Chapter 9

Security Risk Analysis

The fourth phase of a security risk assessment is security risk analysis. The security risk analysis depends on all the previous stages to supply the information required to analyze the security risk to the organization. The security risk assessment phase consists of techniques and approaches for determining individual and overall security risk levels. This process can take many different forms, depending upon the security risk assessment method performed. The security risk assessment process will be discussed here by describing the process in the following three steps:

1. Determining security risk
2. Creating security risk statements
3. Team review of security risk statements

Each of these steps is described in more detail in the sections that follow.

9.1 Determining Security Risk

The overall objective of the security risk assessment analysis process is to determine and convey the security risk to the organization's assets. While a composite security risk will be determined at a later point in the process, the objective here is to determine the security risk to the organization's assets based on threat/vulnerability pairings. The security risk determination therefore is dependent upon the identified threats and vulnerabilities measured, and based on the probability of the threat/vulnerability pair, the value of the asset affected, and the impact that the threat/vulnerability pair will have on the asset.

This information is obtained throughout the data-gathering phase of the security risk assessment. Once the team has all the data in order, the calculation of the security risk can be performed. The basic equation for security risk calculation is

$$\text{Risk} = \text{Assets} \times \text{Threat} \times \text{Vulnerability}$$

For several reasons, this simple equation is merely an illustration of the principle that security risk is calculated based on an understanding of the asset value, the extent of the threat, and the likelihood of the threat exploiting an existing vulnerability. Various security risk assessment approaches have different ways of specifying the security risk equation variables and of calculating their result.

The methods discussed in this book (e.g., data gathering, reporting, evidence collection) should apply equally well to any existing security risk assessment such as NIST 800-30, Facilitated Risk Analysis Process (FRAP), OCTAVE, and proprietary methods. This book is not intended to create yet another security risk assessment method, but to prepare teams and individuals to participate effectively in any security risk assessment effort.

The basic security risk equation simply computes the relationship between asset value, threat frequency, and vulnerability likelihood. This basic equation is illustrated in Figure 9.1. Because each of these areas has been previously covered in this book, it is tempting to think that determining security risk is a simple calculation. However, determination of the value of assets, the frequency of a threat, and the likelihood of a vulnerability is clouded by uncertainty.

9.1.1 Uncertainty and Reducing Uncertainty

When dealing with probabilities of threat and impact, we must recognize that the measurements we use (whether quantitative or qualitative) have an element of

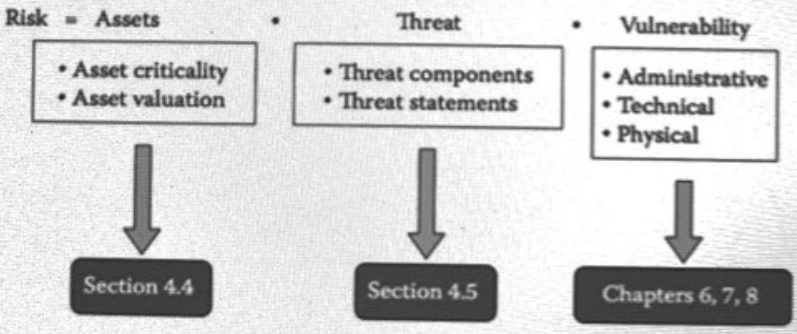

Figure 9.1 Basic security risk equation. The basic security risk equation computes the relationship between assets, threats, and vulnerabilities. Each of these areas has been covered in a previous section of this book.

uncertainty. The agreement as to the probability that a threat will compromise an asset can be an oversimplification of a more complex measurement. The question may now become, "How sure are you about that probability?"

To continue with a security risk assessment, the assessment team must either reduce or accept uncertainty (see Figure 9.2). In this section, we shall discuss various ways of reducing and controlling uncertainty.

SIDEBAR 9.1 Interpreting Requirements

Many security risk assessment projects include a requirement to compare the current security posture against a set of requirements, a regulation, or a standard. These will collectively be referred to as *requirements*.* Despite the intentions of those who develop the requirements, security compliance requirements rarely have a straightforward interpretation. Because of the ambiguity of the language, an interpretation process or professional judgment is required to resolve areas of confusion.

