2 Security Management: Policies, Requirements, and Organizational Issues

2.1 Introduction

Information security is not strictly a technical problem. A wide range of management approaches are needed to develop organizational policies, roles, and training programs before choosing the security technologies. For a comprehensive security solution, enterprises need to develop security management approaches that:

- Establish general security policies and enforce them through controls and audits
- Develop organizational roles and responsibilities
- Institute security awareness and training programs
- Establish security requirements based on external (e.g., regulations) and internal considerations
- Conduct risk analysis to understand the vulnerabilities
- Develop circumventions and policies to mitigate risks

The main responsibility of managers is to develop an approach that secures the enterprise assets by developing and enforcing proper policies, controls, and audits that are essential without creating unnecessary bureaucracies. This chapter discusses these issues and exposes the reader to different aspects of security management and presents a quick summary of several security management approaches. Many articles on different aspects of security management can be found in the *Information Security Management Handbook* edited by H. Tipton and M. Kraus (4^{th} and 5^{th} ed., Auerbach, 2000). In addition, the basic concepts of audits and controls are introduced here. Chapter 13 treats this topic in more detail.

Chapter highlights

- Key steps in managing security:
  - Establish security policies (what needs to be secured, at what level, penalty for breach).
  - Develop procedures and guidelines to enforce the policies.
  - Select technologies to enable the above.
  - Put someone in charge.
  - Raise awareness of employees, managers, customers, because even the best-crafted policies and procedures fail if no one is aware of them (few read security documents).
- Policies are assertions about who can do what, when, and the consequences of non-compliance.
- Many policies need to be specified and enforced. British standard 7799 provides good guidelines for types of policies.
- Security requirements must consider internal (organizational) as well as external (legal, governmental) requirements.
- Risk assessment and mitigation is a key component of security management
- Several tools for automated risk analysis are becoming available
- Attack trees are a good approach for threat and vulnerability analysis
- Risks can be mitigated through policies as well as tools.
- Security technologies that support privacy, integrity, authorization, authentication, accountability, and availability (PIA4) mitigate risks.

2.2 Developing Policies, Controls, and Audits

2.2.1 Developing Security Policies

Security policies are assertions, contained in one or more documents, about how security controls will be implemented in an organization. Security policies are not technology-specific and serve the following goals to mitigate risks:

- Protect confidential and proprietary information from theft and unauthorized exposure/modifications.
- Reduce or eliminate legal liabilities to employees and third parties.
- Prevent waste of computing resources.

For example, a company can have a policy stating, "If an unauthorized employee is found to access, modify, or pass along highly confidential information, then he/she will be immediately terminated from employment." At Nissan, employees were sending sexually explicit emails to their colleagues. They were warned but continued this practice, resulting in their termination from Nissan. The employees sued, claiming email privacy, but Nissan won the case because it had an explicitly documented policy indicating that the employees did not have privacy rights on corporate email.

It may be a company policy to mitigate some risks by using *insurance*. For example, a company may feel that loss of some resources may result in a financial loss that should be covered through an insurance policy.

Policies like these clearly communicate company guidelines to employees and serve to protect the employees plus the corporation. Although security policies are needed for partners and customers, they are essential for internal controls. In particular, policies must be stated explicitly for the managers at all levels, in order to avert misuse. Many of these policies are part of company-wide auditing and accounting controls. Policies also drive the security measures taken in different areas (see Figure 2-1). Extreme measures and controls should be introduced for extreme situations. The main idea is that the more extreme the policies, the higher the costs are to maintain them. We will discuss auditing in a later chapter.

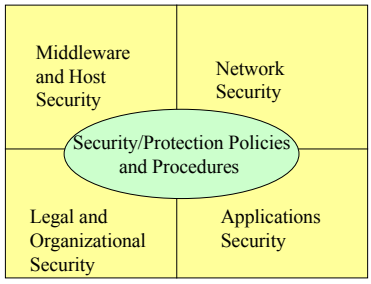

Figure 2-1: Security Policies and Procedure as Central Players

Policies, like many other things, have a life cycle – they are developed to mitigate risks, enforced through implementation of procedures, and monitored for compliance. The British Standard 7799/ISO 17799 also can be used in developing a security policy (see the sidebar, "British Standard 7799/ISO 17799)" and also the short case study in section 2.2.2). The "Site Security Handbook, RFC 2196" is also a good starting point to develop a security policy. Published in 1997, RFC 2196 is an update of the earlier "Site Security Handbook - RFC 1244" that was developed in 1991.

Good security policies are hard to find. A security policy is useless unless people read and understand the document. A good security policy is a complete, comprehensive and understandable document that clearly specifies:

- Who can use what resources
- Proper versus improper use of the resources
- Granting of access and use with indication of when an access should expire. (A consultant I know used a system three years ago and still has access to it, although he did not need to have access after the initial two-week consulting assignment.)
- System Administrator privileges
- User rights and responsibilities
- What to do with sensitive information
- Desired security configurations of systems

Once developed, the policies need to be translated into a practice by having them read and signed off by all employees. Unwritten or informal assertions that have not been read and signed off do not lead to good practices. Even well-written and signed-off policies do not guarantee compliance. It is always easier to write policies than to follow them, especially if you do not have to comply with them. It is easy to require difficult-to-guess passwords – but then how can you remember them? Good auditors know this, and so they check for sticky notes on the monitor, and look under the keyboard for passwords.

