



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Risk Assessment
A Practical Approach
by
Dan Hlavac

GIAC Security Essential
Practical Assignment
Version 1.4b
December 5, 2003

© SANS Institute 2004, Author retains full rights.

Introduction	3
Risk Assessment vs. Risk Management.....	3
<i>What is Risk</i>	3
Risk Assessment Methods.....	4
<i>Quantitative vs. Qualitative</i>	5
<i>Affects of Mitigating Controls</i>	7
<i>Data Classification</i>	7
Practical Risk Assessment.....	8
<i>Eleven items that need to be addressed</i>	9
<i>Enterprise Model</i>	10
<i>Education Process</i>	10
<i>Acceptable quantifiable levels for business impact and likelihood</i>	11
<i>Baselines for key areas</i>	12
<i>Predetermined risk tolerance levels</i>	12
<i>Risk acceptance procedures</i>	13
<i>Central repository for mitigating controls</i>	13
<i>Effectiveness of mitigating controls</i>	13
<i>Cost, ease of implementation, and impact of controls</i>	14
<i>Solution patterns</i>	14
<i>Risk Monitoring</i>	14
Conclusion.....	15
References	16

© SANS Institute 2004, Author retains full rights.

Introduction

“Risk assessment” and “risk management” have become some of the latest and greatest buzz-words in our industry. These concepts are not really new, they’ve been around for decades, but they have been getting a lot more attention lately. Most of this attention comes from the business areas that are writing the checks for the security solutions used by an organization. Assessing and managing information risk are the means by which businesses hope to show a return on their investment and/or a logical rationale for some of the decisions made regarding mitigating controls.

With all of the attention on information “risk assessment” and “risk management”, it is important to separate the hype from reality. The pages that follow will focus on risk assessment and the eleven primary items that need to be addressed when creating and executing risk assessment. These eleven items provide a clear understanding of how risk assessment fits into an organization’s operational strategy and provides a better understanding of what risk assessment should be used for. Before we look at these items, we should take a cursory overview of basic risk concepts.

Risk Assessment vs. Risk Management

There are many definitions of risk assessment, but the common theme is the “analysis of risk”. To produce this analysis, the standard formula used throughout the industry is identifying assets and comparing them to the threats and vulnerabilities relative their importance. Most of the tools for analyzing and addressing risk use a derivative of this formula:

$$\text{Risk} = \text{Threat (or threat level)} \times \text{Vulnerabilities} \times \text{Asset (or Impact)}$$

Most realize that this is NOT an actual formula that should be used for quantitative analysis. Rather, it is a guideline that illustrates the components for risk assessment.

What is Risk

To help reduce ambiguity regarding the definitions of this discipline, several key items should be addressed. First, “risk” as defined by Dictionary.com is *“the possibility of suffering harm or loss; danger”*. When discussing risk for the information security practitioner or professional, Symantec indicates risk is *“A threat that exploits a vulnerability that may cause harm to one or more assets”*. (Symantec, 2003).

A “threat” is an entity capable of causing harm or loss. According to Symantec, a “threat” is “a circumstance, event, or person with the potential to cause harm to a

system in the form of destruction, disclosure, data modification, and/or Denial of Service (DoS).” The “threat level” is the likelihood of a threat attempting to manifest itself in a given situation. Finally, a “vulnerability” is the means by which a threat has the potential to manifest itself.

Risk assessment is the analysis by which the level of risk is derived. Risk management on the other hand is the method by which risk is reviewed, mitigated, accepted, or transferred. There is a very strong distinction between the two. Assessment is the identification of risk while management is the method for doing something with that knowledge. For example, there was one fast food company that may or may not have identified the risk of having their coffee extra hot. To my understanding, it was done because they wanted to ensure that the coffee was still hot when the customer arrived at work. I will assume that the company thought they understood the risks at the time, but did not count on such a substantial law-suit after it spilled on someone. To better manage that risk, the company now serves the coffee at a much lower temperature and added additional warning labels. In Information Technology, the risk of not encrypting customer data may be very large for a variety of reasons. However, management of that risk may come in the form of intrusion detection or advanced auditing and alerting.

