

REDUCING SUBJECTIVITY IN QUALITATIVE RISK ASSESSMENTS

Robert Smock

June 2002

SANS Security Essentials

GSEC Practical Requirements (v.1.3) (December 2001)

ABSTRACT

One of the acknowledged drawbacks inherent in all types of risk assessments is the difficulty in adequately comparing risks for the purpose of risk-mitigation decision-making. This drawback is exacerbated in qualitative analysis because of the subjectivity of the assessment process and the lack-in-rigor of the descriptive units of measure. Quantitative analysis, typically using dollars as a unit of measure, provides some structure for such a comparison, but can be complex and difficult to do. This study describes a simple approach in establishing a framework for providing this comparison capability while at the same time, reducing the subjectivity of the assessment process. While the focus of this study is on qualitative risk assessment, it is equally applicable to other methodologies.

The approach described below involves the establishment of a risk scoring methodology for those risks identified during the course of a qualitative risk assessment. While simple, it has the effect of removing a certain amount of the subjectivity inherent in such analyses by using operational impacts as a criteria for assigning value to risk. These values can then be used to appropriately prioritize the risks for mitigation or acceptance. In practice, the risk scoring approach tends to add credibility to the assessment process and typically improves communication between the analyst and the business decision maker regarding the relative nature of the identified risks. The bottom line benefit of this process is an improved security posture through improved awareness of risks, typically leading to reduced risk through increased mitigation.

BACKGROUND

The most fundamental aspect of protecting information and information technology (IT) resources is the need to assess the risk to those resources in the environment that they are used or will be operating. It is widely recognized that risk cannot be effectively or efficiently reduced until what kinds of and how much risk is known. A large proportion of those whose job it is to safeguard computing resources will – or should - spend a significant amount of their time and effort assessing risk.

The goal of assessing risk is to identify those issues that may negatively impact the resources operating in a particular environment (OMB). However, the objective for any organization that undertakes a risk assessment is to ultimately act on the findings of the risk assessment, i.e. to mitigate the risks identified through the assessment process. Ideally, an organization strives to implement risk mitigations that result in the greatest reduction in risk for the least amount of cost. The key here is the measurement of the risk,

i.e. how much risk exists. Without an accurate measurement of risk, actions taken to reduce risk will be, at best, only a guess. With only a small portion of IT budgets today being allocated to security, such a “hit and miss” approach to reducing risk and increasing security is unacceptable.

Today, the security industry and practitioners generally agree that there are at least three primary types of assessment used for identifying and measuring risk: qualitative, quantitative, and knowledge or rule based (Ruthberg, Tipton, Bonyun). *Please note that for the purposes of this study, we will not make the fine distinction between assessment and analysis, and will use these terms interchangeably (C&A).* Each type of assessment has its own benefits and drawbacks, and each allows the security practitioner to gather, analyze, and report risk in varying degrees of depth & thoroughness. When applied to information and information technology resources, each is applicable and acceptable as conditions – and customers – dictate.

RISK ASSESSMENT FUNDAMENTALS

All risk assessment methodologies require the same basic fundamental information to begin: a definition of the scope of the risk assessment, i.e. what resources will be assessed, the assignment of a value or importance to the resources included in the scope of the assessment, the identification of all possible threats to those resources, and the identification of all existing vulnerabilities in those resources. Once this information is gathered, threats are then matched (paired) with those vulnerabilities that would allow the threat to impact the resource. Given this information, risk can now be “measured”.

Based on the above information, the assessment of the amount of risk to the resources identified in the scope of the assessment is then simply a matter of plugging “values” into the risk equation (Tipton, Krause, Ozier), which can be generalized as

$$\text{Risk} = \text{impact} \times \text{likelihood} \times (\text{threat} \times \text{vulnerability})$$

where *impact* is the amount of damage that would be suffered by the identified resources based on the importance of the resource, *likelihood* is the probability (i.e. chance) that this damage will occur as a result of the listed threat-vulnerability pairing, *threat* is the event which would cause the damage, and *vulnerability* is the weakness in the resource which would allow the threat to cause the damage.