In the past, one centralized security policy was enough. At present, policies are needed for Internet access, email use, remote access from telecommuters, and other controls. Specifically:

- Remote access policies are needed to define how (dial-up, public Internet, virtual private network) and by whom the corporate resource will be accessed from remote sites.
- Internal use policies are needed to specify when and how the internal users will access the system resources. This also specifies whether the users can download and create local copies of sensitive data.
- User account policies define who can request an account at what authority level.
- PIA4 policies specify the privacy, integrity, authorization, authentication, accountability and availability (PIA4) of most valuable assets. For example, privacy policies indicate what information needs to be private and confidential; integrity policies may indicate what type of steps must be taken by whom to recover from loss of data integrity in a corporate database; and availability policies may specify how long a critical resource can be unavailable before emergency measures are triggered.
- Firewall management policies define who can access and modify the firewall rules and the controls to assure that the firewall is not modified by unauthorized personnel.
- Disaster recovery policies may indicate what actions are taken by whom to recover from a disaster. For example, as indicated in the opening case study about the World

Trade Center Disaster in chapter 1, Merrill Lynch was able to resume its business later in the same day because it had a good disaster recovery plan with detailed policies about how to handle disasters. The plan included priorities for business activities, names of personnel who would take over in case of a disaster, and even provisions for housing and feeding employees during the disaster period. This disaster recovery plan went into action within minutes after the incident, and Merrill was operational later that day.

British Standard 7799 (ISO 17799)

This standard is internationally known for policy specifications. The standard, known as BS7799 (also ISO 17799), is a set of recommendations organized in 10 major sections. These sections cover policies for areas such as business continuity planning, system access control, system development, physical and environmental security, personal security, security organization, computer and network security, and asset classification and control. For additional information, see www.securityauditor.net.

2.2.2 Short Case Study: Providing a Roadmap for IT Security by Using ISO 17799

St. Jude Medical is a $1.4 billion medical equipment manufacturer with operations around the globe. It is highly dependent on its computer resources and needed a policy that set the security expectations for its 3,500 computer users. David Stacy, the global IT security manager at St. Jude, faced the challenge of defining and implementing the company's first information security policy. He relied on the ISO 17799 to define the high level policies that were intended to have a long life.

Stacy wrote a security outline around eight of ISO 17799's ten sections. Among the areas adopted were requirements for a written security policy; IT asset management; access control; provisioning for services providers, security management and maintenance; disaster recovery and business continuity planning; and compliance review procedures. The physical and personnel security sections were not included because they had little practical value to his environment. After developing a draft in three weeks, Stacy and St. Jude's IT and security administrator spent about five months reviewing and revising the document. Stacy and his team presented company executives with the security policies document that was approved and sent to managers for implementation.

The role of a standard that can serve as a template for defining security policies was considered very valuable in this effort.

Source: Walsh, L., "ISO 17799 offered St. Jude Medical a road map for formalizing IT security," *Information Security Magazine*, March 2002.

2.2.3 Establishing a Control Environment

2.2.3.1 Overview

It is not enough to develop elegant policies and then forget about them. Controls and audits are needed to assure *continued* compliance with company policies. In general, policy statements and manuals are boring documents that collect dust in most offices. In a previous job, our group was handed a 300-page policies and procedures manual that we never had a need to open. A control environment is needed to actually implement the policies so that the policies are actively followed and not embedded in some manuals. In addition, audits are needed to assure that the polices continue to be followed. Audits and controls, discussed in more detail in chapter 13, are a good mechanism for implementing and monitoring important corporate policies.

Controls represent the combination of manual and automated measures that protect information systems and ensure that they perform according to corporate policies and procedures. Controls consist of all the policies, procedures, standards, and methods that ensure the safety of the organization's assets, the accuracy and reliability of its accounting records, and operational adherence to corporate standards. Weak controls can lead to major losses (see the sidebar, "Control Problems Cause a Loss of $750 Million").

Although controls have been used in IT for a number of years, there are some differences in the digital age. First, the focus of controls in the past was on internal systems and the employees who used these systems – management was considered a "trusted party." However, the numerous frauds incurred by managers of large corporations in recent history have changed all that. Second, the control of information systems in the past was an afterthought that was addressed only towards the end of implementation. In the digital age, organizations are critically dependent on information systems; thus vulnerabilities and control issues must be identified and addressed as early as possible. Finally, new information technologies (increased use of XML, Web Services, wireless systems, and application servers) raise new control issues that have not been traditionally thought of in the traditional control procedures.

In the past, most controls were defined in terms of application-specific and general controls.[1] This thinking needs to be refined a bit. Let us use the framework shown in Figure 2-2 to separate the application-specific, IT infrastructure, and management/process-specific controls. This framework and its components will be explained later (chapter 7), but the main idea is that the modern business applications and services rely on a deep technology stack that consists of numerous components from different suppliers. All these components need to be secured, controlled and audited. Specifically, a combination of controls are needed that take a complete systems view and include application controls, IT infrastructure controls, and administrative and process controls.

[1] Laudon and Laudon, "Management Information Systems," Prentice Hall, 8th Edition, 2003 (chapter 14).