Risk Assessment Methods

Methods for assessing risk come in many shapes and sizes. Generally, they all have the same approach – figure out the impact of an event and then determine the likelihood of that event. For example, the International Organization for Standardization (ISO 17799) considers risk assessment as: “...systematic consideration of :

- a) the business harm likely to result from a security failure, taking into account the potential consequences of loss of confidentiality, integrity or availability of the information and other assets
- b) the realistic likelihood of such a failure occurring in the light of prevailing threats and vulnerabilities, and the controls currently implemented (ISP/IEC 17799:2000(E). pp. ix).”

Some assessment methods include other factors, such as assurance level – how comfortable are you with the analysis you have just provided, etc. For the purposes of this paper, we will continue to follow the basic principles.

A risk assessment should produce a value for the data regarding acceptable losses of confidentiality, integrity and availability. Determining each of these items is tricky at best. If an organization only has two types of data: sensitive and non-sensitive, then risk assessment may be easier. In fact, industry best practices would be easier to implement and there would be fewer exceptions. However, most organizations implement new applications or processes that

require an assessment falling somewhere in the middle. This is the point at which an organization wants to ensure the controls are adequate but do not cost too much. For this to occur, an organization should probably have a list of available solutions and understand which combinations of solutions are better than others. Again, the best solution will depend on the value of the data and how much the company is willing to spend to protect it vs. the money they could lose if it wasn't protected. One of the most difficult questions an organization must answer is how much money they will lose. Monetary losses can be calculated for some items, such as production losses, but trying to calculate losses from bad press is much harder.

Another problem is there are very few methods that measure the effectiveness of controls. Which combination of controls is better? Information security is not like an automobile which can be rated for safety based on whether it has a one stage or two stage airbag, side impact beams, side impact airbags, disc brakes, etc. It can better be compared to physical security which uses locks, CCTV, motion detectors, etc. Clearly more motion detectors and cameras are better than fewer. But how many are enough? As many as required to watch every part of the premises you want to watch? Is an organization simply recording information or is there an effort to actively monitor the data with dedicated personnel? How much is at risk if you do not actively monitor something? Where is the trade off? Some industries have security patterns established aligned with industry standards. For example, the casino industry has realized they have a very large threat level. As an industry, they choose to monitor every table and nearly every participant for fraud. But again, even in that industry, there is a struggle to keep up as threats change.

Quantitative vs. Qualitative

Most assessment methods can be divided into two categories - quantitative and qualitative. The former has an emphasis on numerical calculations, while the latter relies more on "best guess" ratings. To be fair, most tools have combinations of both, but usually place an emphasis on a certain approach over the other.

For example, the National Institute of Standards and Technology (NIST) calls for a program that includes a need for metrics to be quantifiable, readily available, using a repeatable process, and useful for tracking performance and resources. The metrics can be of three types: implementation of security policy, effectiveness of security services, and impact of security events to the business. They have a very detailed description of the organizational structure needed, the roles and responsibilities of each person or group, and the maturity of metrics process within an organization. However, their approach only leads the organization to defining and tracking goals. It does not give any insight regarding how those goals can be measured, only that they should be measured.

The problem with the NIST example, and with many others, has not gone unnoticed. Donn Parker once indicated that trying to assign numbers to something as complex and unpredictable as the criminal mind is irrational (Briney). However, this has not stopped the industry from trying to assess the seemingly un-assessable. For example, CCTA Risk Analysis and Management Method (CRAMM) uses a scale from 1 to 7 to compare assets to threats and vulnerabilities. (CRAMM website, 2003). This assumes that a set of professionals can determine the likelihood of an event even when taking into consideration the mindset of irrational people.

There are many debates over the validity of any quantitative approach. It seems reasonable to assume that you cannot quantify those things for which you do not have a measurement. I have not found any statistics regarding the probability of an event happening when taking into consideration business context, data criticality, and the various puzzle pieces of an information security solution. Statistics for this type of information is very difficult to find. In addition, organizations should not rely solely on the data from its own organization. This is similar to only basing insurance on what has happened to you, and only you, in the past. Rather, insurance is based on empirical evidence from millions of others. Some events can be predicted, such as the likelihood of getting the latest big virus if your systems are not patched. But trying to predict the behavior of irrational people, as Donn Parker would suggest, is next to impossible (Briney, 1999).