It is interesting to note here one of the basic truths of risk assessment that is commonly overlooked or forgotten when using risk measurement as a means for decision making, even by those within the security industry. The generalized equation above is purposely expressed as an exercise in multiplication to demonstrate the implicit fact that if any of the variables on the right side of the equation are or can be made to equal zero, then there is no resultant risk, i.e. the risk to the resource being assessed is also equal to zero. Therefore the goal in mitigating risk is to take actions which will drive one or more of the four variables to zero, or more practically, to as close to zero as possible (University of Houston).

METHODOLOGY COMPARISONS

In order to understand the evolution of the risk comparison framework that will be described, we will briefly review the primary risk assessments methodologies, paying particular attention to their respective differences and similarities, strengths and weaknesses.

Quantitative and qualitative risk assessments are basically conducted using similar techniques. Their differences lie in the robustness and preciseness of the values used in the risk equation described previously, and in the resulting measurement of risk. After identifying all possible threats to a given set of resources, and all possible existing weaknesses (vulnerabilities) in those resources, the threats and vulnerabilities are “paired”, combining those threats that are capable of causing damage by exploiting the given weakness. After all possible threat-vulnerability pairs are identified, and the resultant risk is measured, plans then are developed to mitigate the risk present with each specific threat-vulnerability pairing. It is when the resultant risk is measured that the differences between the two methodologies become apparent.

Knowledge-based risk assessments are generally considered a third type of risk assessment methodology, although that is not entirely correct. In practice, either quantitative or qualitative metrics, which will be described in some detail shortly, may be used in conjunction with knowledge-based analyses. The difference is in how the threat-vulnerability pairings are derived. In a knowledge-based scheme, an assumption is made about the nature and use of IT resources, that there are a generally accepted “standard” set of threats and vulnerabilities inherent in the use of IT resources regardless of location or function (Childs).

Given this assumption, a set of “rules” or guidelines – requirements – can be developed which, if implemented, will act to mitigate risk by eliminating or reducing the various assumed threats and vulnerabilities. Varying levels of security (e.g. strong, moderate, or weak) can be established by varying the robustness by which the “requirements” are implemented. Likewise, risk can be identified by measuring the robustness of the implementation of the individual requirements.

METRICS

Quantitative analysis of risk typically assigns specific dollar values (e.g. the purchase or replacement cost) as a way to measure the relative importance of the resources being assessed, i.e. the more important a resource, the higher the dollar value. Such analyses also typically measure risk as a dollar amount that can be expected to be lost over a given period of time. Quantitative analysis uses finite probabilities for the potential of occurrence (likelihood) of a specific threat exploiting a specific vulnerability. This likelihood is usually calculated from actual historical data, typically the number of occurrences of the same or related events over a specific period of time.

In contrast, qualitative analysis of risk assigns descriptive values (e.g. high, moderate, low) as a way to indicate the relative importance of the resources being assessed. Likewise, similar descriptive values (e.g. high, moderate, low) are used to indicate the relative danger (i.e. risk) to the resources being assessed, and the potential for the occurrence (i.e. likelihood) of a specific threat exploiting a specific weakness.

As stated previously, either type of metric is applicable to the knowledge-based methodology, although more often than not, qualitative metrics are used because they are more easily applied to the individual rules or requirements.

BENEFITS AND CONSTRAINTS

In practice, because quantitative analysis is based on mathematics and statistics supported by objective metrics, such analyses are considered to be more rigorous. Metric values are typically expressed in dollars, which the business oriented decision makers are more likely to grasp and understand with regards to the concept of risk and risk mitigation being a function of cost or dollars lost. Using a standard metric (e.g. dollars) for similar risk assessments across multiple platforms, facilities, or business units provides a basis for comparison of risk and prioritization of risk mitigation activities.

The down side to quantitative analysis is that it many times is difficult, complex, and time-consuming to conduct and interpret. Standard metrics are not available for every identifiable risk scenario, i.e., objective probability-of-occurrence metrics do not exist for every conceivable threat-vulnerability pairing, making the assessment more difficult. Such analyses can become complex because what standard metrics that do exist are not always applicable in all cases. For example, how does one assign a dollar value to a human life, an irreplaceable art object, or human thoughts and ideas.

