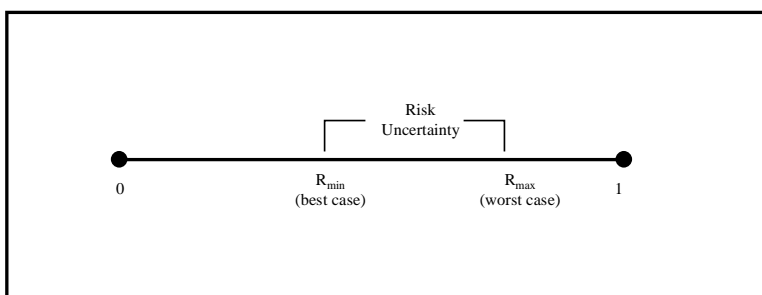
The model described above is useless if there is no way to figure out what the uncertainty actually is for a particular risk. Fortunately, there are several techniques that can help us gain an understanding of the interval, if not precisely measure it. We can employ existing risk assessment tools themselves to assist in developing a value for the uncertainty associated with their own measurements.

For each independent source of risk, available risk analysis tools allow a risk value to be computed. Additionally, for each source of risk, evidence can be amassed to generate estimates of the degree of confidence or assurance one has regarding the calculated risk value. For example, if one were interested in the risk from a data-mangling virus, he could develop estimates for the impact or adverse consequence of such a virus, the probability that the system would encounter such a virus, and the probability that, if encountered, the virus would result in the adverse consequence. He could also accumulate evidence to support his claim that the above estimates were sound, i.e., that they derived from a sound basis.

The following methods are examples of ways to determine a *confidence interval* and *confidence coefficient* for a particular risk. Most organizations trying these experiments are likely to find that there is quite a large uncertainty associated with their risk measurements.

3.4.1 Best-Case Worst-Case Comparison Method

Figure 2 shows the best case and worst case estimates for R as a normalized value plotted on a scale from 0 to 1, where 1 represents the maximum possible consequence. The uncertainty is depicted as the distance between R_{min} and R_{max} .



**Figure 2: Determining Uncertainty
Using Best Case-Worst Case Approach**

One way to quantify the range of uncertainty is to perform a “best case” and “worst case” risk analysis. The “best case” analysis will produce a minimum value for risk. Assuming a constant value for consequence,³ this involves assuming very low figures for the probability of attack (threat) and for likelihood of success (vulnerability). In this paper, we will refer to this value as R_{min} . Similarly, a “worst case” analysis will produce a maximum value for risk, or R_{max} , by using

³ In the development of this method, we consider only the uncertainty associated with likelihood. The uncertainty associated with consequence is treated in Section 3.7.

high likelihood values. Best case assumptions might include, for example, that all the mechanisms work as advertised, that the number of threat agents is low, that the consequence is low, and that the probability of success is low. Worst case assumptions would assume the opposite.

This approach is very easy to perform since existing risk tools can be used to obtain the necessary best and worst case measurements. One needs only to compute a value for risk twice, once with best case assumptions and once with worst case assumptions.

3.4.2 Repeated Measurements Method

Another way to assess the uncertainty is to take repeated measurements and calculate the standard deviation. This approach is more suitable when risk assessments can be repeated frequently on targets that are similar. For example, in products that are being built on an assembly line, the organization has the opportunity to make repeated quality measurements, gathering a large quantity of data. By calculating the standard deviation of these measurements, it can generate a value for the uncertainty associated with the defect rate of their product. The practical effect of employing the standard deviation is to fix the confidence coefficient, $1 - \epsilon$, at a particular value, i.e., 95 percent.

This method might be used, for example, in the process of selecting a firewall. By gathering information about vulnerabilities or flaws found in various firewalls, the likelihood of finding such a flaw in each product could be estimated. But the confidence interval could also be determined by exploring the standard deviation of those measurements. If the confidence interval is too large, the measurement is not of much use. But if enough information can be gathered to reduce the confidence interval to a reasonable range, the risk measurement can become quite meaningful.

3.5 Normalizing Assurance

A way to obtain a normalized value for assurance would be to divide the risk confidence interval by the largest possible risk facing the enterprise. The largest possible risk would be that risk resulting from the largest possible consequence multiplied by the highest possible value for likelihood, which, as for any probability figure, is equal to “1.” Therefore the largest possible value for risk has the same value as the largest possible value of consequence, C_{Max} . Applying the best case-worst case method for determining assurance, the normalized value of assurance becomes

$$A_{Norm} = \frac{\text{Worst Case Risk} - \text{Best Case Risk}}{C_{Max}}$$

3.6 Relationship between Confidence Coefficient and Uncertainty

In applying the best case-worst case method described above, the resulting values for assurance can vary greatly depending upon who decides the input values for the two extreme cases. Each person will tend to apply a certain level of confidence in risk estimates, albeit perhaps unconsciously. This level of confidence, expressed as the confidence coefficient, $1 - \epsilon$, represents the probability that the “real” value of R actually falls within the confidence interval defined by 2δ . This interval defines a zone of uncertainty. One person, trying to be absolutely certain that R falls within this zone (equivalent to a value for $1 - \epsilon$ close to 100 percent), will choose to make the confidence interval very large, accepting a very large amount of uncertainty. Another person might be willing to accept a much smaller amount of uncertainty and a smaller confidence interval, effectively applying a much smaller value for $1 - \epsilon$.

This means that uncertainty (or confidence interval, 2δ) and the confidence coefficient, $1 - \epsilon$, are interdependent. To be absolutely certain that R falls within the confidence interval, the interval must include the entire possible range. Similarly, reducing the uncertainty by shortening the interval causes the confidence coefficient to decrease. At the limit, this means that zero uncertainty can only be obtained with zero confidence.

An example should make the relationship clear. First, suppose a company has determined that they are 95% sure that their expected loss to the enterprise will be between \$4M and \$7M. Management decides that they need to do something about security and decrees that they must have a more exact estimate of expected loss. So the risk assessors return with a value of \$5.26M, but express zero confidence in their answer. Whereupon management revises its direction and tells them to produce an answer with 100% confidence. In this case, the assessors determine that they are absolutely sure that the expected loss will fall somewhere between \$0 and \$100M—the latter figure being the point that their catastrophic insurance kicks in.

Obviously, the company must find an acceptable compromise. To make results comparable, the same confidence coefficient must be used across risk estimates. The advantage of explicitly specifying this confidence coefficient, is that people performing the risk analysis are much more likely to produce comparable numbers regardless of whether they are, by nature, conservative or optimistic.

The approaches described above have involved fixing the confidence coefficient, $1 - \epsilon$, at a particular level in order to enable comparisons among risks. Since the level of certainty can be described as a percentage, a natural choice might be the standard statistical significance level of 95%, which would be the value if the “repeated measurements” method were used.

Another possible approach for fixing the confidence coefficient is to give all the participants a qualitative verbal standard intended to represent a given level of certainty. The legal field has produced a number of such standards for evaluating arguments [WILL95b]. These include:

- Substantial Evidence (a considerable amount)
- Preponderance of the Evidence (more than the evidence against)
- Clear and Convincing Evidence (what a reasonable person would believe)
- Evidence Beyond a Reasonable Doubt (no reasonable person can doubt)

Such verbal standards may have more meaning to the people who are performing the risk assessment than a “percentage” of certainty. Since the particular standard chosen matters much less than the fact that everyone use the same one, the substitution of legal standards for quantitative ones should not affect the utility of the results.

3.7 Consequence Uncertainty

Earlier, we made the simplifying assumption that consequence was fairly well known, with little associated uncertainty, and have dealt principally with the contribution to risk uncertainty from likelihood. But consequence is often no easier to assess and assign a value to than likelihood, and the effect upon the risk calculation is exactly the same. In this section, we discuss the factors affecting consequence uncertainty and its effect on assurance.