Like CRAMM, the Information Security Forum (ISF) has a similar approach in their tools. SPRINT, for example, is a tool that uses a more balanced approach between quantitative and qualitative. It uses a scale from 1 to 5 to indicate the impact to the business if there is a loss of confidentiality, integrity, or availability. It then uses a scale from 1 to 5 to indicate the likelihood of a loss of confidentiality, integrity, or availability. The scales are used as guides when answering a series of questions. After the answers are calculated, the analyst should be able to see the items that have the highest impact to the organization and those that have the highest likelihood of an event occurring. If the item being addressed has high marks in both categories, then that item should have stringent controls applied to help mitigate those risks. (Information Security Forum website, 2003).

However, in my research and use of risk assessment tools, I cannot find any methods or tools that help address the following issues:

- quantifiably address the impact to an organization,
- determine the actual likelihood of an event occurring, and
- calculate the reduction of risk due to mitigating controls

Until we begin to provide raw data for these events, most of these models do not provide much value.

In addition, there should be some analysis for items such as implementation and support costs, ease of implementation and support, operating impacts to the organization, etc. Before we can address these issues, the bulleted items above should be attainable.

Affects of Mitigating Controls

Assuming an organization could quantify the cost related to the loss of confidentiality, integrity, and availability AND can determine what loss is acceptable, but how can it determine the likelihood that an event will happen and the cost associated with the event? The likelihood depends on two variables. First, what is the value of the data to someone who should not have it (hackers, disgruntled employees, corporate moles). And second, what is the possibility someone could successfully breach the security controls that are in place and/or are planned to be implemented. If likelihood can be determined, there should be a method that clearly demonstrates how much any set of mitigating controls reduces that likelihood. While this may seem to be a very large task to undertake, it can begin with small assumptions and grow more exact over time. However, it should still be part of a risk assessment approach.

Data Classification

Before we go any further, there needs to be a quick statement about data classification. I have witnessed firsthand the ugly effects of disparate approaches to mitigating risk when data classification and risk assessment are not married. Most people might assume that one is a natural extension of the other. This is correct if the two are integrated into the same operational system. However, in some organizations, data classification grew up in the legal and record retention neighborhoods. For example, data might have been classified as proprietary or sensitive. Does this mean that one set of data should be more protected than the other set? Not likely. By performing a risk assessment on these data types, an organization should have a better understanding of the impact if there was a loss of confidentiality, integrity, or availability. The assessment will indicate the criticality of the data by means of business impact. What does it mean to be “sensitive” anyway? Who made that determination? An organization may have performed an assessment to make those distinctions in the beginning. However, if that is not the case, then after a risk assessment is performed, an organization might see that some proprietary information is just as important as sensitive information, and vice versa.

Some would argue that the risk assessment process provides its own data classification schema. Some would argue that the legal, operational, and security classifications should be separate and used for separate reasons. I am advocating that an organization should be aware of the possible disparities,

aware of the possible handling requirements based on the different categorizations, and have a plan or strategy for ensure that all of the classification types at least co-exist comfortably.

Practical Risk Assessment

The primary purpose of any set of mitigating controls is to reduce the likelihood of an event occurring within acceptable business limits. When information is evaluated, the risk rating can be plotted on a simple chart. The objective of mitigating controls is to reduce the likelihood of an event from occurring (see Figure 1)




Figure 1

However, ongoing support for a risk assessment process needs to account for many other variables, including, but not limited to, the changes to the business impact and the associated costs of the controls (see Figure 2).




Figure 2

In order to have an effective risk assessment process, an organization must have a practical risk management strategy that accounts for all of these changes plus additional items. It is very futile to create a process that is not encapsulated in an enterprise methodology for addressing, assessing, and accepting the risks that are found or that might be incurred. This paper will not address risk management; however it will address the items necessary for implementing a risk assessment process.