Finally, such analyses can be time consuming because of the effort involved in assigning metric/dollar values to every identified resource, to cost every conceivable impact, and to calculate the resultant risk for every conceivable threat-vulnerability pairing. Additionally, such metric/dollar-based analyses then requires the costing of multiple risk mitigation options to form the basis of the risk management decision-making process and allowing appropriate risk mitigation priorities to be set.

In contrast, the benefits of qualitative analysis are that it is quick and simple to conduct, and is intuitive to communicate and understand. Neither extensive research or complex calculations are necessary to apply the typical, simple, 3-tier descriptive indicator (i.e. high, moderate, low) to the importance of assets, to the chance of occurrence, or to the impact of an occurrence of a threat-vulnerability exploitation; subjective, experience-based “estimates” as to the relative nature of the metric will suffice. Results are easily communicated to decision-makers because the comprehension of the relative difference between descriptive metrics such as “high” and “low” are almost universally understood in every scenario.

The most obvious drawback to the qualitative assessment methodology is the rigor and definition of the metrics, i.e. how high is a “high risk”, how much difference is there between a “high” risk and a “low” risk, and what is the difference between one “high risk” and another.

The benefits of a knowledge or rule-based risk assessment methodology are reduced effort on the part of the analyst, a reduced learning curve for new analysts, and a faster turn-around time for completion of analyses. Because the threat & vulnerability aspects of the analysis are implicit, the effort behind identifying and quantifying these variables is dramatically reduced. Once the rule-set is established for an existing set of resources, new analysts require less training before they become effective at applying the rules in a risk assessment. Because of the reduced effort in identifying the risk variables and because the rule set is standardized, risk assessments can be completed in much less time.

The rule-based methodology is not without its drawbacks however. Because threats and vulnerabilities are implied, certain specific threats and vulnerabilities unique to an environment or a set of resources may be overlooked. Because existing, known threats and vulnerabilities are addressed generically through a standard rule-set, the robustness of the mitigation of the risk presented by the specific threat or vulnerability may not be commensurate with the amount of risk that can be tolerated in a specific environment or with a specific set of resources.

RISK PRIORITIZATION

From a practical perspective in today’s business environment – and even more so in a government environment – a single management chain, a single decision-maker has the fiduciary and operational responsibility for multiple sets of resources, with each resource set typically having individual and independent risk assessments. The issue is how does this single decision-maker, assuming the availability of limited resources and funding, prioritize the risks across the multiple sets of resources, ensuring that the limited resources are used to address the most serious risks.

If the risk assessments were conducted using the quantitative, dollar-based methodology described previously, one consideration typically used is the amount of loss expected for each risk; those with more potential for loss receive the highest priorities for mitigation. Of course, the relative importance of each of the sets of assets could be another consideration. However, if the risk assessment were conducted using qualitative methodologies, the issue becomes less clear, e.g. how does one compare one “high” risk with another “high” risk, or how does one know one “high” importance asset versus another.

One method that can be used to overcome the previously described constraints in comparing the results of separate, qualitative risk assessments across different sets of resources is to develop a scoring methodology. By using a simple, stand-alone, numeric scoring algorithm based on generalized operational impacts rather than on specific risk measurement units (e.g. dollars), we can provide a simple, relatively objective basis for

comparing and prioritizing risks for mitigation or acceptance. Increased objectivity is achieved because of the replacement of the subjective, descriptive metrics typically used in qualitative risk assessments with actual, operations-based criteria that support the assignment of values used in the calculation of risk.

RISK SCORING

The method presented here is one possible approach, but the scoring system used and the associated criteria are arbitrary and open to definition as necessary as conditions and customers dictate. The two factors that make any defined scoring system relevant are its applicability to a particular environment and the fact that it is used consistently across all sets of resources, for all risk assessments.

For the purposes of this study, we will use a qualitative, knowledge-based risk assessment methodology as our point of reference. The risk scoring process will be applied against the “rules” established by the knowledge-based methodology.

As with all risk assessment methods, the starting point for the scoring process is the typical gathering of the usual information required to define a risk event as previously described, primarily the threats and vulnerabilities, or in the case of knowledge-based assessments, the “rules” are developed.