Some requirements, such as the "Common Criteria for Information Technology Security Evaluations," have a formal interpretation process. Such a formal process requires procedures for requests for interpretation, draft and formal rulings, and a catalog of previous rulings.

Other requirements simply depend upon the professional judgment of the security risk assessment team and the team leader. The judgment of adequate interpretations is based on the situation and an understanding of the intention of the requirement. Table 9.1 provides an example of the interpretation process.

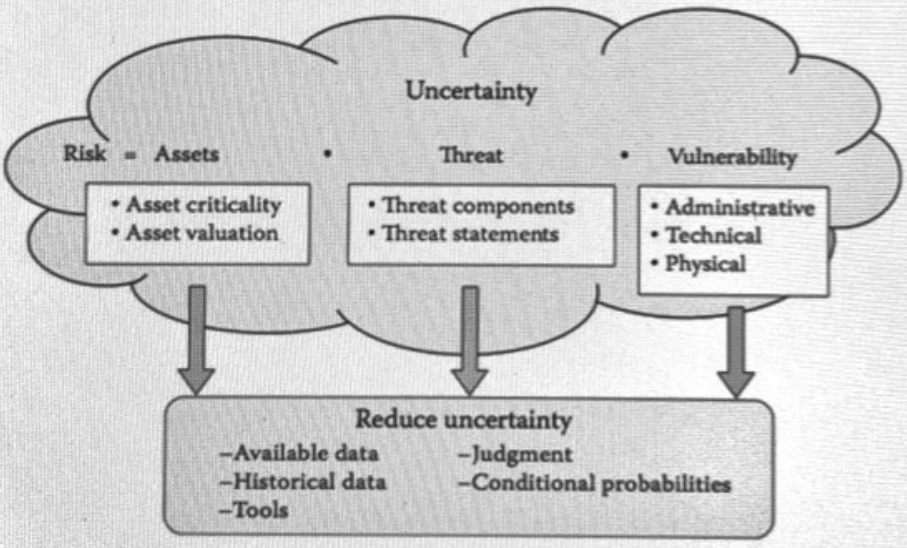

Figure 9.2 Introduction of uncertainty into determining security risk. The basic components of the security risk equation are clouded by uncertainty. Uncertainty may be reduced by gathering additional information or using consensus techniques.

* The term *requirements* is used loosely here to mean any statement within the standard, regulation, or guidance that speaks to the security controls that should be in the information system being assessed.

Table 9.1 Interpretation Process Example

Step	*Discussion/Example*
Requirement	Automatic log-off: Implement electronic procedures that terminate an electronic session after a predetermined period of inactivity. HIPAA 164.312(a)(2)(iii).
Interpretation	The organization must ensure that any session on an organizational information system with protected health information (PHI) is terminated after a reasonable period of inactivity. A reasonable period can be anywhere between 5 and 15 minutes for most sessions. There could be reasonable explanations for why a session would need to be inactive longer (e.g., a batch session that runs long and the system does not recognize the processing as an activity). The termination of the session can take place on any of the following elements of the session:
	• *The user terminal*. This can be a workstation within the organization's buildings or a remote computer.
	• *The media*. This could be a modem pool or a LAN, etc.
	• *The end system*. This could be an organizational information system with PHI.
Discussion	After interviewing the organization's systems administrators, it is found that workstations' default configuration currently locks after 15 minutes of inactivity, but this is not documented in policy and is not enforced.
	Session controllers have the capability to automatically disconnect the session after a period of inactivity but are not configured to do so. For some session users, this could produce a problem because the line looks inactive but a process could still be running. Use of the automatic log-out feature could impact the business mission for some users.
Findings	The organization *does not* currently meet this requirement because it does not automatically terminate idle and inactive remote sessions.

Table 9.1 Interpretation Process Example (Continued)

Step	*Discussion/Example*
Recommendation	Use of the system automatic disconnect feature would likely interfere with operations if it is set to a reasonable limit such as 5–15 minutes. However, the organization should set this control to terminate access after 4 hours of recorded inactivity. Because this control does not adequately protect unattended workstations, the organization should document and implement a policy to lock workstations after 5–15 minutes of inactivity.

Note: The interpretation process involves an interpretation of the requirement, a discussion of the environment and application, a finding, and a recommendation to the organization based on the interpretation.