Listed below are the top eleven items needed for this process to be successful. I have not found any tools or methodologies that include all eleven of these items. An organization does not need to have complete answers or strategies for each of the items, but without addressing them, a risk assessment process is “at risk” for failure.

Eleven items that need to be addressed

- 1) An enterprise risk model
- 2) A clear and decisive education process for the business and technical areas
- 3) Acceptable quantifiable levels for business impact and likelihood
- 4) Baselines for key areas
- 5) Predetermined risk tolerance levels
- 6) Risk acceptance procedures
- 7) Central repository for mitigating controls
- 8) Effectiveness of mitigating controls
- 9) Cost, ease of implementation, and impact of controls
- 10) Solution patterns
- 11) Risk monitoring

Enterprise Model

The model should incorporate the needs of the assessment process, the resources, tools, and organizational impact. It should also align with other business type risk assessment processes already established. In addition, it should incorporate any other data classification schemes.

Most importantly, it should indicate the objectives of the model as a tool for the business areas. Businesses exist because they take risks. The model should establish guidelines for providing appropriate controls to mitigate those risks by providing the right amount of protection or service to the organization without being a burden.

Education Process

Education is the key to any successful implementation of a new model, process, application, or tool. Business partners or customers need to understand what risk assessment can and cannot do for the organization. And, they need to understand what the tools can and cannot do. They also need to understand their part in the process.

For example, if you use a tool, such as SPRINT from ISF, the business areas need to understand that the tool only provides relative criticality ratings based on a subjective discovery process. During the evaluation of “business impact”, ratings are provided for asking some business professionals their opinion. Likewise, “threats and vulnerabilities” are measured by asking security professionals their opinion. Granted, an opinion by a professional is better than a guess by a non-professional. However, everyone should have a clear understanding that the ratings are mere guidelines and not something that should be taken as Truth.

There should also be a clear understanding of the usefulness and limitations of a tool. Not every tool will provide a clear risk assessment for every situation. Most of the tools provide output from subjective and sometimes arbitrary evaluations. Trying to bend the tool to provide an assessment for which it was not designed is akin to taking a low end car off-road driving. The car is not very high quality to start and it was definitely not designed for off-roading.

There also needs to be an education process for the security professionals who will be using the new risk assessment method. This includes using the tool(s), working with the business partners/customers, education of risk mitigation theories if necessary, etc.

Security practitioners, professionals, analysts, or employees need to understand that risk assessment is not meant to be a check-list / cookie-cutter approach to

security. After all, we are really just practicing security, like doctors and lawyers. Solutions can be more secure, but never completely secure. Security professionals must adapt and change to new threats, new impacts, and new technologies. Their expertise is needed to analyze, design, and implement appropriate controls that enable the business to operate efficiently and effectively.

And finally, there needs to be a significant amount of education for information security management, quality control, and/or auditing groups. Clear and accurate measurements are needed that determine if the solutions implemented by the security professionals are “better” than if a risk assessment had not been performed. I have seen situations where the security professionals were already implementing mitigating controls above and beyond best practices. There was never a case presented to those employees demonstrating how a risk assessment would help them implement better solutions. Even after a risk assessment was performed, the same solutions were implemented anyway. Therefore, it is important to demonstrate that risk assessment can help officially validate the need and expense for a solution if used correctly.

Acceptable quantifiable levels for business impact and likelihood

The impact to an organization if there is an event and the likelihood of that event occurring are both very abstract figures. Some organizations can put dollar amounts around impact because they know that when a system is unavailable they will lose x dollars every hour or minute. However, it is more difficult to place a dollar amount on the impact when the president of the organization is interviewed on a major news program because someone stole the social security numbers from their database and posted them on the Internet. Impacts to internal operations are easier to measure than an impact related with bad press or lawsuits. These are not impossible things to measure. However, trying to associate precise figures for these events is difficult.