It may be useful at this point to recall the generic equation for risk as previously described, namely

$$\text{Risk} = \text{impact} \times \text{likelihood} \times (\text{threat} \times \text{vulnerability})$$

where “threat x vulnerability” is also equal to “rule” for knowledge-based assessments.

Again, the above equation is merely a representation of the calculation of risk. From a practical standpoint, once all of the rules/requirements (threat-vulnerability pairs) are established, the equation (Childs) reduces to

$$\text{Risk} = \text{impact} \times \text{likelihood}$$

since we assume the existence of a threat for every rule, and every rule essentially equates to a potential vulnerability. Therefore the risk of exploitation of each rule, i.e. the chance of damage being caused because of a particular rule being broken or a specific rule not being implemented, is simply the product of the estimated amount of damage (impact) that can be caused and the estimated likelihood (probability) of the exploitation taking place.

For the purposes of this study, we will use a simple scoring scale of one (1) to five (5), with one being the least and five being the most. Using the simplified risk equation established above we will use a combination of professional judgment and pre-defined criteria to establish values for impact and likelihood. The calculation will result in a risk

score of one (1) to (25), where low scores represent low risk and high scores represent high risk.

The actual establishment of what constitutes a “low” value for risk and what constitutes a “high” value is again, highly subjective and arbitrary, and should be tailored as the needs of the environment and customer dictates. It bears repeating that the two factors that make any defined scoring system relevant are it’s applicability to a particular environment and the fact that it is used consistently across all sets of resources, for all risk assessments.

It also bears pointing out that the subjectivity and arbitrariness of the low/high definitions is tempered by the experience and training of the risk or security analyst, and the definitions should only be established after the environment – the business needs, the mission/objective, asset importance, potential threats, etc. – is analyzed and a thorough understanding is established.

Done correctly, the overwhelming intuitiveness of the scoring method becomes obvious, as even the most casual observer – or business manager – can understand the simple correlation between high scores and high risk. The correlation between risk level and score is illustrated in the following table:

		Likelihood				
		1	2	3	4	5
Impact	5	5	10	15	20	25
	4	4	8	12	16	20
	3	3	6	9	12	15
	2	2	4	6	8	10
	1	1	2	3	4	5

Table 1 – RISK LEVEL

The risk score for a given vulnerability/rule is determined by multiplying the impact and likelihood values. The table illustrates the corresponding level of risk by the color associated with the risk score; red – high risk, yellow – moderate risk, green – low risk. The goal of risk management is to make changes (eliminate vulnerabilities, implement controls, etc.) to reduce the risk score to the low-risk range. The coverage of the risk level rating colors on the table (e.g. more green scores, less yellow scores) can be adjusted according to an organizations tolerance to risk.

Again, what constitutes the various levels and associated scores/colors of risk can be altered as necessary to achieve the objective of appropriately measuring the risk in a given environment, as long as the same representations are used for all resources in a given environment, and in all risk assessments conducted on the resources in that given environment (GAO).

SCORING CRITERIA

In an effort to remove the subjectivity inherent in such a qualitative assessment methodology and to add rigor to the risk measurement/scoring process, we must develop criteria to use as a basis for determining the values for impact and likelihood in our risk

calculation. Again, the specific criteria used is only important in as far as it is applicable to the environment in which the resources being assessed operate, and that the criteria is used equally for all resources.

The primary consideration in developing such criteria should be that it accurately reflect the effect that a particular threat-vulnerability exploitation, or for our purposes, a rule violation, would have on the ultimate mission, purpose, and objective of the resource being assessed (GAO).

By using a combination of the established criteria and the information gathered by the analyst regarding the operating environment, and based on the experience of a particular analyst, a subjective determination can be made as to which criteria best describes the effect that a given rule/requirement would have on the ability of the resource being assessed to carry out and achieve its primary function.

The criteria to be used for the purposes of this study are represented in the following tables. Table two (2) represents the criteria for assigning IMPACT values – how much damage can be done. Table three (3) represents the criteria for assigning LIKELIHOOD values – what is the probability that the damage will occur, i.e. what is the chance that the risk will be realized.

