A Practical Approach to Measuring Assurance

George F. Jelen
G-J Consulting
gjelen@acm.org

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

Abstract[1]

Assurance has been defined as "the degree of confidence that security needs are satisfied"[2]. The problem with this definition is that, unless one has a way to specify security needs in some measurable way, assurance can not be expressed in a measurable way either. The definition leaves the practitioner with the challenge of determining what "security needs" are, whether or not they have been "satisfied," and how to determine "confidence." In this paper, we define assurance as ***"a measure of confidence in the accuracy of a risk or security measurement."*** *A critical feature of the view of assurance presented here is that it is orthogonal to the measurement of risk and security. High assurance ratings have traditionally been associated with high security and low risk. Our definition permits high assurance to be associated with low security and high risk as well. It also provides a way of deciding whether or not the assurance one has is sufficient.*

1. Introduction

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 opinion about what assurance is, how it might be measured and communicated, and how one might combine and compare the assurance gained from various sources.

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 take advantage of quantitative risk measurement methodologies and to employ them in such a way as to yield a rough measure of assurance that permits 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 exactly 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 then 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—that which this paper calls *assurance*.[2]

1.2 Overview

Assurance has been defined as "the degree of confidence that security needs are satisfied" [2]. The problem with this definition is that, unless one has a way to specify security needs in some measurable way, assurance can not be expressed in a measurable way either. The definition leaves the practitioner with the challenge of determining what "security needs" are, whether or not they have been "satisfied," and how to determine "confidence." In this paper we propose specific ways to interpret these

[1] This research was partially supported by the National Security Agency under Contract Number MDA904-97-C-0223. It borrows heavily upon an earlier study by the authors and published as an ARCA report [1].

[2] For a thorough treatment of the relationship between evidence and assurance, see [1].

terms in such a way as to make determining assurance more tractable.

In order to render the definition useable, we believe it is necessary to associate the terms in the definition with a measurement scale. For our purposes, the exact nature of the scale is not particularly important: the scale could employ numeric or "fuzzy" values, and it could be absolute or relative. For example, one might define "security need" as limiting annual loss expectancy from some cause to less than $100,000, or as restricting the extent of data corrupted by a virus to an amount no greater than that from user error. However, whether our statement of need is expressed in absolute (as in the first example) or in relative terms (as in the second), it ultimately requires a means and a scale of measurement. Determining whether or not virus-caused data corruption exceeds that from user error presumes some way of measuring both.

Once a measurement scale has been selected, "security need" can be expressed as a threshold value on that scale and a measurement of the actual level can be made. Simply comparing this measurement to the threshold value determines whether the need has been satisfied. However, the degree of confidence in this determination depends on the accuracy of the measurement. A large uncertainty in the measurement will impose a corresponding lack of confidence that the need has been satisfied.

In order to accommodate these factors, we offer the following definition of assurance:

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

This definition allows the concept of assurance to be associated with anything that can reasonably be considered a security or risk measurement method, and is not restricted to measures of specific properties like correctness or robustness.

In this paper, we argue that assurance is an integral part of the risk and security management processes. Our formulation is intended to be universal enough to embrace any assurance method. Within the security community, the heavy reliance on various criteria and evaluation processes to generate assurance has sometimes led practitioners to lose sight of their purpose, and therefore alternate methods for obtaining assurance have often been neglected. Because the method presented in this paper permits at least a rough measure of assurance, it is easier to understand and compare the costs and benefits of alternate sources of assurance evidence.

The paper will show how such a measure of assurance can be made by structuring the risk inquiry to yield information about uncertainty. We show how risks, along with their uncertainty, can be compared and evaluated. We then suggest ways of determining a course of action if the risks are found to be unacceptable.

2. Discussion

The view of assurance advocated in this paper assumes some means of measuring risk or security in which there exists uncertainty in the measurement. Since there will never be a way to measure either risk or security without introducing uncertainty, this approach to assurance should be universally applicable. In defining assurance in this way, we intend to imply nothing regarding the quality or the accuracy of any current risk measurement methods or tools. We merely note that such methods and tools exist and are in use. Our method will not and can not improve their accuracy, but it may raise the level of awareness in the uncertainty that surrounds their use.