Many factors contribute to consequence uncertainty. Since most of the more worrisome events have never occurred, the full consequences remain incompletely known. A fully described event (a threat-vulnerability pair) will typically specify the who, what, and how of an event, but would probably not specify a time or place. Yet, the consequence associated with an event can vary greatly depending upon when the event takes place. Time of day, time of year, temporal association with other enterprise occupations, can all greatly influence the severity of the same event’s consequence. For example, power outages or denial of service attacks would probably be felt much more by an enterprise during its normal working day than in the middle of the night. An enterprise whose business is seasonal would likely suffer greater consequences during their peak months than during the business’ off-season. And consequences to organizations could be much greater during critical periods—during a war for a military organization, during a custody battle for a large business, in the middle of a major competitive procurement proposal effort for a consulting firm, etc. By definition, the range of uncertainty (i.e., confidence interval) must include these variations.

In order to consider how the uncertainty surrounding consequence values might be reduced, it is necessary to understand how these values are calculated in the first place. Usually, consequence values are determined through the use of techniques that solicit and harmonize the independent judgments of a number of persons knowledgeable about the mission and survival limits of the enterprise. Knowledge of the mission is important because the evaluation of consequences should always be made from a mission perspective [JELE95]. Understanding the enterprise’s survival limits is also important, since, from the perspective of the enterprise itself, the “death” of the enterprise is usually considered the ultimate consequence. Considering the enterprise’s mission, these persons develop a set of consequences that they feel most threaten the enterprise. This list is then arranged in priority order. From the prioritized list, and considering specific threat and vulnerability data, a postulated set of events that would produce these consequences is then generated.

Prioritizing the consequences, employing some kind of consensus process, is usually not particularly difficult, but gaining agreement on a numerical value for this consequence can be quite daunting. Many of the most dire consequences are extremely difficult to evaluate. What value does one place on the loss of a human life, for example? At best, several assumptions have to be made in order to produce a meaningful number. Fortunately, for most purposes, it is not necessary to obtain a number that has any real world significance. In the majority of cases, some

arbitrary value, like “10,” can be assigned to the most serious consequence, and all other consequences can be assigned numbers relative to that arbitrary value.

The general approach that we have outlined for evaluating the uncertainty surrounding likelihood—namely, the best case-worst case approach—is equally applicable to judging the uncertainty surrounding consequence. Consider the statement, “The loss that would result from a total shutdown of an enterprise’s main computer network would be approximately \$1.5M per day.” To make such a statement, a number of assumptions, at least partly backed up by data, are required. Assumptions might involve such items as expected business per day, access to backup data, availability of key employees, etc. In order to produce a value for uncertainty, it would be possible to first evaluate such a statement under universally optimistic assumptions to generate a best case value, and then to evaluate the same statement under universally pessimistic assumptions to produce a worst case figure. The difference between them provides a measure of consequence uncertainty. And, assuming that one can compute best case and worst case values for both likelihood and consequence, the computing of a value for assurance amounts to multiplying the best case consequence and likelihood together to yield a best case value for risk, multiplying worst case consequence and likelihood together to yield a worst case value for risk, and then subtracting one from the other to obtain a confidence interval for risk, which is our measure of assurance.

4. RISK PLANE DISCUSSION

People make decisions about security risks all the time. Some decisions amount to “bet the company” choices, while others have much smaller potential consequences. These decisions are often made without considering all the factors because they are either unknown or simply too complex to understand. Organizations in this situation may have a “false confidence” that they are secure, when in fact this confidence is based on an inadequate understanding of the risks. This section describes how the analysis of assurance described above provides some of the information necessary to make a more informed decision about security risks.

There is, of course, a large body of knowledge about risk management, which deals with this issue in great detail. The point of this section is merely to show the role that assurance can play in the risk management process. The approach described here only touches on how the costs associated with investing in security safeguards factor into the decision to accept the risk or not.

A popular way of displaying different risk-causing events is by way of a *risk plane*. In a risk plane, unfavorable events are plotted on a two-dimensional graph in which consequence serves as one axis and likelihood as the other. Figure 3 also depicts how the consequence uncertainty interacts with the likelihood uncertainty to define an area of risk uncertainty.⁴




Figure 3: Combining Consequence and Likelihood Uncertainty

As typically employed, events are plotted as single points on the graph. In Figure 4, We have departed from that practice. We have chosen, instead, to plot six, very different, stylized event classes as rectangles whose vertical and horizontal dimensions represent, respectively, the

⁴ A slightly different method of displaying virtually the same data appears in John Carroll’s book on computer security. His book includes a figure containing what he calls a “Plane of Uncertainty,” in which severity uncertainty is plotted against frequency uncertainty. See [CARR95], p. 482.

uncertainty surrounding the consequence and likelihood of the event. For the purpose of this illustration, we will assume that, in each case, the consequence associated with each event is known fairly well, and that any uncertainty or lack of assurance is attributable to the likelihood factor. By displaying events in this manner, the appropriate action on the part of a decision maker is much clearer.

4.1 Case Studies

The following six cases discuss each of the six different possibilities for events depicted in the risk plane in Figure 4.

Case 1 represents a situation in which both the consequence and likelihood of the event are known, with high assurance, to be low. Any example of this case tends to sound silly as soon as it is voiced, precisely because it is highly unlikely and of low consequence, even if it were to occur. An example might be a passing comet. Since both the consequence and likelihood are known to be low, the risk is very low and no mitigation action is warranted.

Case 2 is similar to Case 1 except that, in this case, the likelihood of the event is high. An example of an event of this class might be an occasional power interruption in an enterprise that is very disciplined in its practice of backing up its files. Since the consequence is low, there is little incentive to spend very much on additional safeguards, but if there were a moderately effective safeguard that could be put in place at low cost, such as an AMPS, it might be worth doing.




Figure 4: Example Events Plotted on a Risk Plane

Case 3 is different. Here, as before, the consequence is low, but the likelihood could be anywhere. An example of this case might be a hardware failure of a key subsystem that is

covered by warranty. If the likelihood were known to be low, one would ignore the problem, as with Case 1, and if the likelihood were known to be high, as in Case 2, one would only invest in relatively inexpensive safeguards. One reasonable approach in this case would be to assume the worst case, act as if the likelihood were high, and invest in cheap fixes if there are any. Another defensible approach, particularly if there are no cheap fixes, would be to gather more evidence in an effort to shrink the confidence interval, i.e., raise the level of assurance. Which decision is the more appropriate would depend upon the relative costs of evidence vs. safeguards.

Case 4 presents a similar situation to Case 2, except that in this case, it is the likelihood that is low and the consequence that is high. An extreme example might be the risk of the enterprise being hit by a meteor. Clearly, if this happened, the result would be disastrous, but the odds of its occurring during the lifetimes of the next several generations is small. In this particular example, one would almost certainly take no action whatsoever, but in more normal examples, one might invest in some inexpensive safeguards.

Case 5 represents the situation in which the event in question is known, with high assurance, to be of high consequence as well as high likelihood, implying very high risk. An example of this situation might be a virus attack against an open, unprotected network containing all of the enterprise's information assets. In a case such as this, mitigation is clearly called for, and would be avoided only if the "fix" were either prohibitive in cost or technically unfeasible.

Finally, Case 6 represents the situation for which the consequence of the event is high but the likelihood is not known. Harm caused by a malicious insider is an example of this situation, since, clearly if there were one, the consequence would be quite high because almost any insider, if intending to do harm, can do a great deal. The problem is in knowing whether or not you have or will have such a person. In this case, mitigation efforts may be in order, either to reduce the probability that one exists—i.e., reduce the threat, or to limit the harm that such a person could perpetrate—i.e., reduce the vulnerability. But it probably makes equal sense, since the consequence and risk are so high, to expend some additional effort at narrowing the assurance interval and thus more precisely determine the true likelihood. These examples and the indicated actions for each case are summarized in Table 1. The possible actions are explored further in Section 5.

Case	Likelihood	Consequence	Assurance	Example	Action
1	Low	Low	High	Passing comet	Ignore
2	High	Low	High	Power interruption	Fix if cheap
3	Unknown	Low	Low	Warranty-protected HW failure	Get more information or fix
4	Low	High	High	Meteor	Fix if cheap
5	High	High	High	Unknown Virus	Fix if at all possible
6	Unknown	High	Low	Malicious Insider	Get more information or fix

Table 1: Summary of Different Risk Cases

4.2 Factors Involved in Decisions

This paper emphasizes the role that uncertainty plays in making the decision to take some action to reduce a risk. The uncertainty is by no means the only factor that must be considered. Cost, schedule, complexity, and even practicality are other considerations.