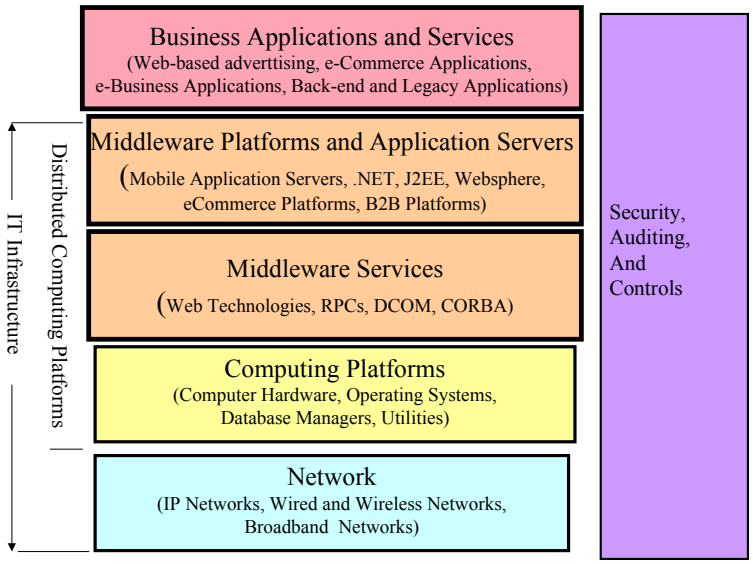

Figure 2-2: IT in the Digital Age

2.2.3.2 Application Controls

These controls are specific and unique to each computerized application, such as payroll, purchasing, and accounts receivable. Concerned with the higher layers of Figure 2-2, these controls include automated and manual procedures that ensure that only authorized data are completely and accurately processed by that application. For example, a control may be specified that no supplier be paid more than $10,000 per month. This control may be enforced by programming the accounts payable program so that it flags any check higher than $10,000 for a supplier.

Many application controls in the past concentrated on batch processing on mainframes. Widespread use of distributed systems with increased use of XML, Web Services, and mobile systems creates new application control issues. It is best to postpone this discussion to Chapter 13 after we have had an opportunity to review these technologies in Part III and IV of this book.

2.2.3.3 IT Infrastructure Controls

These controls, traditionally known as "general controls," span multiple technologies and are pervasive. They affect all applications supported by the organization's information technology infrastructure as shown in Figure 2-2. On the whole, these controls apply to all computerized applications and consist of a combination of hardware, systems software, and middleware services that create an overall control environment. Examples of these IT controls are corporate network access controls (i.e., who can access the network), database administration controls (i.e., who can have DB administrator authorities), and middleware controls (i.e., who can start up and shut down middleware platforms such as e-commerce servers). Once again, rapid advances in IT infrastructure raise control issues (see Chapter 13).

2.2.3.4 Administrative and Process Controls

These are formalized standards, rules, procedures, and control disciplines to ensure that the organization's IT infrastructure and application controls are properly executed and enforced. These are in fact "meta" controls – the controls for other controls. The most important administrative controls are (1) segregation of functions, (2) written policies and procedures, and (3) supervision. Process controls are concerned with the processes that exist in modern organizations. For example, the systems development process needs to be reviewed at various points to ensure that the process is properly controlled and managed. Computer operations controls apply to the work of the computer operations department and help ensure that automated and manual procedures are consistently and correctly applied to the storage and processing of data.

2.2.3.5 Costs and Benefits of Controls

Good controls are essential for secure operations and are necessary for compliance with laws and regulations. However, too many controls can be expensive and may lead to loss of productivity if used excessively. If all the control mechanisms at every level for everything in a digital firm were implemented, they could make the system economically or operationally unfeasible. Some common sense and cost/benefit analysis are needed to determine which control mechanisms provide the most effective safeguards without sacrificing operational efficiency or cost. The criteria may include importance of asset, efficiency, and risk mitigation (see chapter 13).

Control Problems Cause a Loss of $750 Million

In February 2002, the loss of $750 million in foreign exchange shocked the financial world. The loss was suffered by Allfirst, a subsidiary of Allied Irish Banks (AIB). This huge loss reminded many of the 1995 scandal in which a Singapore trader had hidden $1.4 billion in losses.

While the details of these losses are somewhat technical, requiring understanding of trading in foreign exchanges, the basic problem turned out to be lack of controls and oversight. In the case of Allfirst, the suspected trader was John Rusnak, who had worked with Allfirst for more than seven years. After a long investigation, AIB concluded that Rusnak had engaged in unauthorized trades and had falsified reports by taking advantage of his position. According to the AIB report, "Mr. Rusnak circumvented the controls that were intended to prevent any such fraud by manipulating the weak control environment at Allfirst." The weak controls included no separation of responsibilities, weak enforcemnt of reporting procedures, and faulty software. This fraud caused indictment of Rusnak, termination of many top managers at Allfirst, and a complete overhaul of control procedures.

Source: Laudon, K., and Laudon, J., *Management Information Systems*, 8th ed., Prentice Hall, 2003

2.2.4 Audits for Evaluating Controls

An audit, although scary to most of us, is a formal evaluation of one or more components of an organization against some goals. There are different types of audits. Examples of audits are income tax audits, financial audits, and efficiency audits. Two types of audits are of particular interest to us: IT audits and information security audits.

2.2.4.1 IT Audits

IT audits, also known as EDP audits, determine the effectiveness of controls needed for efficient and secure operations. An IT audit identifies the controls that govern individual information systems and assesses their effectiveness. To conduct such an audit, the auditor must acquire a thorough understanding of the applications, the IT infrastructure and the administrative procedures of the organization. During the audit, the auditor usually goes through numerous documents and interviews key individuals who use and operate a specific information system, concerning their activities and procedures. Different types of controls, ranging from application controls to administrative controls, are examined. The auditor typically traces the flow of sample transactions through the system and performs numerous tests, using automated audit software. The IT audit produces an audit report that lists and ranks all control weaknesses and estimates the probability of their occurrence. It then assesses the financial and organizational impact of each threat.