2.1 A note on the examples

In the examples in this paper, since risk tends to be a more easily measurable quantity than is security, we have generally assumed that the measurements upon which the determination of assurance are based will be of risk rather than of security. 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, this is not meant to imply that a risk-based technique is required. As effective ways are found to "measure" security, the same approach to the calculation of assurance can be applied.[3]

2.2 Distinguishing assurance from security and risk measurements

In the paper, "What Color is Your Assurance," assurance is said to be "something said or done to inspire confidence" [3]. But one can inspire confidence in some given *thing* or in a *statement* about the thing. The notion of assurance advanced in this paper deals with the confidence we have in the statement, not the thing itself. In the security-assurance realm, there are really two questions:

1. How secure am I?
2. How confident am I of my answer to Question 1?

[3] There is a least one current effort, of which the authors are aware, aimed at defining measures or metrics for security. The Security Metrics Action Committee of the Process, Assurance and Metrics Working Group, part of the System Security Engineering Capability Maturity Model (SSE-CMM) initiative, is currently engaged in such an effort.

The first question deals with the notion of security, while the second deals with the notion of assurance as we define it. If 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.

A critical feature of the view of assurance presented here is that it is orthogonal to the concepts of risk and security. Keeping these dimensions separate helps to establish a clear distinction between them and thus reduces the confusion caused by any overlapping of their meanings. High assurance ratings have traditionally been associated with high security and low risk. Our definition permits high assurance to be associated with low security and high risk as well.

Confusing the assurance dimension with the security and risk dimension makes it very easy to miss the interesting cases where the two dimensions lead to seemingly conflicting conclusions. For example, consider a network component with a large number of security mechanisms, but no information to indicate whether or not they are correctly configured. Here, it is unclear whether the many mechanisms add or reduce the amount of assurance. By making the two orthogonal and independent, sources of added security can be considered separately from sources of added assurance. It is thus easier to recognize that it would probably be more useful to gather more specific information about the configuration 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, a security-based measurement would work equally well.

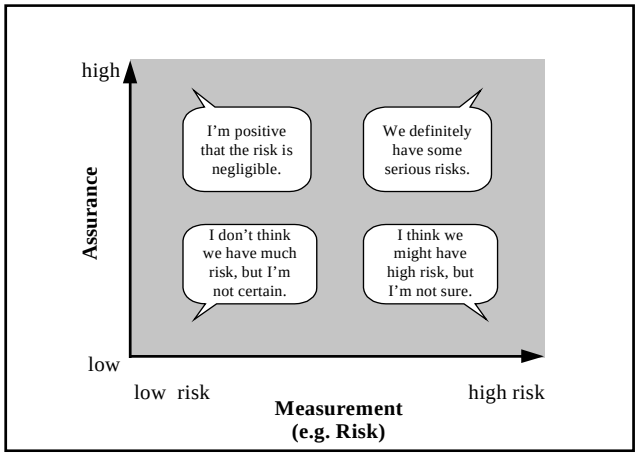

Figure 1: Assurance and Measurement Dimensions Are Orthogonal

2.3 Negative evidence

Our definition of assurance deliberately accommodates the concept of "negative" evidence. Negative evidence is information that tends to establish the existence and magnitude of specific insecurities of and risks to the system. Examples include such items as successful penetrations, criminal records, or OSHA violations.

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 generally results in higher risk estimates.

3. A Measure of Assurance

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

3.1 Risk definition

As defined in this paper, the amount of assurance ultimately depends upon the confidence in the accuracy of the measurement of the risk or security 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.

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" [4]. Expressed mathematically, *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 outcomes. 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 unwanted 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, likelihood can be expressed as:

$$Likelihood = Threat \bullet Vulnerability$$

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.

$$Risk = Likelihood \bullet Consequence$$

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.

3.2 Risk measurement uncertainty

Risk measurements have a great deal of uncertainty associated with them because 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 risk assessor cannot know *exactly* what numbers to select, and must instead make an estimate based upon his 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.

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.

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, indicating higher assurance.

3.3 A mathematical model

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 $1 - \varepsilon$.

