

Information Security and Risk Management

Thomas M. Chen
Dept. of Electrical Engineering
SMU, Dallas, Texas

Introduction

It is easy to find news reports of incidents where an organization's security has been compromised. For example, a laptop was lost or stolen, or a private server was accessed. These incidents are noteworthy because confidential data might have been lost. Modern society depends on the trusted storage, transmission, and consumption of information. Information is a valuable asset that is expected to be protected.

Information security is often considered to consist of confidentiality, integrity, availability, and accountability (Blakley, McDermott, and Geer, 2002). Confidentiality is the protection of information against theft and eavesdropping. Integrity is the protection of information against unauthorized modification and masquerade. Availability refers to dependable access of users to authorized information, particularly in light of attacks such as denial of service against information systems. Accountability is the assignment of responsibilities and traceability of actions to all involved parties.

Naturally, any organization has limited resources to dedicate to information security. An organization's limited resources must be balanced against the value of its information assets and the possible threats against them. It is often said that information security is essentially a problem of risk management (Schneier, 2000). It is unreasonable to believe that all valuable information can be kept perfectly safe against all attacks (Decker, 2001). An attacker with unlimited determination and resources can accomplish anything. Given any defenses, there will always exist a possibility of successful compromise. Instead of eliminating all risks, a more practical approach is to strategically

craft security defenses to mitigate or minimize risks to acceptable levels. In order to accomplish this goal, it is necessary to perform a methodical risk analysis (Peltier, 2005). This chapter gives an overview of the risk management process.

Background

Risk management may be divided into the three processes shown in Figure 1 (NIST, 2002; Farahmand, Navathe, Sharp, and Enslow, 2003; Alberts and Dorofee, 2002; Vorster and Labuschagne, 2005). It should be noted that there is not universal agreement on these processes, but most views share the common elements of risk assessment and risk mitigation (Microsoft, 2004; Hoo, 2000). Risk assessment is generally done to understand the system storing and processing the valuable information, system vulnerabilities, possible threats, likely impact of those threats, and the risks posed to the system.




Figure 1. Steps in risk management.

Risk assessment would be simply an academic exercise without the process of risk mitigation. Risk mitigation is a strategic plan to prioritize the risks identified in risk assessment and take steps to selectively reduce the highest priority risks under the constraints of an organization's limited resources.

The third process is effectiveness assessment. The goal is to measure and verify that the objectives of risk mitigation have been met. If not, the steps in risk assessment and risk mitigation may have to be updated. Essentially, effectiveness assessment gives feedback to the first two processes to ensure correctness. Also, an organization's environment is not static. There should be a continual evaluation process to update the risk mitigation strategy with new information.

Risk Assessment

It is impossible to know for certain what attacks will happen. Risks are based on what might happen. Hence, risk depends on the likelihood of a threat. Also, a threat is not much of a risk if the protected system is not vulnerable to that threat or the potential loss is not significant. Risk is also a function of vulnerabilities and the expected impact of threats.

Risk assessment involves a number of steps to understand the value of assets, system vulnerabilities, possible threats, threat likelihoods, and expected impacts. An overview of the process is shown in Figure 2. Specific steps are described below.




Figure 2. Steps in risk assessment.

1. *System characterization*: It is obviously necessary to identify the information to protect, its value, and the elements of the system (hardware, software, networks, processes, people) that supports the storage, processing, and transmission of information. This is often referred to as the information technology (IT) system. In other words, the entire IT environment should be characterized in terms of assets, equipment, flow of information, and personnel responsibilities.

System characterization can be done through some combination of personnel interviews, questionnaires, reviews of documentation, on-site inspections, and automated scanning. A number of free and commercial scanning tools are available, such as Sam Spade, Cheops, CyberKit, NetScanTools, iNetTools, Nmap, Strobe, Netcat, and Winscan.

2. *Threat assessment*: It is not possible to devise a defense strategy without first understanding what to defend against (Decker, 2001). A threat is the potential for some damage or trouble to the IT environment. It is useful to identify the possible causes or sources of threats. Although malicious attacks by human sources may come to mind first, the sources of threats are not necessarily human. Sources can also be natural, for example, bad weather, floods, earthquakes, tornadoes, landslides, avalanches, etc. Sources can also be factors in the environment, such as power failures.

Of course, human threats are typically the most worrisome because malicious attacks will be driven by intelligence and strategy. Not all human threats have a malicious intention; for example, a threat might arise from negligence (such as forgetting to change a default computer account) or accident (perhaps misconfiguring a firewall to allow unwanted traffic, or unknowingly downloading malicious software).