9.1.1.1 Review Available Data

If any data is available on specific behavior—use it. Possible sources of data include crime statistics, analysis from previous security risk assessments, or knowledge regarding related industries. Be sure to adjust the data up or down for the specific circumstances of the organization and environment. For example, when determining possible employee theft, the team should review records of past employee theft at the organization or at similar organizations.

In some cases, the team may be able to find official statistics or actuarial information. Examples of this type of information are given extensively in Chapter 8. Information gathered may have to be extrapolated for the unique environment of the assessed organization. Examples of extrapolated information include reported frequency plus an estimated unreported frequency or official statistics modified by local conditions.

9.1.1.2 Examine Historical Data

Past events provide some indication as to the likelihood of similar events. For example, if employee theft occurred three times last year and no significant safeguards have been put in place since then, it may be reasonable to assume that employee theft has an annual rate of occurrence (ARO) of 3.0.

9.1.1.3 Use Judgment

Although other sources of information may provide some insight as to a number or rating that a specific threat should receive, it really comes down to the use of professional judgment. The security risk assessment team should document their reasons for each likelihood rating and be prepared to defend it to the security risk assessment sponsor.[1]

Various techniques exist for using judgment to reduce uncertainty. These include the Delphi technique, a decision-making process that polls experts individually and gradually works toward a group consensus, and bounding the problem:

- Bounding the Problem—When attempting to determine security risk component factors through the use of judgment, it is useful to first bound the problem with best- and worst-case scenarios. The team developing the values should use reasonableness when considering the best and worst cases.
- Develop a Probability Distribution—Once the problem has been bound, the team should develop a range of values with probabilities for each value. This is called a probability distribution.

Table 9.2 provides an example of reducing uncertainty through bounding the problem and developing a probability distribution. For example, it is reasonable to consider that a power outage would last a minimum of 5 minutes and a maximum of 2 days.[2] Most power outages are estimated to last less than 1 hour. Experts are polled to determine probable occurrences of various power outage times.

The use of probability distributions adds another level of accuracy in the estimation or measurement of security risk to an organization's assets. This additional information allows the use of statistics, probability functions, and mathematical modeling. However, the complexity of these calculations necessitates the use of tools.

Table 9.2 Probability Distribution for Duration of Power Outage

Upper and Lower Bound	*Range of Values*	*Probability Distribution*
5 minutes	5–10 minutes	7%
	10–15 minutes	8%
	15–30 minutes	34%
	30–60 minutes	18%
	1–2 hours	14%
	2–4 hours	8%
	8–16 hours	6%
	16–24 hours	3%
2 days	1–2 days	2%

Note: The use of judgment can be improved through the use of a range of values and probability distributions.

9.1.1.4 Use Tools

There are many tools available that can assist in the process of reducing uncertainty. Several examples of these tools include Microsoft Excel,[3] @RISK,[4] RiskWatch,[5] and BDSS.[6] Each of these tools has multiple capabilities to assist in the process of performing a security risk assessment. Microsoft Excel has built-in statistical functions; many plug-ins, such as StatPro and others, allow for the extension of Excel functions to include higher-level statistical modeling such as regression analysis and forecasting. @RISK allows for the modeling of uncertainty in security risk assessment decisions by replacing uncertain values with probability distributions. RiskWatch is security risk analysis software that can address complications, such as unavailability of information over time, and has built-in values for many standard vulnerabilities and threats. BDSS (Bayesian Decision Support System) integrates the concept of uncertainty into the risk calculations. This is just a small sampling of the many risk assessment tools available. The security risk assessment team should decide if such a tool would be useful to their effort.

9.1.1.5 Use Conditional Probabilities

Some threat/vulnerability pairs should be considered with respect to the chain of events that must occur for the threat/vulnerability pair to ultimately impact the security of the organization's assets. In such a chain of events, each event must occur for the next event to be considered. These are referred to as "conditional probabilities." It is the probability of the final event taking place (considering all the other events in the chain), not the first event, that is used as the probability for the threat/vulnerability pair.

When determining the probabilities of a threat/vulnerability pair having an impact on the organization's assets, consider the chain of events that must take place for the threat to exploit the vulnerability. For example, consider the threat/vulnerability pair of an ex-employee gaining access to your routers and switches by dialing in and using the administrator password. The probabilities of such an event may seem difficult to determine because there are many factors involved, but, by considering the chain of events, determining the probabilities becomes more tractable.