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

For a given probability, $1 - \varepsilon$, high precision of the estimate would 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, $1 - \varepsilon$, is denoted as the "*confidence coefficient*" of the interval [5]. 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, $1 - \varepsilon = 0.95$, $R^* = \$5.5$ M, and $\delta = \$1.5$ M.

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, where a small confidence interval is representative of high assurance.

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 idea is discussed later in the paper.

3.4 Investigating the uncertainty

The model described above is useless if there is no way to determine what the uncertainty actually is for a particular risk. Fortunately, there are several techniques that can help us gain an approximation 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 demonstrate that the above estimates were sound.

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

3.4.1 Best-case worst-case comparison method. One way to quantify the range of uncertainty is to perform a "best case" and "worst case" risk analysis. The "best case" analysis produces a lower bound figure for risk. This involves assuming minimum estimates for consequence, for the probability of attack (threat), and for likelihood of success (vulnerability). In this paper, we refer to this value as R_{min}. Similarly, a "worst case" analysis produces an upper bound value for risk, or R_{max}, by using high consequence and likelihood estimates. 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.

Given a risk assessment tool or method, this approach is quite easy to perform since the tool or method 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.

Figure 2 shows the best case and worst case estimates of 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}.

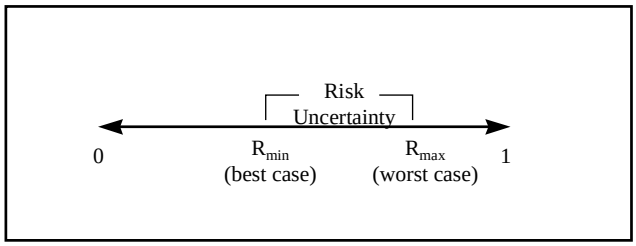

Figure 2: Determining Uncertainty Using Best Case-Worst Case Approach

3.4.2 Repeated measurements method. Another way to quantify the uncertainty is to take repeated measurements and calculate the standard deviation. This approach is more suitable when measurements of risk components (i.e., threats, vulnerabilities or consequences) can be repeated frequently on targets that are similar. For example, in products that are being built on an assembly line, an 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 the product. The practical effect of employing the standard deviation is to fix the confidence coefficient, $1-\varepsilon$, 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 encountering these vulnerabilities in each product can be estimated. But the confidence interval can 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.

In general, the repeated measurement method is more easily applied to the uncertainty associated with vulnerability than with other of the risk components. Fortunately, it is not necessary that the same method be used for all components: the two methods can be used together as part of the same risk assessment. Specifically, the best case-worst case method could be used to derive the uncertainty associated with threat, and the repeated measurement method could be used to derive the uncertainty associated with vulnerability. Both methods yield a lower bound value and an upper bound value for

their respective components. The two lower bound figures and the two upper bound figures can be multiplied to produce lower bound and upper bound values for likelihood—the difference between them representing the uncertainty associated with likelihood.

3.5 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 his or her risk estimates, albeit perhaps unconsciously. This level of confidence, expressed as the confidence coefficient, $1 - \varepsilon$, represents the probability that the "real" value of R actually falls within the confidence interval defined by 2δ, the zone of uncertainty. One person, trying to be absolutely certain that R falls within this zone (equivalent to a value for $1 - \varepsilon$ 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 only a very small amount of uncertainty and a smaller confidence interval, effectively applying a much smaller value for $1 - \varepsilon$.

This means that uncertainty (or confidence interval, 2δ) and the confidence coefficient, $1 - \varepsilon$, 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. Dissatisfied with this uncertainty, management revises its direction and tells them to produce an answer with 100% confidence. In this case, the assessors find that they are absolutely sure that the expected loss will fall somewhere between \$0 and \$100M—the latter figure being the point at which the company's 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, $1 - \varepsilon$, 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 [6]. 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 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.

3.6 Consequence uncertainty

When attempting to generate estimates of risk uncertainty, it is important not to neglect the contribution from consequence. 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 upon 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 its 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 [7]. Understanding the enterprise's survival limits is also important, since, from the perspective of the enterprise itself, the total demise 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.

One of the general approaches that we outlined for evaluating uncertainty—namely, the best case-worst case approach—can be used to judge 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. The Role of Assurance in Risk Management

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.