Malicious human attackers are hard to categorize because their motivations and actions could vary widely (McClure, Scambray, and Kurtz, 2001). Broadly speaking, human attackers can be classified as internal or external. The stereotypical internal attacker is a disgruntled employee seeking revenge against the organization or a dishonest employee snooping for proprietary information or personal information belonging to other employees. In a way, internal attackers are the most worrisome because they presumably have direct access to an organization's valuable assets and perhaps have computer accounts with high user privileges (e.g., Unix root or Windows admin). In contrast, external attackers must penetrate an organization's defenses (such as firewalls) to gain access, and then would likely have difficulty gaining access with root or admin privileges. External attackers might include amateur "hackers" motivated by curiosity or ego, professional criminals looking for profit or theft, terrorists seeking destruction or extortion, military agents motivated by national interests, or industrial spies attempting to steal proprietary information for profit. External threats might even include automated malicious software, namely viruses and worms, that spread by themselves through the

Internet. It might be feasible to identify major external threats, but a possibility always exists for a new unknown external threat.

3. *Vulnerability analysis*: Threats should be viewed in the context of vulnerabilities. A vulnerability is a weakness that might be exploited. A threat is not practically important if the system is not vulnerable to that threat. For example, a threat to take advantage of a buffer overflow vulnerability unique to Windows95 would not be important to an organization without any Windows95 computers.

Technical vulnerabilities are perhaps the easiest to identify. Vendors of computing and networking equipment usually publish bulletins of bugs and vulnerabilities, along with patches, for their products. In addition, several Web sites such as Bugtraq (<http://www.securityfocus.com/archive/1>) and CERT (<http://www.cert.org/advisories>) maintain lists of security advisories about known vulnerabilities. It is common practice to use automated vulnerability scanning tools to assess an operational system. Several free and commercial vulnerability scanners are available, such as Satan, SARA, SAINT, and Nessus. These scanners essentially contain a database of known vulnerabilities and test a system for these vulnerabilities by probing. Another method to discover vulnerabilities in a system is penetration testing which simulates the actions of an attacker (NIST, 2003). The presumption is that active attacks will help to reveal weaknesses in system defenses.

Not all vulnerabilities are necessarily technical and well defined. Vulnerabilities might arise from security management. For example, human resources might be insufficient to cover all important security responsibilities, or personnel might be insufficiently trained. Security policies may be incomplete, exposing the system to possible compromise. Other vulnerabilities might be related to system operations. For example, suppose old data CDs are disposed in trash that is publicly accessible. It would be easy for anyone to retrieve discarded data.

4. *Impact analysis*: The impact of each threat on the organization depends on some uncertain factors: the likelihood of the threat occurring; the loss from a successful threat; and the frequency of

recurrence of the threat. In practice, these factors may be difficult to estimate, and there are various ways to estimate and combine them in an impact analysis. The impact analysis can range from completely qualitative (descriptive) to quantitative (mathematical) or anything between.

It would be ideal to estimate the exact probability of occurrence of each threat, but a rough estimate is more feasible and credible. The likelihood depends on the nature of the threat. For human threats, one must consider the attacker's motivation, capabilities, and resources. A rough estimation might classify threats into three levels: highly likely, moderately likely, or unlikely (NIST, 2002).

The loss from a successful threat obviously depends on the particular threat. The result may include loss of data confidentiality (unauthorized disclosure), loss of data integrity (unauthorized modification), or loss of availability (decreased system functionality). In financial terms, there is direct cost of lost assets and indirect costs associated with lost revenue, repair, lost productivity, and diminished reputation or confidence. Some losses may be difficult to quantify. Qualitative impact analysis might attempt to classify impacts into broad categories, such as: high impact, medium impact, and low impact. Alternatively, quantitative analysis attempts to associate a financial cost to a successful threat event, called a single loss expectancy (SLE). If the frequency of the threat can be determined (e.g., based on historical data), the product called annualized loss expectancy (ALE) is the product of the SLE and frequency (Blakley, McDermott, and Geer, 2002; NBS, 1975):

$$\text{ALE} = \text{SLE} \times (\text{annual rate of occurrence}).$$

5. *Risk determination*: For each threat, its likelihood can be multiplied by its impact to determine its risk level:

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

The most serious risks have both high likelihood and high impact. A high impact threat with a very low likelihood may not be worthy of attention, and likewise, a highly likely threat with low impact may also be viewed as less serious. Based on the product of likelihood and impact, each threat may be

classified into a number of threat levels. For example, a simple classification might be: high risk, medium risk, or low risk. Other classification approaches are obviously possible, such as a 0-10 scale (NIST, 2002).