2.2.4.2 Information Security Audits

These audits determine how the privacy, integrity, and availability of an organization's information is being achieved and what can be done if it is not. This is also a formal assessment of how effectively the organization's security policies are being followed. Several companies have well written security policies that no one follows. On the other extreme, some companies have many security technologies but no security policies. Both are needed, as we will see. The security audits formally examine and evaluate the management as well as the technical aspects of a company's security system and the relevant policies. A security audit may concentrate on:

- Specific applications, especially those that handle financial transactions
- Firewalls that control the access to system resources
- Computer systems that host sensitive data and applications
- Corporate networks, intranets or extranets for secure access
- Crucial policies and procedures
- Physical support systems such as fire alarms, exits, electrical power systems, and emergency systems
- All or a mixture of the above

Due to its scope, security entails more than just a meeting; it involves planning, detailed investigations and interviews with key personnel, and an audit report that summarizes the findings.

See chapter 13 for a detailed discussion of security audits.

2.2.4.3 Security Policy as a Basis for Audits

A security audit is essentially an examination of how effectively an organization's security policy is being implemented. When performing an audit, the auditors typically start with an organization's security policy as a basis for assessment. The auditors

determine whether the policy document covers all the basic components of security elements – so the document should be comprehensive. Good policies that are complete and also written in plain English are very rare.

Unfortunately, many organizations do not have a written security policy. In these cases, the auditors typically define one or recommend one after the audit is complete. The "Site Security Handbook, RFC 2196" or the British Standard 7799 are good places to start.

Good auditors not only measure security policy compliance, but they also scrutinize the stated policy itself. In particular, they try to determine: is the policy too difficult to follow; does it concentrates on trivial or irrelevant issues; does it accurately reflect the practices; and do the employees know about it? A good security audit should result in a new and improved version of the policy.

2.3 Organizing and Training for Security

Beyond policies, controls, and audits, several organizational issues such as roles, responsibilities, and training programs are needed as part of a security management approach.

2.3.1 Organizing for Security: Roles and Responsibilities

A key element of a security management program is to put someone in charge of security. In most cases, information system security has not been a separate job but instead a job function that is attached to other jobs. For example, I have found database administrators in charge of database security, network administrators responsible for network security, and IT managers responsible for IT security. It is better to appoint:

- An information security coordinator for each business unit (BU) to coordinate the security activities of the BU.
- A Chief Security Officer for overall coordination at the enterprise level.

It is also important to separate duties; a single individual should not have complete control of everything. In addition, rotation of responsibilities to keep a fresh perspective is important. It is not a good idea to keep one person in one position too long. Rotation of responsibilities has some plusses and minuses. On the plus side, it keeps interest up, makes more people aware of the security issues as they are rotated through various responsibilities, and assures that security people do not bypass procedures for friends. For example, I have noticed that security guards in buildings do not check IDs diligently once they know that the person works there. It is possible for a person to have a revoked ID but still go in and out of buildings because he/she is a familiar face. Changing security guards periodically takes care of this problem. The main drawback of job rotation is that it may reduce job efficiency as new people need to be trained and are going through the learning curve.

Results of this process are usually represented in several documents with tables and charts. The following table shows an example of the main resources (physical site, databases, computers, networks) to be protected, the level of protection needed (based on

requirements), the person who will be responsible for the protection of the resource, and the training/awareness needed. Other columns can be added to this table.

Table 2-1: Sample Security Management Decision Table

	Level of protection Needed	Person Responsible	Training/awareness Needed
Resource1			
Resource2			
Resource 3			

2.3.2 Security Awareness

Security awareness is a vital aspect of security management. Even the best-crafted policies and procedures fail if no one is aware of them. In most cases, these policies and procedures are documented in thick security documents, and very few people actually read the security documents. Thus a good security program must include the following steps (see Peltier [2000]):

- Develop the security policies and procedures. Security policies include what needs to be secured at what level, and what is the penalty for breach. You also need to establish procedures and guidelines to enforce the policies and select technologies to enable the above.
- Make the people fully aware of the policies and procedures. This awareness should encompass employees and managers, as well as the customers.
- Organize security awareness days and drills. These events, if done correctly, can be very useful in explaining the importance of security and the roles played by policies, procedures, and staff in maintaining a vigilant security system.
- Institute random checks (after hours) to see if security policies are being followed. In an organization that I worked in, they checked to see if material marked proprietary and confidential was left openly in rooms and also if the desktops were properly "locked" (the screen saver required an ID and password to access the system). The next day, we had a green or red sticker on our doors to indicate if our areas were properly secure.
- Publish security alert newsletters to inform people about some real examples of security breaches and the actions taken. Once again, in my own experience, the security newsletter was published every quarter or so, was a one-pager, and had some interesting (some humorous) examples of security breaches or attempts. The newsletter indicated very clearly the actions taken by the company – in many cases termination of employment.

2.3.3 Security Training

Security training is an important element of a security program. To identify training needs, it is a good idea to document key security threats and then assess current levels of awareness about the threats. The "gap" between what is needed versus what is known should drive the training program.

It is essential to get a corporate buy-in before embarking on an ambitious security training program. Let us face it – while security is important, it is not very interesting to be trained in security and then to read the wonderful security documents. I have slept through several of them. Some type of buy-in and corporate incentive is typically needed to get the general organizational population engaged in security. Here are some ways to conduct security training somewhat successfully:

- A required course (lecture or computer-aided) that the employees have to take as a condition of continued employment.
- Provide incentives (e.g., part of performance appraisal) for security training. This method works quite well and I have used it several times as a manager. Employees who do not take a required training course in, for example, business ethics or security have something said about it in their yearly performance review.
- Self-awareness tests/questionnaires that are sent to employees who have plenty of time to do this.