Figure 5: A Simplistic Approach to Risk Decisions

Figure 5 shows a very simplistic approach to making risk decisions for the purpose of describing the role of uncertainty in the process. A decision maker could simply select a value of R representing a threshold value above which risks become unacceptable. This results in drawing an imaginary diagonal line across the risk plane. Events below the line are ignored while events above the line are accorded some action (as described in Section 5).

Another approach is to prioritize the risks in terms of overall risk (calculated as the product of consequence and likelihood). The decision maker can then deal with as many of the most serious risks as possible and can look to other factors, such as cost, to decide whether or not to address them.

In this framework, however, events are not simply plotted as points on the risk plane, but as areas of uncertainty. In cases where the uncertainty rectangle crosses a predetermined threshold, as in Figure 5, there is no way to be certain as to whether the actual risk falls above or below the threshold. The process for comparing these risks with risks that have a smaller amount of uncertainty is discussed in Section 5.

5. DECIDING A COURSE OF ACTION

Consider the decision maker who must decide whether to accept a risk or make the investment to mitigate the problem. The usual approach for dealing with risks judged to be unacceptable is to identify an appropriate safeguard and implement it. In many cases, this action will reduce the risk to an acceptable level. However, there are some situations in which the uncertainty associated with a risk is so large that there is no way to tell what effect an additional safeguard would have.

In practice, these situations occur frequently. For example, consider a company that is concerned about its susceptibility to the risk of a virus attack. The company employs an undocumented shareware virus tool, so there is a great deal of uncertainty associated with estimates of their residual risk. If the company determines that the risk of viruses is still unacceptable, they are forced to do something.

If the risk can be expressed in monetary units, i.e., dollars, then the risk, which is now equivalent to expected loss, can be compared with the cost of the safeguards necessary to mitigate the expected loss. So long as the mitigation cost is less than the expected loss (probability of loss or likelihood, multiplied by the severity of the loss or consequence), then the reasonable decision is to invest in mitigation. When the mitigation cost is greater than the expected loss, the sensible decision is to accept the risk. It makes sense, then, to invest in safeguards until the expected loss is reduced to the point that it becomes less than the mitigation cost. All of this assumes that the decision maker has sufficient information on which to base her decision.

If this is not the case, the decision maker is left with two choices for trying to deal with the risk. She could add security mechanisms, such as well documented virus checkers, procedures, or training. Alternatively, she might decide to improve assurance by obtaining additional information about whether or not the risk is negligible or serious. Either method could result in a lower value for R^* , the calculated risk value. Adding a security mechanism could do this by raising at least the best case risk value; and adding assurance, by narrowing the gap between best case and worst case estimates. These two alternatives are further explained and contrasted below.

5.1 Adding Mechanisms

The most common way of dealing with unacceptable risks is to add another security mechanism, which attempts to reduce the likelihood or consequences of that event. This approach is certainly appropriate when the uncertainty associated with a particular event is very small. For example, if a disk drive manufacturer has measured the mean time to failure of the devices experimentally, there will be little question about the assurance related to the risk assessments, since the uncertainty is low.

In this type of situation, there is little value in gathering additional information to provide the risk assessor with a better basis of estimate. No amount of additional information will reduce the risk to an acceptable level. Therefore, adding a mechanism seems to be the best approach. Of course, adding a mechanism may introduce new uncertainty if there is little information about it, and this new additional uncertainty may negate any benefit gained by adding the mechanism.

5.2 Adding Evidence

As shown above, where there is little uncertainty in a risk estimate, obtaining more information offers no appreciable benefit. No amount of additional information will do anything about the risk itself. However, in many cases, the uncertainty associated with a risk estimate is huge, since the assessors do not have good information on which to base their estimates about threat, vulnerability, and consequence. This section discusses ways to reduce this uncertainty.

The uncertainty results from the assessor's lack of relevant information on which an estimate can be based. Therefore, the obvious solution is to seek additional information. The simple approach of assembling evidence already available can be quite inexpensive and yet provide a great deal of information to the assessor. If this is insufficient, if there is no available evidence, or if one has the luxury of planning an assurance strategy, the assurance argument approach discussed later in this paper may be useful.

Regardless of the way that the information is organized, it will tend to increase an assessor's confidence in his risk estimates. For example, a risk assessor provided with details of the track record of a firewall may decide that the worst case is not really as bad as previously thought. This lowers R_{max} , shrinking the confidence interval from the high end. This, in turn, shifts the midpoint of the confidence interval, R^* , representing the best estimate of overall risk, to a lower value.

However, increasing assurance is not always a cause for rejoicing. If a risk assessor is provided with "negative" evidence that a firewall has been successfully penetrated many times, he will conclude that his earlier best-case estimates were too high and lower them. This acts to shrink the confidence interval from the low end, raising R_{min} . In effect, he will be more certain that the risk is high. Although perhaps discouraging, this is useful information that helps the decision maker make an informed choice.

6. INGREDIENTS OF ASSURANCE ARGUMENTS

A logical way of assembling evidential data in order to derive values for risk and assurance is by way of an *assurance argument*. There are four elements that can be used to structure an assurance argument. Table 2 lists these elements and provides a brief description. Basically, assurance arguments consist of a set of claims based on a logical framework, supported by evidence, and bounded by a set of assumptions. These arguments are nested in the sense that each argument is composed of lower-level supporting arguments and evidence. The cycle of generating lower-level claims and supporting evidence continues until it is reasonable to assume the claim without further evidence. This stopping point for assurance arguments is called the *assumption zone*.

The assurance argument should be documented to form an “assurance package” delivered to the consumer. The assurance package might take the form of a brochure, fact-sheet, white-paper, security documentation, or certification package. Most products and systems should be accompanied by some sort of assurance package, although the packages may vary considerably in their level of detail.

Argument Element	Description
Claims	Statements that something has a particular property
Evidence	Empirical data on which a judgment can be based
Reasoning	Statements which tie evidence together to establish claim
Assumption Zone	Limit of an argument where claims are accepted without evidence

Table 2: Assurance Arguments Are Built from These Elements.

6.1 Claims

An assurance argument is based upon claims about specific properties of an enterprise component. *Claims* are statements that associate subjects with their attributes or properties. Specifically, a *claim* is a statement that something has a particular property. The reason that these statements are called *claims* is that they may or may not be substantiated by any evidence.

6.1.1 Subjects

Subjects are the things about which one wishes to make a claim. A subject could be the entire enterprise; it could be the process, the people, the environment, or the technology comprising the enterprise; or it could be subgroups of people, subprocesses, specific aspects of the environment, or specific components comprising the technology. There are a wide variety of ways to break down subjects, the choice among them being situation dependent. However, we have found that breaking subjects down into their supporting processes, people, environment, and technology is useful across a wide variety of subjects.

Process consists of any activities that establish, affect, or maintain the security of an enterprise. Examples of processes include clearing users for access to the system, escorting maintenance personnel, reviewing audit logs, releasing magnetic media, scanning the system for viruses, using the system, administering the system, handling written logs, monitoring the system, managing the configuration, and assessing risk. Processes that are controlled, defined, mature, optimized, and repeatable, are much more likely to contribute to better security and lower risk. Evidence that can support process arguments includes process documentation, process metrics, and past performance information.

People include users, administrators, maintenance personnel, security officers, operators, organizations, and anybody else who could affect the security of the enterprise. People who are capable, experienced, knowledgeable, reputable and trustworthy are considered to be much more likely to perform without introducing security vulnerabilities. Thus, evidence about the education, training, past performance, experience, and background can be very useful in supporting arguments about people.

Environment includes geographical location (e.g., country, terrain), structural considerations (e.g., doors, windows), the physical setting (e.g., locks, protected network hardware, and locked computer rooms), and the organizational culture. An environment that is controlled, quality-centered, and stable is more likely to reinforce a security focus. Evidence that can establish these properties includes site security plans, physical architectures, blueprints, structural design analysis, and results of physical penetration tests.