Value	IMPACT
5	Any loss of ability to perform the primary function of the resource during critical processing events/time
	Unrecoverable loss of ability to perform the primary function of the resource
4	Reduced ability to perform the primary function of the resource during critical processing events/time
	Long-term recoverable loss of the ability to perform the primary function of the resource
3	Reduced ability to perform the primary function of the resource
	Inability to detect or reconstruct the scope of loss of the primary function of the resource
2	Short-term recoverable loss of the primary function of the resource
	Reduced ability to detect or reconstruct the scope of loss of the primary function of the resource
1	Accepted Industry standard

Table 2 – IMPACT
What is the damage that will be done if the risk is realized, i.e. what damage is caused by a threat exploiting the vulnerability/rule.

The objectivity of the impact criteria can be enhanced by ensuring that each item in the table can be related to at least one of the fundamental objectives of security (GAO), namely, detection – the ability to identify an attempted or successful exploitation,

prevention – the ability to avoid a successful exploitation, and recovery – the ability to restart or continue operations despite a successful exploitation. A comprehensive impact table may reflect various levels of impact for each of the three objectives.

Value	LIKELIHOOD
5	More than 1 occurrence per year
4	No more than 1 occurrence per year
3	More than 1 occurrence in lifetime of the resource
2	No more than 1 occurrence in lifetime of the resource
1	Potential for at least 1 occurrence but controls are in place

Table 3 – LIKELIHOOD
What is the probability that the impact/damage will occur, i.e. what is the chance that the risk will be realized, that a threat will exploit the vulnerability/rule.

At this point, given the information collected, the established criteria, and a measure of professional judgment and experience, the analyst now simply assigns values for the impact and likelihood for each rule in the knowledge base, and then calculates the resultant risk for each rule. The score is then correlated with the assigned level of risk as determined from table 1.

PRIORITY RANKING OF RISKS

Using the described procedure results in a risk score that readily facilitates the prioritization of the risks for mitigation or acceptance.

Risks with a “low” score (1-6) that falls into the green portion of table 1 can typically be considered low risk and accepted without further mitigation, or planned for future, long-term mitigation.

Risks with a “moderate” score (8 – 12) that fall within the yellow portion of table 1 can typically be considered as presenting enough of a risk to be considered for immediate or near-term mitigation.

Risks with a “high” score (13 – 25) that fall within the red portion of table 1 can typically be considered as being a high risk, presenting an immediate threat to continued secure operations, and should be addressed immediately.

Again, the range of the “high”, “moderate”, and “low” scores, i.e. the coverage of the risk level rating colors on the table (e.g. more green scores, less yellow scores), can be adjusted according to an organizations tolerance to risk.

EXAMPLE

At this point, it may be useful to process through an example of the described risk scoring methodology.

For purposes of this example, the basis of our qualitative, knowledge-based risk assessment process consists of the following arbitrary rule set:

1. Must log security relevant events
2. Must review logs regularly
3. Must not use group user Ids which prevent individual accountability
4. Must suspend an account indefinitely after 3 failed logon attempts
5. Must prevent use of trivial passwords
6. Must expire passwords after an appropriate period of time as defined by system category
7. Must encrypt password files, private data, and other sensitive data if concern is confidentiality
8. Must implement password scheme to enforce minimum password length and character content
9. Must implement malicious code checking
10. Must implement virus detection & eradication

Assuming a fictional set of resources, we apply the criteria from tables 1 and 2 against this rule set and calculate the resultant risk.

Rule	Item	Impact	Likelihood	Risk score	Risk Level
1	Log events	3	2	6	Low
2	Review logs	3	3	9	Moderate
3	Accountability	2	2	4	Low
4	Failed logons	1	2	2	Low
5	Trivial passwords	5	5	25	High
6	Password expiration	3	4	12	Moderate
7	Encryption	3	2	6	Low
8	Password content	4	4	16	High
9	Mal-ware checking	3	1	3	Low
10	Anti-virus	4	5	20	High

Table 4 – EXAMPLE

For each “rule” (item, column 2) in the knowledge-based risk assessment, impact and likelihood are determined by the analyst, based on their experience and expertise, from the established criteria for these values in tables 2 and 3. The values are multiplied to determine the risk score (column 5). The risk score is then matched to the appropriate level (column 6) as determined in table 1.