First, determine the events that must occur for the threat/vulnerability pair to impact the system. In this case, the employee would have to be terminated, have knowledge of the passwords, and attempt to access the system before the passwords are changed.

- Conditional Events for Example Threat:
  1. Employee is terminated.
  2. Employee has knowledge of passwords.
  3. Employee has desire to gain access.
  4. Passwords are not changed before employee attempts access.

Second, determine the probabilities of each event (see Table 9.3):

Table 9.3 Conditional Probabilities

Event Number	*Event*	*Probabilities*	*Discussion*
1	Employee is terminated	30 (30 times per year someone is terminated)	The organization (of roughly 1000 employees) has terminated an average of 30 employees a year over the last 5 years. Of those employees, 10% were terminated for cause.
2	Employee has knowledge of passwords	3 sys. admin. terminated per year 0.3 sys. admin. terminated for cause	Only system administrators have passwords to routers and switches. 11% of the terminated employees were system administrators (for-cause termination was proportional).
3	Employee has desire to gain access	0.75 sys. admin. 0.15 sys. admin. for cause	15% of terminated employees desire to gain access, but 25% among sys. admin. 50% of those terminated for cause desire to gain access.
4	Passwords are not changed in time	0.075 sys. admin. 0.015 sys. admin. for cause	Passwords are typically changed prior to the terminated employee leaving the building, but occasionally (10% of the time) personnel performing this duty are busy on other tasks and cannot get to it until the end of the day.
Annual expected breach		0.09	9% chance of this happening per year.

Note: Determining the probability of an ex-employee gaining access to the organization's routers and switches by dialing in and using the administrative password can be a rather difficult figure to develop. The use of conditional probabilities together with some known data reduces the complexity and uncertainty of such an estimate.

1. Terminated Employees—Continuing with this example, start with how many employees are likely to be terminated this year. In the absence of any other knowledge, such as a planned layoff or merger, historical records will provide a fairly accurate measurement. In our example, the organization of 1000 employees has terminated an average of 30 employees per year over the

last five years. Most terminations were friendly, such as the employee changing careers or moving away, but about 10 percent of the time the termination was for cause. At this stage we are considering 30 ex-employees per year.

2. Terminated System Administrators—Not all terminated employees would even have access to the router and switch passwords. Only system administrators have such knowledge. Of the terminated employees, 11 percent were system administrators. At this stage, we are considering 3.3 system administrators: 3.0 system administrators terminated and 0.3 terminated for cause.
3. Terminated System Administrators with Desire to Gain Access—Not all terminated system administrators would even want to gain access to the routers and switches of their ex-employer. After all, if they get caught, their career (and freedom) could be jeopardized. Here we must use some judgment. Based on experience and intuition, let us say the team comes up with the judgment that 25 percent of the system administrators would have the desire to attempt access. This may seem like a high number, but consider that (a) they might just be "checking" the system to see if the passwords were changed, (b) anyone who would catch them would likely be their friend and may not turn them in, and (c) the system administrators' knowledge of the system allows them to believe that they can get around without anyone detecting their presence. The percentage for system administrators terminated for cause increases to 50 percent because the same reasons (a)–(c) apply, but now we have reason (d): they are out to hurt the organization for hurting them. At this point in the chain, we are now considering 0.75 terminated system administrators with desire to access the routers and switches and 0.15 system administrators terminated for cause with the same desire.[7]
4. Terminated System Administrators with Desire and Means to Gain Access—Fully realizing the dangerous situation of terminated employees with sensitive passwords, this organization has procedures in place to change all passwords on systems if the passwords are known by an employee being terminated. The procedure calls for the system administrators to change the passwords as a part of the termination procedures. However, because system administrators are understaffed and overworked, the procedure is not always completed prior to the employee's leaving the building. About 10 percent of the time, the personnel performing this task cannot get to it until the end of the day. It is only in these circumstances that the terminated system administrator can access the system. At this final stage, we are now considering 0.075 terminated system administrators with the desire and means to access the routers and switches and 0.015 system administrators terminated for cause with the same desire. This gives us an annual expected breach of the system through this chain of events of 0.09. Put another way, there is a 9 percent chance of this scenario happening.