2.3.4 Example: Government Agencies Beef Up IT Security

Due to the criticism of the federal government's security practices and policies, several agencies are evaluating themselves and beefing up IT security. Each year, the House Committee on Government Reform's Subcommittee on Technology, Information Policy, Intergovernmental Relations and the Census assigns security grades to several government agencies. The grades are based mainly on a set of criteria that mainly require that each agency inventory all its IT assets and be able to assess the security of each.

The Department of Justice (DOJ) is one of the agencies that received a failing grade on its Fall 2003 report card on IT security. In response, the DOJ has made a number of changes, including the establishment of a department-wide IT security staff that answers directly to the CIO. That group has organized a security council within the department that comprises the top security officials from each of Justice's dozens of organizations. The representatives to the security council belong to agencies such as the United States Attorney's Office; the Bureau of Alcohol, Tobacco, Firearms and Explosives; and the U.S. Marshals Service. The IT Security Council is responsible for implementing and overseeing all the security programs in the DOJ. This type of centralized IT security control is normal in large enterprises but is very new to federal agencies. So far, the results have been encouraging. The DOJ took the following steps to improve IT security:

- Created centralized security staff within the office of the CIO
- Built an IT Security Council to manage all security programs
- Established a department-wide automated security-evaluation and risk-mitigation program

In addition to DOJ, other departments are also adopting similar approaches. The Environmental Protection Agency (EPA), for example, has developed an automated security evaluation and remediation application for testing the security of each machine and monitoring the needed remediation process. The Department of Transportation (DOT) implemented a comprehensive vulnerability assessment and remediation package that performs continuous scans, instead of the traditional monthly or quarterly assessments. Both the EPA and the DOT have improved their grades in the 2003 report card.

Source: Fisher, D., "Agencies Beef Up IT Security," *eWeek*, January 5, 2004.

2.4 Establishing Security Requirements

To protect IT and physical assets, security requirements are established as the first step in building a security architecture. These requirements must include external as well as internal factors.

2.4.1 External Factors that Drive Security Requirements

Several national and international factors drive security initiatives in enterprises. For example, government regulations and consumer/customer attitudes towards security change due to major national or international events (e.g., the September 11, 2001 attack). These factors, not within the control of organizations, include:

- National and international emphasis such as homeland security
- Imminent or possible intruder/attacker threats on certain days (e.g., 4th of July weekend in 2002)
- Privacy and confidentiality laws and legal requirements imposed by the government (e.g., the heavy regulations in the healthcare industry)
- Consumer/customer attitudes towards security and privacy
- Threat models – "ankle biters versus national enemies"; protection against hacking versus assaults based on general industry (e.g., banks are more prone to attacks by thieves and military sites are more prone to espionage attacks)

2.4.2 Organizational Requirements

Many security requirements are specific to specific industry segments and organizations within the industry segments. For example, the airline industry has different security requirements than banks, and large international banks have different security requirements than small local banks (if there are any left!). The main idea is: how does security support your organizational goals? The main business driver for business security is the growing reliance of businesses on IT. The organizational requirements should:

- Clearly note the business drivers for security.
- Classify assets to be protected according to relative importance (not everything needs to be protected at the same level).
- Establish realistic survivability and intrusion tolerance requirements.
- Consider tradeoffs between QoS requirements and survivability/intrusion requirements.
- Keep budgetary and policy restrictions in mind because solving all problems can be expensive. The main challenge is: how do you balance costs versus benefits?
- Consider cultural situations while managing security in a global and multi-cultural environment.

2.4.3 Tradeoffs between Security and Availability

An important aspect of security requirements is understanding tradeoffs between the level of security and safety needed for various resources. Security and availability requirements can be specified by users or system administrators. For our purpose, we specify these in terms of two dimensions: system security and system availability (see Figure 2-3). The basic idea is that a highly protected system is highly secure and 100% available; i.e., it is intrusion-tolerant. Intrusion may be intentional or unintentional and involves a combination of security and fault-tolerance. There are several types of intrusions such as the following:

- Nothing modified – i.e., the intruder only looks but does not change anything
- Denial of service – i.e., the intrusion disallows service to legitimate users
- Modification to the system so that it behaves differently than intended
- Damage (permanent or temporary) to the system so that it is not recoverable
- Introduction of viruses

The protection policy chosen against possible intrusion threats can be represented as a matrix: (S, A) where S represents the security level chosen and A the availability. The security S is provided at the following levels:

- Level 0: No security specified
- Level 1: Authorization and authentication of principals
- Level 2: Auditing and encryption (Privacy)
- Level 3: Non-repudiation and delegation

Note that the security levels are inclusive (i.e., if you choose level 2 security it implies levels 0, 1, and 2). Availability A can be represented in terms of replications (more replications increase system availability):

- Level 0: No replication (i.e., only one copy of the resource is used).
- Level 1: Replication is used to increase availability. The resource is replicated for a fail-safe operation.
- Level 2: FRS (Fragmentation, Redundancy, Scattering) is used. FRS schemes split a resource (e.g., a catalog is broken into 4 fragments), replicate it, and scatter it around the network to achieve high availability and intrusion tolerance. Details about FRS can be found elsewhere [Deswarte 1988, Deswarte 1991].

Figure 2-3 shows how the security and availability levels can be mapped to protected systems. A protection policy can be chosen for each component of a system. For example, network access can have a different protection policy than a database. You can choose, for example, that a wireless network needs a protection policy (security level 2, availability level 1) while a database only needs a protection policy (security level 1, availability level 0).

Obviously, there are costs (time, efforts) associated with enforcing a given protection policy. Some protection policies are easier to enforce than others because COTS solutions are available. The policies involving FRS (Fragmentation, Replication, Scattering) require considerable effort because FRS schemes are still being developed.