4.1 The risk plane

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 (see Figure 3). The figure also shows how the consequence uncertainty interacts with the likelihood uncertainty to define an area of risk uncertainty.[4] By

[4] A slightly different method of displaying virtually the same data appears in John Carroll's book on computer security. His book

plotting consequence and likelihood as ranges, and thus risk as an area, a much better intuitive understanding of this uncertainty is conveyed.

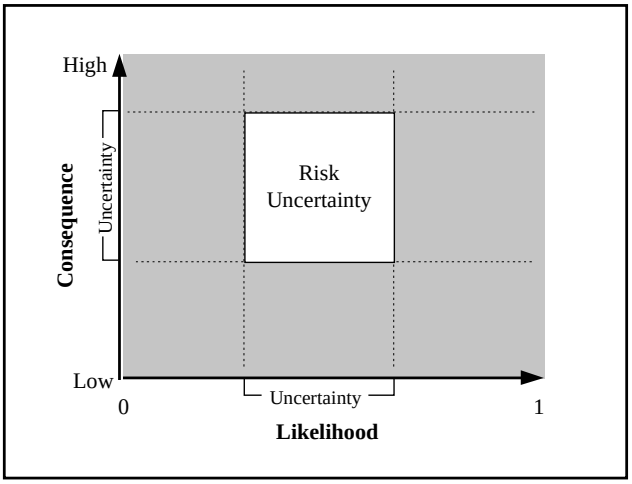

Figure 3: Combining Consequence and Likelihood Uncertainty

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 risk uncertainty is by no means the only factor that must be considered. Cost, schedule, complexity, and even practicality are other important considerations; but the lower the risk uncertainty, the easier it is to balance these other factors.

Figure 4 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 a constant risk curve through the risk plane. Events below the curve are then ignored while events above the curve 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). Starting at the top of the prioritized list, 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 paper, however, events are not simply plotted as points on the risk plane, but as areas of uncertainty. In cases where the risk threshold curve passes through the uncertainty rectangle, as in Figure 4, there is no way to be certain whether the actual risk falls above or below the threshold. In such cases, it is often unclear whether it would be more useful to collapse the rectangle by adding security mechanisms or by obtaining more assurance evidence.

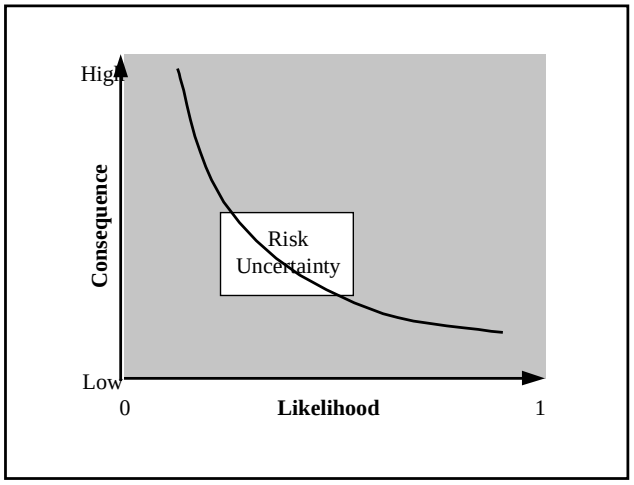

Figure 4: A Simplistic Approach to Risk Decisions

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. This action will often reduce the risk to an acceptable level.

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.

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,

includes a figure containing what he calls a “Plane of Uncertainty,” in which severity uncertainty is plotted against frequency uncertainty. See [8], p. 482.

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 uncertainty is so great that the company decides that the risk of viruses is still unacceptable, they are forced to take some action.

In these cases, 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 more detailed information about the particular virus checker in question. Either method should result in a reduced confidence interval, 2δ, and greater assurance. Adding a security mechanism can 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 risk is reasonably well known, i.e., 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 has been well established.

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

There are other situations in which the uncertainty associated with a risk estimate is very large. Here, the appropriate action may be different since the assessors do not have good estimates of 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 accurate estimates of threat, vulnerability and consequence can be based. Sometimes it merely requires assembling the evidence that is already available. This can be quite inexpensive and yet provide a great deal of information to the assessor. In other cases, it is necessary to seek additional information or evidence. Adding evidence serves 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 greater than originally thought. Although perhaps discouraging, this is useful information that helps the decision-maker make an informed choice.