9.2 Creating Security Risk Statements

Between asset valuation, threat frequency, vulnerability probability, and impact affect, there are many values or numbers of which to keep track. If the security risk assessment team is using a tool, the tool can be used to keep track of and report the values and numbers. If the team is using a process without an automated tool, then an approach is required to track these values.

One such approach is the creation of security risk statements. A security risk statement is a method of presenting related information in the expression of a security risk. Three examples of security risk statements are given in Table 9.4. In the first example, the security risk statements are informal language expressions combining the threat agent, vulnerability, vulnerability target, policy violated, and the asset exposed. This simple set of security risk statements is useful in smaller-scale assessments where there are not numerous security risk statements to be made. Notice that these statements lack the ability to express the impact of the security risk, the likelihood of the scenario, existing security controls, overall security risk, and recommended solutions.

The second approach to developing security risk statements incorporates all of these components into a single row within a spreadsheet. This more complex approach has the advantage of documenting all of the constituent components of a security risk statement while still providing a reasonably understandable and compact format for what could be a complex set of information.

As with the other approaches presented in this book, these are just a few of the many approaches currently in use by information security professionals. Again, if the team is using a tool or other method, it is likely that the tool or method

Table 9.4 Example Risk Statement 1

Threat Agent	*Vulnerability*	*Vulnerability Target*	*Policy Violated*	*Asset Exposed*
A competitor	may social engineer	the sales office	to reveal	key customer lists
A hacker	may exploit known vulnerabilities	in the remote authentication protocol	to disrupt	remote authenticati on services
An intruder	may gain access	to the telephone closet	to eavesdrop on	sensitive conversations

Note: A security risk statement is a method of presenting related information in the expression of a security risk. This table provides several examples of security risk statements using sentence constructs for threat agents, vulnerabilities, policy violated, and asset exposed.

provides its own approach for consolidating and presenting this information (see Table 9.5).

9.3 Team Review of Security Risk Statements

Because of the large amount of data generally compiled during the data-gathering stage, it is a good idea for the security risk assessment team to divide up the task of creating security risk statements. Generally, the statements can be divided up along the areas of study, that is, administrative, physical, and technical. Further division can be accomplished by subdividing the technical areas according to systems or subgroups of the systems.

Team members should work alone or in small groups (e.g., two people) to create the security risk statements covering the data assigned to them. Once the draft statements are complete, the team leader should compile the complete list and distribute copies to all team members. The next task is for the entire team to review the draft statements and arrive at a consensus for the statements and the data values contained in the statements.

9.3.1 Obtaining Consensus

Arriving at a consensus for the elements within the security risk statements is an important step in the security risk analysis process. This step ensures that all members of the team have a chance to express their findings. Furthermore, obtaining a team consensus on the security risk statements allows all team members to gain a perspective of the overall security risk of the organization through a better understanding of all of the elements.

While obtaining consensus on these statements, the team should be wary of too much overlap. The following advice on obtaining security risk statement consensus may prove useful during this exercise:

- Avoid Overlap—While reviewing draft security risk statements, the team may find that some security risk statements may completely overlap or duplicate others. In this case, the statements should be reduced to a single statement. In fact, the team should institute some type of ordering of the statements (e.g., according to subject) so that all duplicates may be found.
- Group Like Findings—There are likely to be some components of the system with many findings. For example, it is typical to find many vulnerabilities within the external interfaces provided by the information system that can be addressed through the latest software patch. Although a security vulnerability tool will produce volumes of information on each of these findings, it is more useful to the customer to group these findings into relatively few findings with common recommendations, such as install latest patches.

Table 9.5 Example Risk Statement 2

				Impact (1–4)				Likelihood			
Ref. No.	*Area*	*Vulnerability*	*Threat Source*	*Conf.*	*Integrity*	*Availability*	*Existing Controls*	*Attempt (1–4)*	*Exploit (1–4)*	*Risk Recommendation*
8	Physical security	Telephone demarcation: access to telephone closet by guests	Cleared customers, visitors, and maintenance	Access to internal networks; access to internal phone conversations		Disruption of telephone service	Area restricted to cleared customers	3: Remote	1: Easy	1 Add locking cabinets to telephone and network connections in shared area
				1: System Disruption				C: Occasional		
32	Admin. security	Hackers may gain unauthorized access to information system through password guessing and social engineering	Hackers, social engineers, customers	Access to sensitive information	Threatens integrity of resources and sensitive information	Interruption of service, monetary losses, credibility losses	Employee awareness of sensitive information	2: Possible	2: Moderate	1 Create and enforce strong password policies, consider two-factor authentication
				1: Serious Risk. Exposure of sensitive information				8: Probable		Create and apply a strong security awareness program