Similarly, if you ask five security professional a seemingly easy question – what is the likelihood of someone stealing my credit card number if I submit it over the Internet without encrypting it? – and you’ll likely get a wide array of answers. Some tools help address these issues by providing a set of questions or variables to populate. However, they only indicate what variables should be used, not how to achieve the answers to those variables.

To have a useful risk assessment process, an organization needs to accept a certain level of ambiguity regarding the numbers associated with the assessments. Remember, the assessment should not produce a line in the sand for which the security professional should adhere. Rather, the assessment should serve as a starting point or a guide. Better analysis tools, methods, and data will be available in the future, but they are not here yet.

Baselines for key areas

What is the impact to an organization if there is a loss of confidentiality, integrity, or availability? Depends on who you ask. The business areas of the organization need to determine the criticality of their data. It should not be the responsibility of a security analyst to make this determination. If an organization cannot make this decision, then it should not waste time, money or resources on risk assessment because they will not provide value.

The baselines need to come from two areas: business and systems. The business area should help determine the criticality of their various data sets. The systems or IT area should come to the table with probability ratings for various data sets in a given context (business and technical environment). This can be tricky because the likelihood of an event occurring is related to the criticality of the data. The primary point is not that the systems area has a defined set of parameters and an associated probability factor. Rather, the business areas need to see consistency regarding the probably factors.

To establish a baseline, the criticality factors of the assets should be associated with the level of business area in your organization. For example, customer data should be assessed once or twice. When the data is planned to be used in a new application, the criticality is already determined, so the security professional can focus on any new threats and vulnerabilities affecting the likelihood of an event. This will prevent each security professional from contacting business management every time the data is used in a different application. This does not mean that criticality will not change, but that will be addressed by ongoing monitoring.

Predetermined risk tolerance levels

This is a fairly easy concept but very difficult to determine - where does an organization want to be in the rating grid (see Figures 1 and 2)? What are they willing to spend? In order to implement an appropriate solution, it would be very beneficial if the business areas / customers indicated what level of safety is required. This is often determined by the criticality of the data, but not always. It is also often determined by how much money is in the budget, but not always. If the business area or customer could indicate that they want to be in the "low" section of the rating grid and they had x amount of dollars and time, it would be easier to design mitigating controls. Not easy, just easier.

Risk acceptance procedures

Occasionally (maybe more often), there are certain mitigating controls that the business area does not want to implement because of time, money, or complexity. At which point, the security professional will need to properly communicate the risks associated with that decision to everyone who needs to know. This can be very tricky. However, the business areas or owners of the data need to be responsible and held accountable for accepting risks. When information is used by multiple parts of the organization, who is the owner? I would suggest forming a high level management group that has the authority to accept risk on behalf of all the parties involved.

Central repository for mitigating controls

Solutions are very rarely derived in a vacuum. Most organizations have a reusable list of mitigating controls that are used in various situations. Larger organizations may have more controls available. It is very helpful to the security professional to have a list of acceptable controls that can be used when developing a secure solution. For example, if a network session needed to be encrypted, it would be helpful to know an organization supported Microsoft's implementation of SSL but not SSL to an IBM Host.

Some may argue that a security professional should know what is, or is not, available in an environment. However, in larger companies, controls that are available may change frequently. A repository also helps the new analysts and any consulting associates that may be employed.

Effectiveness of mitigating controls

Assuming that we have relatively valid ratings for business impact and likelihood, the objective of a security professional is to reduce the likelihood to an acceptable level (See Figures 1 above). This means the goal is to establish a combination of security solutions that are appropriate for the given risk level within acceptable costs. We might be able to say two firewalls are better than one, but to what extent does having router ACLs, a firewall, and an intrusion detection system reduce the likelihood of a successful penetration? I am not suggesting that these numbers cannot be derived, albeit the numbers will be very subjective. I am suggesting that if an organization is going through the pain of a risk assessment process and the business areas are indicating to acceptable risk level, then the security professionals ought to ensure they can extrapolate the amount by which each set of controls reduces the risk in question so as to arrive at the desired level.