It should be noted that the security and availability levels are being suggested here as a basis for discussion and analysis. Different security and availability levels can be introduced, if needed. The cells of Figure 2-3 can be used to represent the security requirements in terms of the type of protection level needed for different objects in the

system. For example, you can decide that the customer database needs a security level of 2 and an availability level of 1. Thus this figure can be used as a "scatter chart" that represents security requirements.

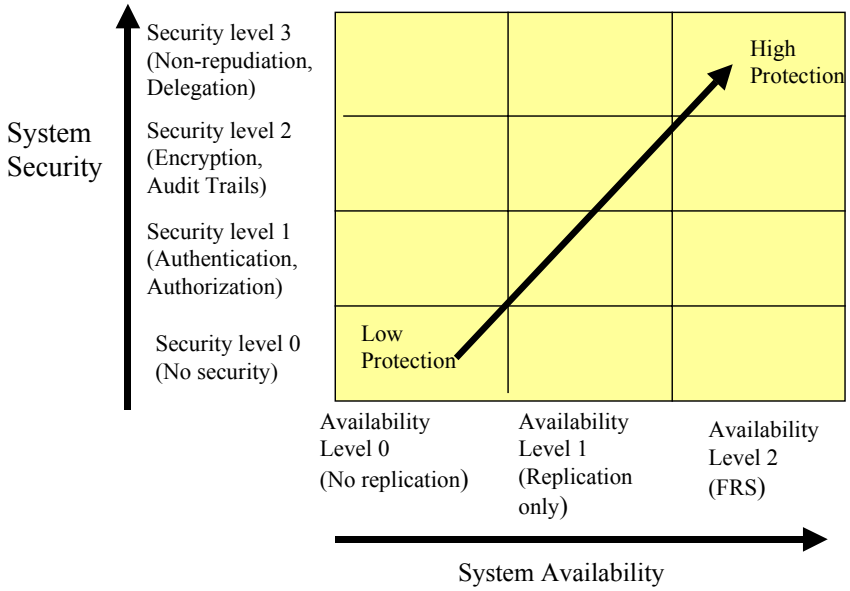

Figure 2-3: Protection Policies

What is FRS (Fragmentation, Replication, Scattering)?

The basic idea is that there is a tradeoff between security and availability. You can, for example, have a highly secure catalog at one site, but this does not improve the availability. You can use replication to increase availability, but this increases security threats (more copies are exposed).

The FRS schemes are supposed to address both issues. The FRS scheme [Deswarte 1988, Deswarte 1991] consists of three things. First, all sensitive information is cut into several *fragments* such that no significant information is contained in any fragment. Next, *redundancy* might be introduced by creating multiple copies of the fragments. Redundant copies tolerate accidental or purposeful destruction or alteration of fragments. To provide additional security and availability, the fragments along with their copies may be *scattered* among the different sites of the distributed system. This helps because if one site is compromised, the others can be used to successfully access information.

2.5 Risk and Vulnerability Management

2.5.1 Risk Management

According to Dictionary.com, risk is the "possibility of loss or harm." There are two elements of this simple definition. First is the "possibility" that represents uncertainty, and the second is the "loss" that signifies undesirable results. High-risk events are those with high possibility and greater loss. For example, police officers have a high risk job due to their frequent encounters with criminals. There are many other risks, some high, some low. For example, there is always a risk that your house may burn down, you may lose your job, or your loved ones may be involved in a serious accident. In our daily lives, we learn to live with a variety of risks and attempt to deal with them the best we can. In several cases, we attempt to minimize and manage risks by avoiding high-risk activities and buying appropriate insurance.

Establishing security can be viewed as an exercise in risk management – i.e., managing the risks associated with threats to the system. The basic idea of risk assessment is:

- What could happen (threat event)?
- If it happened, how bad could it be (threat impact)?
- How often could it happen (threat frequency, annualized)?
- How certain are the answers to the first three questions?

Based on assessment, a risk-mitigation approach needs to be developed. This entails asking:

- Are the risks acceptable?
- What can be done?
- How much will it cost (annualized expense)?
- Is it cost effective?
- Can risk be transferred (as through an insurance policy)?

The effort needed to mitigate risks should be proportional to the risk assessed. You do not want to spend millions of dollars to protect an asset that is worth two thousand dollars. It is a good idea to estimate a total expected loss, as follows:

$$\text{Total Expected Loss (TEL)} = L1xF1 + L2xF2 + \ldots. LnxFn$$

where:

L1, L2,,, Ln are losses expected with events 1 to n

F1, F2,,,, Fn are frequencies per year

TEL is typically annualized; i.e., if you lose $1M for an event that can happen once every 10 years, then annual TEL = $100K. The loss can be tangible (quantitative) or intangible (qualitative). The loss (L) for an event can be measured in terms of asset value (A) and exposure (E) per event, such that L = AxE. Consider, for example, a restaurant of value $10M that is being insured for fire. There is a 50% chance that a fire in the kitchen will burn the restaurant. Thus L = 10x0.5 = $5M. If this restaurant was to be insured for fire, the insurance agent will pay special attention to the kitchen because the L value for the kitchen is very high.

A great deal of literature exists on qualitative versus quantitative aspects of risk analysis. Elements that need to be estimated include:

- Asset value
- Threat frequency
- Threat frequency exposure
- Safeguarding cost
- Safeguarding effectiveness

Initial attempts at quantifying everything have not succeeded. Qualitative measures of Low, Medium, and High are commonly used. The main problem is that these measures can be highly subjective (your low may mean my high). A possible solution is to assign ranges; e.g., L < 10K and H > 100K.