The risk level reflects the priority of that risk. High risks should be given the most attention and most urgency in the next process of risk mitigation. Medium risks should also be addressed by risk mitigation but perhaps with less urgency. Finally, low risks might be acceptable without mitigation, or may be mitigated if there are sufficient resources.

Risk Mitigation

It may be safely assumed that any organization will have limited resources to devote to security. It is infeasible to defend against all possible threats. In addition, a certain level of risk may be acceptable. The process of risk mitigation is to strategically invest limited resources to change unacceptable risks into acceptable ones. Risk mitigation may be a combination of technical and non-technical changes. Technical changes involve security equipment (e.g., access controls, cryptography, firewalls, intrusion detection systems, physical security, antivirus software, audit trails, backups) and management of that equipment. Non-technical changes could include policy changes, user training, and security awareness.

Given the output from the risk assessment process, risks can be assumed or mitigated. Risk assumption refers to risks that are chosen to be accepted. Acceptable risks are generally the low risks, but a careful cost-benefit analysis should be done to decide which risks to accept. When risk mitigation is chosen, there are a number of different options (NIST, 2002):

- Risk avoidance attempts to eliminate the cause of risk, for example, eliminating the vulnerability or the possibility of the threat. For example, common software vulnerabilities may be remedied by applying up-to-date patches. So-called deterrent controls seek to reduce the

likelihood of a threat. Preventive controls try to eliminate vulnerabilities and thus prevent successful attacks.

- Risk limitation attempts to reduce the risk to an acceptable level, e.g., by implementing controls to reduce the impact or expected frequency. For example, firewalls and access controls can be hardened to make it more difficult for external attackers to gain access to an organization's private network. Corrective controls reduce the effect of an attack. Detective controls discover attacks and trigger corrective controls.
- Risk transference refers to reassigning the risk to another party. The most common method is insurance, which allows an organization to avoid the risk of potentially catastrophic loss in exchange for a fixed loss (payment of insurance premiums).

An overview of the steps in risk mitigation are shown in Figure 3. The steps are described below.




Figure 3. Steps in risk mitigation.

1. *Prioritize actions*: The risks with their corresponding levels identified through the risk assessment process will suggest what actions should be taken. Obviously, the risks with unacceptably high levels should be addressed with the greatest urgency. This step should identify a ranked list of actions needed to address the identified risks.

2. *Identify possible controls*: This step examines all possible actions to mitigate risks. Some controls will be more feasible or cost effective than others, but that determination is left for later. The result from this step is a list of control options for further study.

3. *Cost-benefit analysis*: The heart of risk mitigation is an examination of trade-offs between costs and benefits related to every control option (Gordon and Loeb, 2002; Mercuri, 2003). This step recognizes that an organization's resources are limited and should be spent in the most cost effective manner to reduce risks. A control is worthwhile only if its cost can be justified by the reduction in the level of risk. Not every cost may be easy to identify. Hardware and software costs are obvious. In addition, there may be costs for personnel training, time, additional human resources, and policy implementation. A control might also affect the efficiency of the IT system. For example, audit trails are valuable for monitoring system-level activities on clients and servers, but might slow down system performance. This would be an additional cost but difficult to quantify.

4. *Select controls for implementation*: The cost-benefit analysis from the previous step is used to decide which controls to implement to meet the organization's goals. Presumably, the recommended controls will require a budget, and the budget must be balanced against the organization's other budget demands. That is, the final selection of controls to implement depends not only on the action priorities (from step 1) but also on all competing priorities of the organization. It has been reported that companies spend only 0.047 percent of their revenue, on average, on security (Geer, Hoo, and Jaquith, 2003).

5. *Assign responsibilities*: Ultimately, implementation will depend on personnel with the appropriate skills. The personnel might be available within an organization, but for any number of reasons, an organization might decide to delegate responsibilities to a third party.

6. *Implementation*: In the final step, the selected controls must be implemented by the responsible personnel.

Effectiveness Evaluation

Effectiveness assessment is the process of measuring and verifying that the objectives of risk mitigation have been met. While risk assessment and risk mitigation are done at certain discrete times, the process of effectiveness evaluation should be continuously ongoing. As mentioned earlier, there are two practical reasons for this process in risk management.

First, risk assessment is not an exact science. There are uncertainties related to the real range of threats, likelihood of threats, impacts, and expected frequency. Similarly, in the risk mitigation process, there are uncertainties in the estimation of costs and benefits for each control option. The uncertainties may result in misjudgments in the risk mitigation plan. Hence, an assessment of the success or failure of the risk mitigation plan is necessary. It provides useful feedback into the process to ensure correctness.