Technology refers to the combination of hardware, software, and communications that is used to automate enterprise processes. Examples of technology might include cash registers, access control software, encryption devices, networks, file servers, trusted workstations, or word processors. Some useful properties for technology include that it be documented, correct, fault tolerant, and tested. Relevant evidence might include schematic diagrams, certification reports, architectures, test results, problem reports, and testimonials.

Supporting arguments can also help establish claims for the entire *enterprise*. *Enterprise* is the term we use to represent the composite of the people, process, environment and technology. Many different claims may be made about an enterprise. Some of the most security relevant of these claims are analyzability, correctness, completeness, and strength. Analyzability implies that the enterprise is not overly complex and is structured such that it can be understood. A correct enterprise indicates that the enterprise accurately performs as specified. Enterprise completeness indicates that all threats to the enterprise are addressed and all the security policies are implemented. A strong enterprise is capable of withstanding attack. Supporting evidence could include corporate security policies; organizational charts; and historical data, including results of previous enterprise-wide risk assessments.

6.1.2 Predicates

In addition to subjects, claims also contain *predicates*. *Predicates* are the things one wishes to say about the subjects. The predicate can contain an attribute or property assigned to the subject by the statement or claim, or it can address some specific threat or vulnerability, such as indicated by the statement, “The environment has been rendered completely safe from fire.”

At the highest level, the claim that one would wish to make about any subject would address risk directly, i.e., “The risk associated with the subject is less than some value (basically the amount one is willing to lose).” In those cases for which the application is unknown, the claim statement is not able to address risk, but would address likelihood instead, i.e., “The likelihood of a loss from this subject is less than some amount.” The difference between these two statements is that the first considers consequence and the second one does not. Since, in the second case, the application is not known, the consequence associated with various events can not be assessed.

In order to determine a value for risk, one must determine consequence, and in order to place a value on consequence, one has to know the mission of the enterprise. Another way to state the difference between the two above statements is that the first is application dependent, whereas the second is not. If and when the specific application is known, and the mission of the enterprise understood, the first statement would be the more appropriate. In those cases where the subject about which the claim is being made could have many different and unknown applications, such as a security product under development, the second statement is the more appropriate.

One form of lower-level predicates associates specific properties with a given subject to form a claim. We call these “*property predicates*.” A *property* is a characteristic trait relevant to establishing assurance. Table 3 lists some examples of security relevant properties. This list is not intended to be exhaustive. Not all properties apply to all components of an enterprise. Also, these properties may be interpreted differently when they are used to describe different things. For example, *strength* takes on a slightly different meaning when it refers to technology than when it refers to an environment.

Claim predicates can also involve the negation of a risk event. We call these “*risk event predicates*.” An example might be, “The software has not been subverted.” As with claims involving properties, those involving risk events can be substantiated with positive or negative evidence, or with supporting subclaims.

Properties	Description
Analyzable	Capable of being checked
Attentive	Alert, vigilant, observant, watchful
Capable	Having required or wanted skills or faculty
Complete	Having all the necessary parts or providing a total solution
Consistent	Uniform and steady
Controlled	Kept within defined limits
Correct	Free from error, defect, or fault with respect to a higher level specification
Defined	Described by a fixed set of parameters
Documented	Committed to writing
Easy-to-Use	Capable of being put in service, performed, or maintained without difficulty
Effective	Produces the desired result
Efficient	Performs with a minimum of waste, expense, or unnecessary effort
Experienced	Having performed similar events or activities
Fault Tolerant	Tolerant of mistakes or errors
Knowledgeable	Possessing requisite information
Learning	Capable of acquiring knowledge as result of experience
Managed	Operates according to a plan; an organized effort
Mature	Well seasoned; time-tested
Measurable	Capable of having dimensions, quantity, or capacity ascertained
Optimized	Designed and built with efficiency in mind; fine tuned for performance
Predictable	Anticipated, expected, foreseen, prepared for
Profiled	Described in a way prescribed by a standard criteria
Quality Focused	Very concerned about quality issues
Rated/Evaluated	Tested against a standard
Recoverable	Able to be repaired or brought back from harm
Repeatable	Capable of being performed, experienced, or produced again
Reputable	The estimation in which a person or organization is held by the public
Robust	Able to continue; resistant to undesirable change
Scaleable	Scope or granularity can be adjusted to meet changing circumstances
Stable	Unwavering; not subject to excessive variation
Strong	Capable of enduring or being defended
Successful	Possessing a high rate of past success
Tested	Subjected to a regimen of testing
Trustworthy	Deserving of confidence that a responsibility will be fulfilled
Well Understood	Universally comprehended across the entire enterprise

Table 3: Examples of Assurance-Relevant Properties

6.1.3 Subordinate Claims

A claim about a particular subject can be backed by supporting claims about the same subject or by similar claims about the various components that comprise the subject. For example, the claim, “The enterprise has a mature process for verifying and validating security,” could be backed up by the subordinate claims that:

- the process employed is defined,
- the process is analyzable,
- the process is measurable,

- the process is complete, and
- the process is tested.

Alternatively,⁵ the same claim could be supported by the subordinate claims that:

- the process for identifying verification and validation targets is mature;
- the process for defining a verification and validation approach is mature;
- the process for performing verification is mature;
- the process for performing validation is mature; and
- the process for providing verification and validation results is mature.

The first method holds the subject constant and decomposes the predicate; the second holds the predicate constant and decomposes the subject. Although it is possible to decompose both the subject and the predicate in the same decomposition step, this is not recommended. For more information on building an argument, see Section 7.

6.2 Evidence

In order to substantiate claims, we assemble information that helps to demonstrate their truth. *Evidence* is empirical data on which a judgment or conclusion can be based. Anything that contributes to the believability of a claim can be considered as evidence. Good evidence tends to be measurable, repeatable, and testable. Design analysis results are an example of evidence that helps to support a correctness claim. Other examples of evidence include analysis results, design documentation, or background investigations. Even circumstantial evidence can contribute to the believability of a claim, even though it may not be directly related. For example, the qualifications of the designers would constitute circumstantial evidence of the quality of the design.

Evidence can be consolidated. For example, the Trusted Product Evaluation Process (TPEP) considers a great deal of evidence and produces a product evaluation rating that summarizes it in a standardized way. This sort of packaging can be extremely helpful in reducing the quantity of evidence presented to consumers.

As with subjects, evidence can be thought of as having properties. Some examples of evidence properties are correctness, completeness, and analyzability. Perhaps the most important property of evidence is its relevance to whatever argument one is attempting to make. It may be 100% correct, absolutely complete, totally analyzable, and there may be large amounts of it, but if it has little or no bearing upon the claim that one is trying to make, it is of little use.

Evidence is generally of one of three types: descriptive documentation, analytic results, or historical data. The first type consists of documentation regarding the way a subject is intended to work. Examples of this type of evidence include business process models, architectures, plans,

⁵ This particular decomposition is based upon the list of base practices for Process Area 3, “Verify and Validate Security” as presented in the *Systems Security Engineering Capability Maturity Model* [SSE-CMM97].

and designs. A second type of evidence is produced by analysis. This type of evidence can be extremely strong; depending on the degree of objectivity, the method used to evaluate, and skill of the analysts. Some examples of analysis evidence include certification reports and security audits. Finally, the third type of evidence demonstrates how the enterprise has performed in the past. Examples of this type of evidence include security metrics, financial information, and customer satisfaction indices.

As was noted earlier, evidence can be positive or negative, in the sense that it can help negate or confirm a risk. Either type, however, reduces the amount of uncertainty surrounding a risk estimate and thus contributes to added assurance. Added assurance is not always reassuring.

In many cases, a single piece of evidence can support multiple claims. Examples might be a system profile under the Common Criteria, or an assessment and rating of an organization's processes under the SSE-CMM. Both of these constitute a number of strong statements about their respective subjects.

Evidence can take the form of a proof but generally does not. More often, it requires a judgment on the part of the receiver as to whether it is relevant, credible and sufficient to substantiate the claim. Unfortunately, these qualities are not measurable, and what one person would accept as relevant, credible and sufficient, another might not.