56	Technical security	Hackers may gain unauthorized access through RPC call services	Hackers or disgruntled employees	Access to sensitive information	Modify sensitive data	Interruption of service, monetary losses, credibility losses	Some logging with minimal review	1: Likely	1: Easy	1 Create firewall rulesets, apply latest patches, harden servers, and install IDS	
				1: Serious Risk. Exposure of sensitive information				A: Frequent			

Note: A security risk statement is a method of presenting related information in the expression of a security risk. Above are several example security risk statements using a security risk assessment approach modified from Department of Defense Methods.

9.3.2 Deriving Overall Security Risk

Lastly, the security risk assessment team should derive an overall security risk. The overall security risk measurement should be consistent with the statement of work and the ranges used to describe individual security risks. For organizations regulated by information security laws such as the Health Insurance Portability and Accountability Act (HIPAA), the overall security risk should indicate a level of compliance. For all other organizations, the overall security risk level should indicate a relative security risk, for example, Moderate Security Risk, and a comparison to others in the same industry. The details provided in such a measurement do not need to be listed. This only has to provide the decision makers an indication of their current security risk in comparison to their security risk tolerance.

Exercises

1. Explain how the simple risk equation (Security Risk = Assets × Threats × Vulnerabilities) is implemented within the NIST SP 800-30 approach.
2. Given the probability distribution in Table 9.2, what statistical method could be used to determine security risk?
3. Compare and contrast any two security risk assessment methods or models.
4. What role does creating a security risk statement play in security risk analysis?
5. What data is currently available to assist in determining the probability of the following scenarios?
   a. A laptop going missing (stolen or lost)
   b. A misplaced thumb drive
6. Requiring encryption on removable media reduces which element of the threat equation?

Notes

1. This last approach for determining the likelihood of an event is one of the main reasons many security professionals opt for the qualitative security risk assessment approach. This is a realization that, even if you choose a number such as 3.0, it is still based on judgment and therefore is subjective.
2. An example where this would be a reasonable assumption is in an organization that has a UPS and 2 days of fuel for on-site generators. A power outage of less than 5 minutes can be handled by a UPS. A power outage of more than 2 days would, in all likelihood, be the result of a major natural disaster and require another level of support.
3. Excel is a trademark of the Microsoft Corporation.

4. @RISK is a trademark of the Palisade Corporation.
5. RiskWatch is a trademark of Expert Systems Software, Inc.
6. BDSS (Bayesian Decision Support System) is a trademark of OPA, Inc.
7. Expressions such as "0.75 terminated system administrators" may seem strange, because we typically discuss humans in whole numbers. However, this expression is interpreted to mean "on average, we would expect 0.75 terminated system administrators to attempt to gain access per year," or put another way, there is a 75% chance that at least 1 terminated system administration will attempt to gain access this year.

References

Common Criteria for Information Technology Security Evaluation. Version 2.1, CCIMB-99-031, August 1999.

Department of Health and Human Services, Centers for Medicare and Medicaid Services, CMS Information Security Risk Assessment Methodology. Version 1.1, September 2002.

Environmental Risk Management Authority. *Approach to Risk: Position Paper on the Approach to Risk, Methodologies Dealing with This and the Technical and Community Information Required for Implementation.* New Zealand, December 2002, ER-OP-03-02 12/02.

Environmental Risk Management Authority, *Preparing Information on Risks, Costs and Benefits for Applications Under the Hazardous Substances and New Organisms Act 1996.* New Zealand, July 2000, ERTG-03-01 07/00.

Krause, Micki, and Harold F. Tipton. *Handbook of Information Security Management.*

National Bureau of Standards. *Guidelines for Automatic Data Processing Physical Security and Risk Management.* Federal Information Processing Standards Publication 31, FIPS Pub. 31, June 1974.