Second, an organization's environment can not be expected to remain static. Over time, an organization's network, computers, software, personnel, policies, and priorities will all change. Risk assessment and risk mitigation should be repeated or updated periodically to keep current.

Future Trends

Today risk management is more of an art than a science due to the need in current methods to factor in quantities that are inherently uncertain or difficult to estimate. Also, there is more than one way to combine the factors to form a risk mitigation strategy. Consequently, there are several different methods used today, and none are demonstrably better than others. Organizations choose a risk management approach to suit their particular needs.

There is room to improve the estimation accuracy in current methods and increase the scientific basis for risk management. Also, it would be useful to have a way to compare different methods in an equitable manner.

Conclusion

Information security is an ongoing process to manage risks. One could say that risk management is essentially a decision making process. The risk assessment stage is the collection of information that is input into the decision. The risk mitigation stage is the actual decision making and implementation of the resulting strategy. The effectiveness evaluation is the continual feedback into the decision making.

Although current methods have room for improvement, risk management undoubtedly serves a valuable and practical function for organizations. Organizations are faced with many pressing needs, including security, and risk management provides a method to determine and justify allocation of limited resources to security needs.

References

Alberts, C., and Dorofee, A. (2002). *Managing information security risks: the OCTAVE approach*. Reading, MA: Addison Wesley.

Blakley, B., McDermott, E., and Geer, D. (2002). Information security is information risk management. In proc. of *ACM Workshop on New Security Paradigms (NSPW'01)*, 97-104.

Decker, R. (2001). *Key elements of a risk management approach*. GAO-02-150T, U.S. General Accounting Office.

Farahmand, F., Navathe, S., Sharp, G., and Enslow, P. (2003). Managing vulnerabilities of information systems to security incidents. In proc. of *ACM 2nd International Conf. on Entertainment Computing (ICEC 2003)*, 348-354.

Geer, D., Hoo, K., and Jaquith, A. (2003). Information security: why the future belongs to the quants. *IEEE Security and Privacy*, 1(4), 24-32.

Gordon, L, and Loeb, M. (2002). The economics of information security investment. *ACM Transactions on Information and System Security*, 5, 438-457.

Hoo, K. S. (2000). *How much is enough? A risk management approach to computer security*. Retrieved October 25, 2006, from <http://iis-db.stanford.edu/pubs/11900/soohoo.pdf>.

McClure, S., Scambray, J., and Kurtz, G. (2001). *Hacking Exposed: Network Security Secrets and Solutions*, 3rd ed. New York, NY: Osborne/McGraw-Hill.

Mercuri, R. (2003). Analyzing security costs. *Communications of the ACM*, 46, 15-18.

Microsoft. (2004). *The security risk management guide*. Retrieved October 25, 2006, from <http://www.microsoft.com/technet/security/topics/complianceandpolicies/secrisk/default.mspx>.

National Bureau of Standards. (1975). *Guidelines for Automatic Data Processing Risk Analysis*. FIPS PUB 65, U.S. General Printing Office.

National Institute of Standards and Technology. (2002). *Risk Management Guide for Information Technology Systems*, special publication 800-30.

National Institute of Standards and Technology. (2003). *Guideline on Network Security Testing*, special publication 800-42.

Peltier, T. (2005). *Information Security Risk Analysis*, 2nd ed. New York, NY: Auerbach Publications.

Schneier, B. (2000). *Secrets and Lies: Digital Security in a Networked World*. New York, NY: John Wiley & Sons.

Vorster, A., and Labuschagne, L. (2005). A framework for comparing different information security risk analysis methodologies. In proc. of *ACM Annual Research Conf. of the South African Institute of Computer Scientists and Information Technologists (SAICSIT 2005)*, 95-103.

Key Terms

Accountability: the assignment of responsibilities and traceability of actions to all involved parties.

Availability: the maintenance of dependable access of users to authorized information, particularly in light of attacks such as denial of service against information systems.

Confidentiality: the protection of information against theft and eavesdropping.

Integrity: the protection of information against unauthorized modification and masquerade.

Risk assessment: the process to understand the value of assets, system vulnerabilities, possible threats, threat likelihoods, and expected impacts.

Risk management: an organization's risk assessment and risk mitigation

Risk mitigation: the process to strategically invest limited resources to change unacceptable risks into acceptable ones.

Threat: the potential for some damage or trouble to an organization's information technology environment.

Vulnerability: a weakness or flaw in an organization's system that might be exploited to compromise security.