6.3 Reasoning

Reasoning is a set of statements that ties together the evidence and supporting arguments to establish a claim. An assurance argument is not a mere collection of the evidence to support a claim. For example, to establish the claim that an operating system is correct, simply providing a formal model and some user documentation is not enough. Instead, the reasoning must point to the evidence, demonstrate its relevance to the claim, and show why all claims have been sufficiently supported.

Assurance derives from the reduction of uncertainty surrounding claims. As uncertainty is reduced, assurance increases. A preponderance of evidence does not necessarily establish assurance claims. The evidence must be shown to be relevant, compelling, and cohesive. To create an assurance package, it is necessary to piece together many pieces of evidence. This process of matching evidence to claims can be very complicated. It may take many different types of evidence to sufficiently establish a claim.

6.4 Assumption Zone

As claims are successively backed by evidence and subordinate claims, which, in turn, are backed by more evidence and subordinate claims, eventually a point is reached such that the evidence is unassailable or the claims are accepted without further justification. This point, called the *assumption zone*, represents the termination of the successive decomposition of claims. The assumption zone represents the point at which claims are accepted without additional evidence and go unchallenged.

To illustrate this point, imagine:

- claiming that a system is free of flaws,

- then claiming that the operating system used is free of flaws,
- then claiming that the operating system design is free of flaws,
- then claiming that the operating system designers were qualified,
- then claiming that the résumés of the operating system designers are accurate,
- then claiming that the hiring process is rigorous,
- then claiming that the hiring process is documented,
- then claiming that the hiring process documentation is complete...

In this example, the documentation of the hiring process is not likely to have a significant bearing on the security of the product or enterprise. For most consumers and producers, evidence to support this last claim lies deep within the assumption zone.

7. STRUCTURING AN ASSURANCE ARGUMENT

Assembling all the pieces of evidence at all different levels of abstraction into a logical argument can be quite challenging. The approach advocated in this framework is to structure an assurance argument in a hierarchical manner. A major objective when constructing an assurance argument is to reach the assumption zone as quickly and efficiently as possible.

We have seen that claim statements form the building blocks of an assurance argument. In this section, we will examine the specific ways in which these claim statements can be composed and decomposed. In general, the composition and decomposition of claims can be achieved in two ways—by subject or by predicate. Subjects can be decomposed relatively simply into people, process, environment, and technology. This decomposition of subjects, we call the *subject tree*. Predicates, on the other hand, can be decomposed by properties or by risk events to form the *predicate tree*.

Building an assurance argument involves working down through both the subject and predicate trees to construct a hierarchy of claims.

7.1 Choosing a Top Level Claim

The top level claim is the root of an assurance argument. All of the evidence, subclaims, and reasoning go towards establishing confidence in this claim. The subject of this claim should, therefore, be appropriately broad and be as inclusive as possible. The security of this top level subject is the overall goal of the assurance argument. Examples include an enterprise, a product, or a work force. Anything could be the subject of a top-level claim, but the artificial narrowing of this top-level subject, such as addressing only the automated system, should be avoided.

The predicate for a top-level claim should also be extremely broad, to be defined by lower level claims. A broad property, such as “security” or “minimal risk” should be claimed. The subclaims, if complete, will define the scope of this predicate.

7.2 The Subject Tree

Frequently, an enterprise is far too large and complex to deal with directly. Risk analyses, therefore, will usually break the enterprise up into smaller enterprises, e.g., divisions of a company, agencies of a governmental department, colleges or campuses of a university, etc. Each of these component enterprises consists of people, processes, environment, and technology. And, in addition to enterprises consisting of sub-enterprises, subjects at the same level within an enterprise are made up of component subjects. Technology is made up of products; processes are made up of subprocesses, etc. This decomposition into component enterprises and component subjects can continue for several levels, but at each level, just as at the top enterprise level, we can talk about the components of risk and of assurance at that level. This decomposition through different levels constitutes the subject tree.

Figure 6 shows how a top-level subject can be decomposed into people, process, environment, and technology for lower level claims. At the next level of the argument, each of these subjects

could be supported by evidence or could be viewed as another top-level subject, which can be further decomposed.

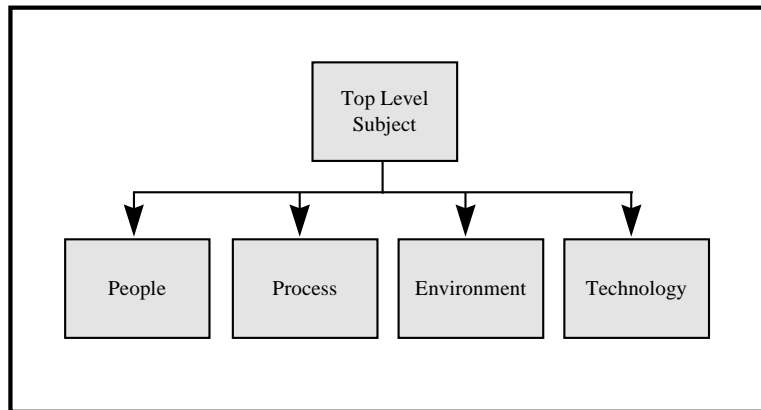


Figure 6: Subject Tree

7.3 Evidence Claims

Evidence can also be viewed as a subject of a claim, and can be supported by claims about the people, process, environment, and technology used to develop or analyze the evidence. This is depicted in Figure 7. Evidence claims can be made about the properties of the evidence, which can then be supported by additional argument. For example, evidence of a particular rating from an evaluation process could be supported with some additional claims about the process used, the expertise of the people performing the assessment, and the technology used to support the assessment.

Evidence claims are, by their nature, a step removed from the main argument. Nevertheless, if an evidence claim at a high level in an argument can be supported by strong evidence, it may be more important than a lower level claim about the enterprise itself.

7.4 The Predicate Tree

In Section 6.1, claims were defined as statements linking subjects with properties or attributes of those subjects. At the highest level, the property might be the association of a value of risk or likelihood of loss to those subjects. At lower levels, the property is usually some attribute like *complete*, *measurable*, *mature*, or *robust* (see Table 3).

As we have seen, the subject, *enterprise*, can be decomposed into people, process, environment and technology. Each of these can also be further decomposed into subunits of people, subprocesses, portions of the environment, or components. This decomposition of subjects, we have called the *subject tree*. Similarly, properties can be decomposed to form a *predicate tree* of subordinate properties. This decomposition of subjects and predicates can be repeated until a level is reached such that the claims can either be adequately supported with evidence, or are at or below the assumption zone.

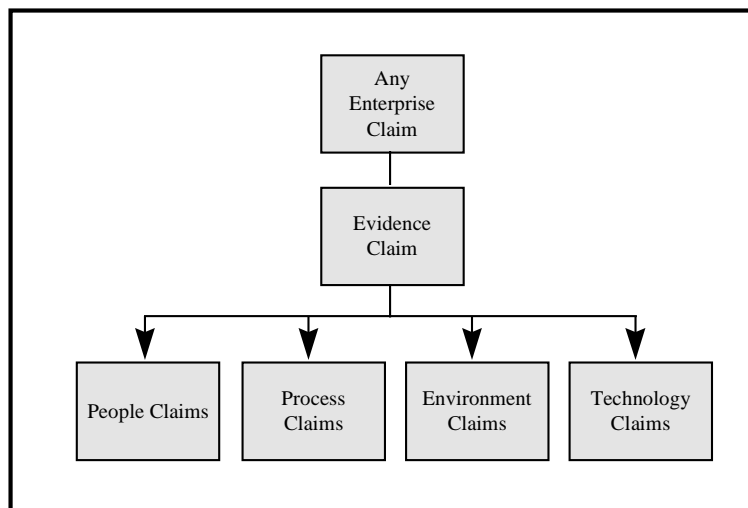


Figure 7: Evidence Claims

Some claims, even though they are backed by quantities of unassailable evidence and sound very positive and of value, can be irrelevant to the risks confronting an enterprise. For example, the claim, “A security process is mature,” would be irrelevant if the enterprise were staffed with persons paid by the company’s competition.

There are two fundamental challenges that confront any security product or system. The first is demonstrating that it does all that is supposed to do, i.e., it performs “at least” the actions stated in claims or specifications. The second is showing that it does not do anything that it is not supposed to do, i.e., it does “no more than” those actions; that is, it must not contain extra functionality or side effects that could be exploited by an attacker [BOEB92]. Whereas claims with property predicates can be used to support the “at least” type of requirement or higher level claim; claims with risk event predicates are required to support the “no more than” type.