Having done all this, the decision-maker probably still does not know how much assurance she may have needed, but she is now in a position to decide whether or not the amount of assurance she has is sufficient, and this we claim is quite good enough.

5.3 Case studies

To better illustrate the dilemma often faced by the decision-maker, Figure 5 plots six extreme, stylized event classes on a risk plane. Each event class is represented as a rectangle whose vertical and horizontal dimensions represent, respectively, the 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.

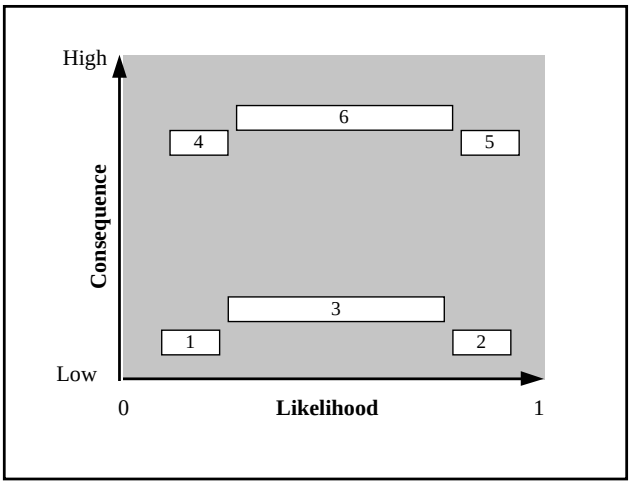

Figure 5: Example Events Plotted on a Risk Plane

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 uninterrupted power supply (UPS), it might be worth doing.

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 additional 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 are 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.

Case	Likeli-hood	Conse-quence	Assur-ance	Example	Action
1	Low	Low	High	Passing comet	Ignore
2	High	Low	High	Power interruption	Fix if cheap
3	Unk	Low	Low	Warranty-protected HW failure	Get more informa-tion or fix
4	Low	High	High	Meteor	Fix if cheap
5	High	High	High	Unknown Virus	Fix if at all possible
6	Unk	High	Low	Malicious Insider	Get more informa-tion or fix

Table 1: Summary of Different Risk Cases

6. Summary

In this paper, we have presented a case for a somewhat modified view of assurance—one that is orthogonal to the concept of security and risk, and closer to the intuitive notion implying the absence of uncertainty. This, we have argued, has the advantage of better accommodating "negative evidence" and of providing a means of expressing the situation in which low security is accompanied by great certainty. We have also presented a practical method of generating some rough measure of assurance given some means of determining risk. Finally, we have shown how a richer understanding of assurance and its basis can be used in the every-day risk management decisions that decision-makers are called upon to make.

Acknowledgments

No paper is the sole product of its authors: this one is the result of many discussions with many people over several years, whose contribution we gratefully acknowledge. We wish particularly to thank the reviewers of our original version who offered cogent comments and several helpful suggestions, most of which we attempted to apply. We feel that the paper has benefited greatly as a result. Finally, we wish to acknowledge and thank the National Security Agency without whose funding support this paper would never have been written.

References

[1] Jeffrey R. Williams and George F. Jelen, *A Framework for Reasoning about Assurance*, Document Number ATR 97043, Arca Systems, Inc., 23 April 1998.

[2] National Institute of Standards and Technology, Invitational Workshop on Information Technology (IT) Assurance and Trustworthiness, March, 1995.

[3] David R. Wichers, Joel E. Sachs, and Douglas J. Landoll, "What Color Is Your Assurance?"

[4] Idaho State University, Glossary of INFOSEC and INFOSEC Related Terms, Vol. II, Version 6, dated August 1996.

[5] Harald Cramér, *The Elements of Probability Theory*, (New York: John Wiley & Sons, 1955).

[6] Jeffrey R. Williams and Marvin Schaefer, "Pretty Good Assurance," *Proceedings of the New Security Paradigms Workshop*, (IEEE Computer Society Press, 1995).

[7] 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.

[8] John M. Carroll, Computer Security, 3rd ed. (Boston: Butterworth-Heinemann, 1995).