Cost, ease of implementation, and impact of controls

Other factors need to be considered in addition to analyzing the amount by which a control reduces the likelihood of an event. For example, implementation and ongoing support costs. The ease of the implementation – how much time and skill is needed (which may affect the implementation or support costs). And, to what extent will this solution impact the operations of the business. Will it affect a few or a lot of employees? To what extent will the employees be impacted? All of these questions need to be considered and addressed when providing mitigating controls based on risks.

Solution patterns

Most businesses have a certain set of business operations that are performed as a matter of routine. The operations usually require a certain set of people to handle certain types of data. While organizations are constantly changing and rearranging, there is often a point at which applications or infrastructures are residing in the middle of their lifecycle. Changes to these applications or infrastructures may or may not be very large. If the changes are not very large, it can be assumed that there is a set pattern for providing a secure solution. That is, the security professional knows the data set, the business context in which the data will be handled, and the associated risks (from the baselines). There have probably already been several similar solutions for that scenario. These solution patterns should be recognized and used as part of an ongoing risk mitigation plan. This does not discount possible changes that may be needed to address additional risk, but the patterns should help alleviate the need for every analyst to re-create a new secure architecture from scratch.

The goal is to establish patterns of secure solutions that do not need to have a full risk assessment performed every time. Changes to threat, vulnerabilities, and business impact will need to be reviewed, but most solutions will default to the implementation previously accepted by the business partners.

Risk Monitoring

Over time, the ratings will change for a variety of reasons, including new impacts to the business from regulations or lawsuits, and of course, changes in threats or vulnerabilities.

The risk assessment process or tool should take advantage of a risk management model that tracks baselines (see above) and correlates three primary measurements:

- 1) the latest impact to the business, including such items as evidence of due diligence
- 2) the latest probability or likelihood of an event occurring (i.e. data used from penetration tests or incident management),
- 3) the off-set cost ratio for mitigating/not mitigating the risk in question

Other data that should be used as part of an ongoing risk monitoring process includes:

- cost of an incident to the business line (actual, projected, labor hrs, etc).
- actual incidents compared to projected likelihood
- cost of controls, including support

Conclusion

There is no substitute for an experienced professional analyzing a situation and producing an appropriate secure solution. Risk assessment methodologies and tools, if implemented correctly, will help that professional by providing some insight regarding data criticality and providing a consistent means by which he or she can attain the necessary requirements.

The industry does not have the enough data readily available to effectively use most quantitative approaches. Subjective approaches are useful only when businesses realize that they are really subjective.

In the end, risk assessment is only as good as the data and opinions provided for the assessment and the professionals analyzing and producing the outcome.

References

Online

Briney, Andy. Information Security Magazine. *Parker's Plan*.

<http://infosecuritymag.techtarget.com/articles/1999/parker.shtml>

CRAMM. Website: <http://www.cramm.com/cramm.htm>

Dictionary.com. Website: (<http://dictionary.reference.com/search?q=risk>)

Horton, Le Grand, Murray, and Ozier. "Managing Information Security Risks - Part 1". *IT Audit*. Vol.3, August 15, 2000. Website:

<http://www.theiia.org/itaudit/index.cfm?fuseaction=forum&fid=250>

Information Security Forum. <http://www.securityforum.org>

International Organization for Standardization, *ISO/IEC 17799:2000(E)*.

<http://www.iso.ch/iso/en/CatalogueDetailPage.CatalogueDetail?CSNUMBER=33441&ICS1=35&ICS2=40&ICS3=>

National Institute of Standards and Technology, *Security Metrics Guide for Information Technology Systems* (July 2003). Website:

<http://csrc.nist.gov/publications/nistpubs>

Symantec – *Security Response: Glossary*. Website:

(<http://securityresponse.symantec.com/avcenter/refa.html>)

Books

Tipton, Harold F. & Krause, Micki. Information Security Management Handbook, 4th Edition. Chapter 15 – Risk Analysis and Assessment. New York (2000).

Fisch, Eric A. & White, Gregory B. Secure Computers and Networks. CRC Press, Florida (2000).

Excerpt from "Quantitative Risk Analysis Step-By-Step"

Ding Tan