Since the second type is far more difficult to support than the second, it is important that, among the set of nested claims, at least one of them specifically addresses the reduction of the likelihood of some event, which, if not dealt with, would threaten the enterprise. Examples of such claims might be “A system within the enterprise is not affected by power outages,” or “There are no industrial spies on the payroll.” These claims might prove a challenge for the evidence producer, but clearly address the likelihood of unwanted events—the first by negating a vulnerability, the second by denying a threat.

Such claims could be substantiated with sub-claims and might themselves be sub-claims to higher level or more sweeping claims. For example, the claim, “There are no industrial spies on the payroll,” might be supported by sub-claims, backed by evidence, that the backgrounds and financial interests of all employees are carefully checked every six to nine months, that all workplaces are under constant video surveillance, and that a cash award of \$100,000 awaits any employee who is able to supply evidence of another employee’s industrial espionage. At the same time, the same claim, “There are no industrial spies on the payroll,” helps to support the

higher claim that key industrial secrets are adequately protected from disclosure to competing companies.

Predicates can take several forms. In addition to previously mentioned “property predicates” and “risk event predicates” (see Section 6.1.2), predicates can address event components, i.e., threats or vulnerabilities, or they can address consequences that result from events. Consider the event: “a collapse of the roof of the main enterprise building due to an accumulation of snow.” The risk event claim statement might be, “The average risk (or expected loss) to the enterprise from a snow-caused roof collapse is less than \$100 per year.” This claim statement could be supported by the following set of sub-claims:

- The roof is capable of withstanding a load of 250 pounds per square foot without collapse (vulnerability predicate).
- The probability that the accumulation of snow on the roof at any time during a given year exceeds 250 pounds per square foot is 0.001 (threat predicate).
- The total loss to the enterprise that could result from a roof collapse is less than \$100,000 (consequence predicate).

Claims statements can be either quantitative, as in the examples above, or qualitative, e.g., “It is highly unlikely that the snow accumulation will ever exceed the load bearing capacity of the roof.”

Figure 8 is a state transition diagram depicting the process for decomposing a predicate for use in lower level claims. Starting with an event or property predicate, one can follow the arrows to find ways of structuring supporting claims. Each level of the argument should only use one type of predicate at a time. Continuing this process, each level of the argument can be structured according to a particular type of predicate. Each argument should, at some point, include an event predicate, to ensure that the argument is actually relevant to security and not an interesting but unrelated assertion.

Threat predicates are predicates concerning the likelihood of a particular threat occurring. As shown in Figure 8, once an argument relies on a set of claims which use threat as a predicate, there is no way to then support those claims by making claims with property, vulnerability, or risk event predicates. The only way to break down threat predicates further is by subject.

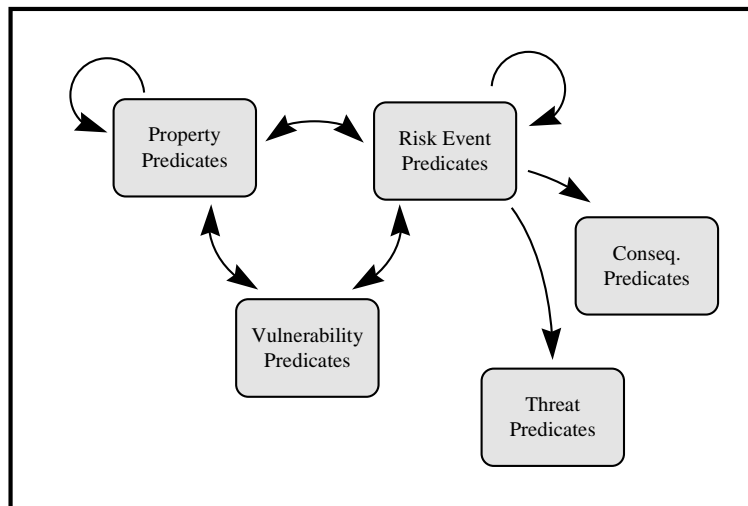


Figure 8: Generalized Predicate Decomposition

Similarly, *consequence predicates* are predicates that target the severity or impact of a particular event. Claims of this nature are the basis of consequence arguments, which are similar but not identical to the assurance arguments discussed above. Consequence statements are backed by evidence and subject to reasoning, and some statements could lie within the assumption zone, in the sense that they are accepted without evidence.

Consequence statements take the form, “The impact (or consequence) to a given subject resulting from a certain event would be ‘X’.” Although, theoretically, such a statement could be developed for any subject within an enterprise, it normally only makes practical sense to make such a statement for the *entire* enterprise. Thus, we will assume that consequence statements will be of the form, “The impact to the enterprise resulting from a certain event would be ‘X’.” The value to be assigned to “X” would be developed directly, based upon available data or evidence. It is generally not helpful to decompose the statement either through a decomposition of enterprise or through a decomposition of the event.

So, how does an enterprise achieve a better estimate of the impact of some untoward event on itself? Certainly, much of an organization’s ability to estimate this derives from its own experiences. Additional relevant “evidence” can be provided from the testimony of other similar enterprises that have suffered a similar event. And finally, consequence analysis is usually very amenable to modeling and simulation methods. Many “what if” scenarios, i.e., “events” in the parlance of this paper, can be modeled and their effects observed.

7.5 Structuring Principles

The process of composing and decomposing subjects and predicates will be difficult in practice, since there are many possible ways to create an argument successfully. For example, one could decide to structure an argument around the components of a system or, alternatively, to structure it around the high level risks to that system. While the structure of any particular argument will be heavily dependent on the details of the enterprise, several principles may guide the argument building process.

First, to avoid unnecessary redundancy, each of the subclaims should be as independent as possible. One way to achieve this is to hold either the subject or the predicate constant as one proceeds from one level in the argument to the one below.

Second, the number of subclaims should not be too high or too low. Given two otherwise equally attractive ways of decomposing a claim, the alternative with a manageable number of subclaims will make the argument easier to comprehend. In general, three to five subclaims will make a nicely structured argument. Fewer subclaims will result in an extremely “deep” argument which is difficult to follow. Similarly, more subclaims will make the argument so broad that it becomes impossible to keep track of all the relevant factors.

Another design principle for building good assurance arguments is that, in so far as possible, subjects and predicates should be at a similar level of abstraction. If high-level subjects are assigned low-level predicates, the resulting claim is difficult to understand. For example, the claim “The enterprise is not susceptible to short circuits,” could be improved by narrowing the breadth of the subject to only the automated system. Similarly, low-level subjects with high-level predicates are also difficult to handle. For instance, claiming that “the hiring process is secure” could be improved by narrowing the predicate to specify a particular attribute of the hiring process, such as “repeatable.”

Each level of the argument must be relevant to the claims above. This principle helps to ensure that effort is not wasted on sophisticated arguments that do not ultimately relate to the overall risk. Checking the relevance of each subclaim is a useful exercise as it helps to increase the understanding of the importance of the argument to overall assurance.

Finally, if the structure of the available information suggests a particular decomposition, following that structure in the argument is likely to be easier than attempting to restructure it. This is simply a pragmatic suggestion that takes advantage of the existing information and its structure. For example, if an organization has achieved SSE-CMM Level Three in some process area, implying well-defined security practices,⁶ it may have assembled a considerable amount of information organized according to the process area that produced it. In this case, it may make sense to structure the claims by process area. The evidence will then naturally support the claims, and the argument will be easier to understand, which is the goal.

⁶ In the context of the SSE-CMM, to say that an organization’s security practices are “well-defined” means that the organization has approved, tailored versions of standard, documented processes, which are used throughout the organization. See the *Systems Security Engineering Capability Maturity Model* [SSE-CMM97].

8. COMPOSING AND DECOMPOSING RISK AND ASSURANCE

This section describes the relationship between risk and uncertainty (and by extension, assurance) at different levels in an assurance argument.