Let us go through an example. Table 2-2 illustrates sample results of a risk assessment for a large department store that sells household goods to over 10,000 customers per day. The probability of a power failure occurring in a one-year period is 10 percent. Loss of business transactions while power is down could range from $10,000 to $100,000 for each occurrence, depending on how long processing is halted. The probability of fire is also 10% but the loss can be much higher. The probability of fraud and embezzlement by management or employees over a yearly period is about 5 percent, with potential losses ranging from $1,000 to $100,000 for each occurrence. Clerical errors occur more frequently (8%) with losses ranging from $10 to $20,000 for each occurrence. The average loss for each event can be multiplied by the probability of its occurrence to determine the expected annual loss. Once the risks have been assessed, analysts can concentrate on the control points with the greatest vulnerability and potential loss. In this case, controls should focus on ways to minimize the risk of fire, and then on clerical errors.

Table 2-2; Risk Analysis of a Discount Store

Exposure	Probability of Occurrence (%)	Loss range / Average	Expected Annual Loss
Power Failure	10%	$10,000-$100,000 ($55,000)	$5,500
Fire	10%	$100,000-$1,000,000 ($550,000)	$55,000
Fraud	5%	$1,000-$100,000 ($50,500)	$2,000
Clerical error	80%	$10-20,000 ($10,000)	$8,000

Steps in Risk Analysis and Mitigation

1. Assess the likelihood of each risk materializing
2. Assess the expected impact of each risk
3. Formulate measures to avoid risks

4. Establish controls and audits to periodically review and address the risks
5. Develop and deploy fallback measures to mitigate the risks if avoidance actions fail
6. Determine the urgency of the risk and take appropriate counter measures

2.5.2 Organizational Computing Models and Risks

The organizational computing models greatly impact the security risks and the policies and procedures needed to manage risks. Traditional computing models, introduced in the 1970s, are:

- Host-dependent (centralized): mainframes serve as the hosts of major applications and databases.
- Hierarchical: the control and authority flows from top (central) sites to lower sites.
- Closed: only known users can use the systems, and the managers can find out who is using the systems and when new users are added.
- Point-to-point: the connections between the various computers (mainframes, workstations, desktops) are well defined because leased point-to-point lines connect the computers together.
- Homogeneous: the computers and the software are largely supplied by one supplier (for example, IBM and/or Microsoft).

Modern enterprise computing models, on the other hand are:

- Decentralized and distributed: the data as well as programs are widely distributed among a variety of computers.
- Flat: there is very little top-down control. Traffic flows between computing devices via central controls.
- Open: the current Internet-based systems are quite open and it is very difficult to know who is communicating with whom.
- Broadcast: on most current networks, instead of the point-to-point connections between multiple users, messages are broadcasted over shared media.
- Heterogeneous: the current systems are very heterogeneous, with hardware and software from a multitude of suppliers.
- Mobile computing: the current systems are increasingly supporting mobile computing through wireless networks.

Naturally, it is easier to manage security in the traditional environments than in the modern Internet-based environments. In reality, most current organizations have a mixture of traditional mainframe-based and newer systems. The main challenge is to decide how to deal with these mixtures. The following guidelines may be used (see Murray [2000]):

- Single user name space and single user logon should be used to hide the multitudes of systems that the users need to access (see Vacca [2000]).
- Use strong authentication to assure that the unknown users in the open environment are who they should be.
- Keep the business service or application (user-visible entities) as high level points of control to keep focus on what is important.

- Use firewalls to localize issues and to segment the activities for control as much as possible.
- Protect keys and security infrastructure as diligently as the resources themselves.

2.5.3 Security Trust Models and Risks

In small companies, it is easy to know who you are communicating with. However, in large e-business environments, the players communicate with several parties in C2B, B2B, B2E, and other configurations over the Internet. In many cases, you are communicating with people outside of your corporate environment, including some you have never met, such as vendors, customers, clients, associates, and so on. Establishing a line of trust in such a setting is difficult and introduces many additional risks. Companies follow a range of *trust models,* such as the following, which dictate how users will exchange security credentials with each other.

- **Direct Trust.** Direct trust is the simplest trust model. In this model, a user trusts the partner because he or she knows the partner directly. This model is simple but does not extend to a large number of users in the Internet environments.
- **Hierarchical Trust.** In a hierarchical system, there are one or more third parties that are trusted by the users. This type of trust "tree" is used in many distributed systems.
- **Web of Trust.** A web of trust encompasses both of the other models and is a cumulative trust model. In other words, you trust my opinion that others are good and honest people only if you consider me to be a trusted person. Systems such as PGP (Pretty Good Privacy) use this model.

2.5.4 Attack Trees for Analyzing and Mitigating Risks

Attack trees are a convenient way to explore potential attacks and thoroughly examine the "attack space." An attack tree is simply a tree that is similar to a logical decision tree used to perform a systematic analysis of the attack space in terms of what is under attack, where the attack could happen, when the attack could take place and how the attack could happen.

To illustrate the key points, let us develop an attack tree against a physical safe (see Figure 2-4). The goal of attackers is to open the safe. To open the safe, attackers have several options: they can pick the lock, learn the combination, cut open the safe, or install the safe improperly so that they can easily open it later. Now you can assign values – I (impossible) and P (possible), in this figure – to the leaf nodes to indicate what needs to be considered next. You can now pursue the nodes that are possible for further evaluation. Let us now evaluate "learn the combination" node and break it into two activities: find the combination written down or get the combination from the safe owner through eavesdropping or other means. Each node becomes a subgoal, and children of that node are ways to achieve that subgoal.