The resulting [risk scores](#) associated with each of our rules can now be used to prioritize our risk mitigation or acceptance activities, allowing us to appropriately prioritize not only between high and low risk, but also between two or more high risks.

SUMMARY & CONCLUSION

The presented study demonstrates that a certain amount of subjectivity can be removed from qualitative risk assessment methodologies to produce results that can be used to appropriately measure and prioritize risks and associated risk mitigations.

All risk assessment methodologies start with the same basic information which must be collected prior to conducting the analysis, namely, what resources will be assessed, the assignment of a value or importance to the resources, the identification of all possible threats to those resources, and the identification of all existing vulnerabilities in those resources.

Depending on the assessment methodology to be used, threats and vulnerabilities are paired or rules/requirements are established, and risk is measured – qualitatively or quantitatively – for each pair/rule using the generalized risk equation where risk is equal to the product of the impact of the event and the likelihood of the occurrence of the event.

For qualitative risk assessments, objective criteria, based on the effect to operations of the resources being assessed, can be established for “impact” and “likelihood”. Through the professional expertise of the analyst, the criteria can be mapped to a risk score that would accurately reflect the relative risk to resources. The risk scores then provide a relatively objective basis for prioritizing individual risks for mitigation consideration.

By adding a simple, stand-alone, numeric scoring algorithm based on true impact to operations, we can provide a simple, relatively objective basis for comparing and prioritizing risks for mitigation or acceptance. The operation is simple because it uses mathematics operations that are no more complex other than simple multiplication. Increased objectivity is achieved because of the operations-based criteria supporting the assignment of values used in the calculations. The scoring scheme is applicable to either threat-vulnerability pairs or individual “rules”, making it applicable to all types of risk assessment.

The benefits of this scoring methodology - reduced costs and improved risk management efficiency and productivity - typically overshadow the drawbacks of the methodology for most IT environments. This methodology can be tailored to any organizations needs by applying either qualitative or quantitative metrics to the results. In today’s competitive, cost-conscious business environment, the benefits of reduced costs and improved productivity make this methodology an attractive option for many organizations in a wide variety of risk environments.

REFERENCES

1. Childs, David, “Information Technology Security System Engineering Methodology”, SANS GSEC Practical Assignment v.1.3, April 2002.
http://www.giac.org/practical/David_Childs_GSEC.doc

2. Ruthberg, Zella, and Tipton, Harold, "Information Security Management Handbook", 1st Edition, Auerbach, 1993, 115 – 137, Bonyun, David, *Techniques of Risk Analysis*.
3. Hutt, Authur, Bosworth, Seymour, and Hoyt, Douglas, "Computer Security Handbook", 3rd Edition, Wiley, 1995, 3.1 – 3.20, Carroll, John, *Information Security Risk Management*.
4. Tipton, Harold and Krause, Micki, "Information Security Management Handbook", 4th Edition, Auerbach, 2000, 247 – 287, Ozier, Will, *Risk Analysis and Assessment*.
5. University of Houston-Clear Lake, "Introduction to Information Security Management", course notes, Fall 2001.
6. C&A Security Risk Analysis Group, "Introduction to Security Risk Analysis & Security Risk Assessment", 2001.
<http://www.security-risk-analysis.com/>
7. United States Government Accounting Office (GAO), "Information Security Risk Assessment, GAO Practices of Leading Organizations", special publication GAO/AIMD-00-33, 2000.
<http://www.gao.gov/special.pubs/ai00033.pdf>
8. Brooke, Paul, "Risk Assessment Strategies", Network Computing Magazine, October 2000.
<http://www.networkcomputing.com/1121/1121f3.html>
9. Federal Deposit Insurance Corporation (FDIC), "Risk Assessment Tools and Practices for Information System Security", FDIC Financial Institution Letters, 1999.
<http://www.fdic.gov/news/news/financial/1999/FIL9968a.HTML>