8.1 Composing Risk

Risks do not simply add as one proceeds up the assurance argument. This is particularly true when the subject is technology. The risk associated with technology is not the mere sum of the risks of its subsystem components. Some risks are likely to be mitigated and new risks—the result of interconnectivity or complexity—are likely to be introduced.

For example, consider the risk associated with a system that involves the connection of two local area nets under the control of some network management software and operated under some procedures tailored to the specific installation. We will assume that each of the two local area nets as well as the network management software had been the subjects of risk analysis and that their risks were assessed to be R_A , R_B and R_S , respectively. We will further assume that the combination of tailored procedures and the network management software mitigates somewhat the specific risks internal to the two local area nets by a factor R_M , and the interconnecting of the two nets also added a complexity risk factor, R_C . We can then express the total risk to the interconnected system, R_T , as:

$$R_T = (R_A + R_B - R_M) + R_S + R_C$$

which, depending upon whether R_M or R_C is larger, would be greater than or less than the simple sum, $R_A + R_B + R_S$.

As indicated above, in order to deal with assurance, all risk or likelihood values that contribute to total risk or likelihood must be computed twice—first to yield a best case figure, and second, to yield a worst case number. Therefore, in the example just presented, minimum and maximum values for R_T could be calculated in the following manner:

$$R_{Tmin} = (R_{Amin} + R_{Bmin} - R_{Mmax}) + R_{Smin} + R_{Cmin}$$

$$R_{Tmax} = (R_{Amax} + R_{Bmax} - R_{Mmin}) + R_{Smax} + R_{Cmax}$$

The assurance interval is then the difference $R_{Tmax} - R_{Tmin}$. This is equal to:

$$[(R_{Amax} + R_{Bmax} - R_{Mmin}) + R_{Smax} + R_{Cmax}] - [(R_{Amin} + R_{Bmin} - R_{Mmax}) + R_{Smin} + R_{Cmin}]$$

which can be written as:

$$R_{Amax} + R_{Bmax} - R_{Mmin} + R_{Smax} + R_{Cmax} - R_{Amin} - R_{Bmin} + R_{Mmax} - R_{Smin} - R_{Cmin}$$

8.2 Combining the Risks From Different Events

In Section 5.2, we saw that the addition of evidence, whether positive or negative, acts to modify the value of R^* and to narrow the range of uncertainty associated with individual risk statements. For example, consider the statement

The expected loss from data entry mistakes on the part of an enterprise’s employees is \$1.5 million a year.

Whereas this statement represents valuable information, it hardly represents the total risk to the enterprise. The risk from this particular event must be combined with risks from other possible events to yield an overall statement of enterprise risk.

It was noted earlier that a value for risk can be obtained by multiplying the consequence (or loss) associated with some event by the product of the probability figures for the threat and vulnerability associated with that same event. To the extent that these events and their associated risks are independent and not mutually exclusive, which is generally true, the total risk can be expressed as the sum of the individual risks, i.e.,

$$R = (r_1 + r_2 + r_3 + \dots + r_n) \text{ where } r_1 = (c_1l_1), r_2 = (c_2l_2), r_3 = (c_3l_3), \text{ etc.}$$

$$\text{where } l_1 = (t_1v_1), l_2 = (t_2v_2), l_3 = (t_3v_3), \text{ etc.}$$

Combining these expressions, we get:

$$R = (c_1t_1v_1 + c_2t_2v_2 + c_3t_3v_3 + \dots + c_nt_nv_n)$$

Applying this technique, a figure representing the total risk associated with any enterprise, consisting of people, process, environment or technology, can be computed. As was pointed out earlier, in those cases in which the application is unknown, consequence can not be determined and the best that can be done is to compute a value for likelihood alone.

8.3 Combining Assurance From Different Events

In order to apply the best case-worst case method of computing assurance outlined in Section 3.4.1, all risk values that contribute to total risk (or, in cases in which the application is unknown, likelihood values that contribute to total likelihood) must be computed twice—first to yield a

best case figure, and second, to yield a worst case number. If estimates for best and worst cases are always given assuming a confidence coefficient of 0.95, results should be reasonably consistent.

For example, suppose that one were interested in the combined assurance associated with four independent events, E1, E2, E3 and E4, whose best case and worst case risks expressed in dollars were as contained in Table 4. Then, assuming that the risk contributions from the four events are independent and not mutually exclusive, the total best case and worst case risks to the enterprise are the sums of the individual best case and worst case risks as shown.

The numbers in the chart would produce values for R^* and δ of \$ 7.0 M and \$ 5.4 respectively. In the case represented by the chart, the level of assurance is not particularly high since, in order to have high confidence that the true value for risk is within the range $R^* \pm \delta$, the magnitude of δ must equal approximately 77% of R^* .

Event	Best Case	Worst Case
Event 1	\$ 0.5 M	\$ 3.5 M
Event 2	\$ 0.2 M	\$ 2.1 M
Event 3	\$ 0.1 M	\$ 1.2 M
Event 4	\$ 0.8 M	\$ 5.6 M
Total	\$ 1.6 M	\$ 12.4 M

Table 4: Combining Four Independent Risk Events

9. RELEVANCE AND UTILITY OF THE FRAMEWORK

In this section we describe the relationship between this paper and three different assurance related security efforts—the Network Rating Methodology, the System Security Engineering Capability Maturity Model, and the Common Criteria. By describing the types of claims that these efforts might support, we hope to start the discussion of how these different approaches to gaining assurance might be compared and combined.

9.1 Network Rating Methodology

Because of a similarity of sources, this paper contains passages that are quite similar to some in the Network Rating Methodology (NRM) document [NRM97]. This paper nevertheless differs from the NRM paper in two significant ways. First, whereas the NRM paper focuses exclusively on networks, the focus of this paper is on entire enterprises, of which information networks are a part. Second, the NRM attempts to define assurance as something that responds to “assurance need,” which the NRM implies can be quantitatively expressed. Specifically, NRM states, “A useful statement of assurance need should explain what the consumer is concerned about (breadth) and how much it will take to satisfy those concerns (depth).” The NRM goes on to state, “It is possible to state these needs in terms of evidence required or by indicating a desired level of assurance.” Stating assurance needs in terms of evidence required is not the same as indicating a desired *level* of assurance.

This paper takes the view that, although the consumer may be able to express her *qualitative* assurance needs, she will be hard pressed to express these needs quantitatively. What this paper presents is an alternative definition of assurance that enables one to quantify the *relative* effect of adding evidence by showing the quantitative effect that evidence can have in narrowing the uncertainty associated with risk and security measurements. Armed with a quantitative understanding, consumers can make informed decisions about whether their assurance needs have been satisfied.

The “matrix” approach adopted by the NRM is, although rigidly defined, quite consistent with the assurance argument structure presented here. Expressing the NRM in terms used in this paper, the top level claim might be:

The security provided by the network is “good enough.”

This claim is broken down into sixteen matrix cells representing predefined subclaims. These subclaims result from breaking the top-level subject (the network) into four areas and the top-level predicate (“good enough”) into four areas. The specific subject areas of the NRM are personnel, operational procedures, physical environment, and technology, which equate rather directly to the people, process, environment and technology subjects of this framework. The predicates of concern within the NRM paradigm are confidentiality, integrity, availability, and authenticity. The framework outlined in this paper, therefore, can easily accommodate the full repertoire of NRM-relevant claims.

9.2 System Security Engineering Capability Maturity Model (SSE - CMM)

The System Security Engineering Capability Maturity Model (SSE-CMM) [SSE-CMM97] “describes the essential characteristics of an organization’s security engineering process that must exist to ensure good security engineering.” Among its objectives is to generate confidence based on the maturity of processes used. It specifically addresses the satisfaction of “assurance needs.” The Model is intended to enable what it calls “capability-based assurance,” which it defines as “trustworthiness based on confidence in the maturity of an engineering group’s security practices and processes.”

A specific claim that could be based on SSE-CMM evidence is:

The security engineering organization’s process is mature.

This claim’s subject, the process, is broken down by the SSE-CMM into ten security-relevant process areas: Specify Security Needs, Verify and Validate Security, Provide Security Input, Assess Threat, Assess Vulnerability, Assess Impact, Assess Operational Security Risk, Build Assurance Argument, Monitor System Security Posture, Administer Security Controls, and Coordinate Security.

The claim’s predicate property, “maturity,” is broken down into various properties related to maturity, including planned, tracked, defined, coordinated, measured, controlled, and improving. The model also discusses evidence for supporting these attributes, including the appraisal method itself. Thus, a hierarchy of subclaims can be constructed according to the methods described in Section 7 of this paper.

These subclaims can then be used to support a wide variety of higher level claims about an enterprise’s security. For example, a claim that a component is free from defects can be supported by an SSE-CMM subclaim that validation process used in building the component was well-defined. Similarly, a claim that a component is resistant to penetration can be supported by the subclaim that the development process is planned and tracked.

One of the Model’s Process Areas is “Build Assurance Argument.” This Process Area is described as consisting of five “base practices,” considered by the Model to be “essential elements of good systems security engineering,” namely:

1. Identify the security assurance objectives
2. Define a security assurance strategy to address all assurance objectives
3. Identify and control security assurance evidence
4. Perform analysis of security assurance evidence
5. Provide a security assurance argument that demonstrates the customer’s security needs are met.

One objective of the present paper is that it be of help in organizing an organization’s efforts toward satisfying this Process Area (PA). It should also prove useful in assisting organizations

establish a connection between “Assess Operational Security Risk” and “Build Assurance Argument” and in showing how some of the same evidence could be used to satisfy both.

9.3 Common Criteria

“The Common Criteria (CC) represents the outcome of efforts to develop criteria for evaluation of IT security that are widely useful within the international community. It is an alignment and development of a number of source criteria: the existing European, US, and Canadian criteria (ITSEC, TCSEC and CTCPEC respectively)” [COMM97]. It is intended to serve as the initial step toward an internationally recognized standard and to open the way to worldwide mutual recognition of evaluation results. Version 1.0 of the CC was published for comment in January 1996. Version 2.0 is scheduled to be available in 1998.

The CC consists of three parts, “Introduction and General Model,” “Security and Functional Requirements,” and “Security Assurance Requirements.” Part 3 contains nine assurance categories from which the assurance requirements for a product or system can be chosen. The nine assurance classes are: “Configuration Management,” “Delivery and Operation,” “Development,” “Guidance Documents,” “Life Cycle Support,” “Protection Profile Evaluation,” “Security Target Evaluation,” “Tests,” and “Vulnerability Assessments.”

The top-level claim that can be based on the CC is:

The information technology has a specific set of security features and assurance evidence.

The particular set of features and assurance evidence is dependent on the particular CC profile used. While the CC contains some guidance for selecting features based on security objective and threat, there is no such guidance for selecting assurance evidence requirements. The top-level claim from the CC can be used in conjunction with other claims to support higher level claims about an enterprise.

The top level subject, “information technology,” can be decomposed as described above. However, the predicate, “has a set of security features and assurance evidence,” is more difficult to break down. Ideally, claims about features should be tied to the particular event that the feature is intended to counter. Similarly, claims about evidence should also be linked to the problem that the evidence addresses. Building the infrastructure that links these claims back to a top-level claim will involve difficult questions about the nature of the risk and the feature or evidence.

When a CC evaluation process is standardized, the results of that process would constitute excellent evidence to substantiate any CC-based claim. In essence, the evaluation stands in for the actual evidence.

9.4 Summary

Although the focus of each of the three described efforts is different, their goals are quite similar, and the framework outlined in this paper can be usefully applied within the context of all three.

The focus of the NRM is on networks; that of the SSE-CMM is on process; and that of the CC is on products or systems. All three, however, represent attempts at creating a meaningful way of describing and comparing the security efficacy offered by different entities. The NRM can apply the framework of this paper as a way of expressing and organizing claims made about the network under consideration that have resulted from the NRM matrix structure. For the SSE-CMM, the framework offers help in assembling evidence to satisfy the “Build Assurance Argument” process areas (PAs) and should prove useful in assisting organizations establish a connection between this PA and the “Assess Risk” PA, and in showing how some of the same evidence could be used to satisfy both. And to users of the Common Criteria, it offers both useful guidance for selecting and collecting assurance evidence to be used within a CC-based evaluation, as well as a means by which the evidence supplied by such an evaluation can be assembled with other relevant evidence into a cogent and consistent assurance argument for an entire enterprise.

Whereas this paper has not provided a means by which one can determine assurance need in the sense of some quantitative or even qualitative statement of the need, it does provide a means of deciding whether or not that need, whatever it is, has been satisfied. Additionally, it has been shown that the framework also provides a logical structure for assembling evidence into cogent arguments that can be employed within at least three very different current security efforts.

APPENDIX A - THREATS, VULNERABILITIES, AND JOINT PROBABILITIES.

If a set of events are mutually exclusive, the probability of one or another of the events occurring is the sum of probabilities of the events occurring individually. Thus, if events A and B are mutually exclusive,

$$P(A \text{ or } B) = P(A) + P(B)$$

If two events A and B are not mutually exclusive, then there is some probability that both can occur. The area of overlap is precisely the joint probability $P(A \text{ and } B)$ or $P(A, B)$. If $P(A)$ and $P(B)$ were simply added as above, this joint probability would be counted twice, since it is part of both $P(A)$ and $P(B)$. Therefore, for events that are not mutually exclusive, the above equation must be modified by subtracting this joint probability. We thus obtain the following

$$P(A \text{ or } B) = P(A) + P(B) - P(A, B)$$

where $P(A, B)$ is the joint probability of events A and B . If A and B are independent events,

$$P(A, B) = P(A) \bullet P(B).$$

With respect to threats, we are normally dealing with events that are not mutually exclusive.

Vulnerabilities are rarely mutually exclusive. They may not even be independent. However, since likelihood is the product of threat and vulnerability and since threats are generally independent, the products of the form $T \bullet V$ are also independent, even if the contributing vulnerabilities, v_1 and v_2 are not.

References

- [BOEB92] W. E. Boebert, "Assurance Evidence," Technical Report Contract 1021-02-91, Secure Computing Technology Corporation, 1 June 1992.
- [CARR95] John M. Carroll, *Computer Security, 3rd ed.* (Boston: Butterworth-Heinemann, 1995).
- [COMM97] "Common Criteria: An Introduction," produced by Syntegra on behalf of the Common Criteria Implementation Board.
- [CRAM55] Harald Cramér, *The Elements of Probability Theory*, (New York: John Wiley & Sons, 1955).
- [ISU96] Idaho State University, Glossary of INFOSEC and INFOSEC Related Terms, Vol. II, Version 6, dated August 1996.
- [JELE95] George F. Jelen, "A New Risk Management Paradigm For INFOSEC Assessments and Evaluations," *Proceedings of the 11th annual Computer Security Applications Conference*, (Los Alamitos, CA: IEEE Computer Society Press, 1995), pp. 261-267.
- [KAHN95] Jay J. Kahn, "A New Perspective on Combining Assurance Evidence," *Proceedings of the 11th annual Computer Security Applications Conference*, (Los Alamitos, CA: IEEE Computer Society Press, 1995), pp. 172-181.
- [NIST95] U.S., Department of Commerce, National Institute of Standards and Technology, *An Introduction to Computer Security: The NIST Handbook*, (NIST Special Publication 800-12), written by Barbara Guttman and Edward A. Roback, p. 92.
- [NRM97] "The Network Rating Methodology: a Framework for Assessing Network Security," dated September 11, 1997.
- [SSE-CMM97] Systems Security Engineering Capability Maturity Model, *Model Description, Version 1.1*, dated June 16, 1997.
- [WICH95] David R. Wichers, Joel E. Sachs, and Douglas J. Landoll, "What Color Is Your Assurance?"
- [WILL95a] Jeffrey R. Williams and Douglas Landoll, "A Framework for Reasoning about Assurance (Version 1.0)," ARCA Document Number ATR 95044, dated November 30, 1995.
- [WILL95b] Jeffrey R. Williams and Marvin Schaefer, "Pretty Good Assurance," *Proceedings of the New Security Paradigms Workshop*, (IEEE Computer Society Press, 1995).
- [WITAT95] National Institute of Standards and Technology, Invitational Workshop on Information Technology (IT) Assurance and Trustworthiness, March, 1995.