In the attack trees, there are AND nodes and OR nodes (everything that is not an AND node is an OR node). OR nodes are alternatives while AND nodes represent steps toward achieving the goal. For example, to eavesdrop on someone for the safe combination, attackers have to eavesdrop on the conversation AND get safe owners to say the combination. Instead of "possible" and "impossible" values to the nodes, you can assign other values: easy, difficult, or very difficult; expensive versus inexpensive; intrusive

versus nonintrusive; legal versus illegal; special equipment required versus no special equipment. Assigning "expensive" and "not expensive" to nodes can help in analyzing if the asset is worth protecting, For example, if the asset is worth $10,000 and it takes $30,000 to steal it and $100,000 to protect it, then a decision has to be made about protection.

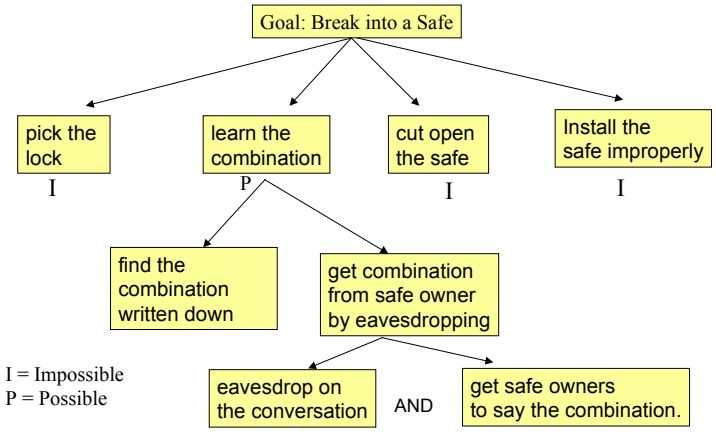

Figure 2-4: Sample Attack Tree for Opening a Physical Safe (Shneier 1999)

An important objective of the tree is to provide a heuristic for systematically considering attacks. When fully developed a tree should aim for completeness. However, since there are many possible attacks, it may not be feasible to put them into a tree. The leaves of attack trees are indicators of attacks. Each attack should be characterized in terms of technical and non-technical capabilities that are required, access needed to the system, risk assumed, and objective of the attack (see Schneier [1999] for additional discussion):

In large systems, it is not possible to draw large attack trees as graphs. In such cases, it is possible to describe the tree in terms of document sections and subsections. The following is an example of an attack tree for denial of service represented as a Microsoft Word document (outline mode). In this case, the tree is defined in terms of what, where, when and how. We have not assigned "possible" or "impossible" values to the nodes.

```
S1.    What. Service Assurance of System Threatened – Denial-of-
Service Attack on System
S1.1.  Where. Individual Component or Subsystem
S1.1.1.Where. End system
       S1.1.1.1. When. During network operation
              S1.1.1.1.1.          How. Passive attack
              S1.1.1.1.1.1. How. Eavesdropping
              S1.1.1.1.2.          How. Active attack
              S1.1.1.1.2.1. How. Illegal logon or system entry as
              user or root to cause denial of service
              S1.1.1.1.2.2. How. Dial-port flooding
              S1.1.1.1.2.3. How. Shutdown
       S1.1.1.2.     When.  During development …
S1.1.2.              Where. Router
S1.1.2.1.            When. During network operation
S1.1.2.1.1.          How. Passive attack
S1.1.2.1.1.1. How. Eavesdropping (preparing for DoS attacks)
S1.1.2.1.2.          How. Active attack
S1.1.2.1.2.1. How. Attacks on routing protocols
S1.1.2.2.            When. During development …
```

2.5.5 Automated Tools for Risk and Vulnerability Analysis

Many automated tools are commercially available at present, and more are announced on a regular basis, to analyze the existing computing environment in detail to identify potential risks. These tools, commonly known as *security audit tools*, scan the computing platforms as well as the interconnecting networks for possible viruses and weaknesses. After the analysis, the tools produce a variety of reports and analysis. Examples of these tools are:

- SAFESuite from Internet Security Systems (www.iss.net)
- Kane security analyzer from www.intrusion.com
- Cerberus Internet Scanner from www.cerberus.infosec.com
- NeuSECURE from Guardnet (http://www.guarded.net)
- Cortgeo from Trigeo (http://www.trigeo.com/)
- Network Audit from Visionel (http://www.visionael.com)
- QuVision from Q1Labs (http://www.q1labs.com)
- Security Threat manager from http://www.open.com).

Some research in automated tools for proactive analysis and automated generation of attack trees is underway. See, for example, the paper "Automated Generation and Analysis of Attack Graphs" by Sheyner, et. al. (http://www.cs.wisc.edu/~jha/jha-papers/security/oakland_2001.pdf).

Vulnerability Management Services Going Proactive

IT managers have traditionally reacted to specific security problems as and when they were discovered by the IT security industry. According to research from the Yankee Group (www.yankeegroup.com), this has changed and most IT managers now use proactive management and warning services on a regular basis. The Yankee report, entitled "Vulnerability Management: Processes Strengthen IT's Security Performance," states that proactive services help managers to reduce system downtime by allowing forward planning. Yankee's research also found that vulnerability-managed services are evolving beyond managed firewall and intrusion detection services by offering services that are aligned with business goals.

Source: *information security news*, January 2004, http://www.infosecnews.com/

2.6 Developing Circumventions and Mitigating Risks

There are several ways of mitigating risks in a security environment. For example, you can specify policies to clarify roles and responsibilities. In addition, a variety of tools and technologies such as encryption, password protections, digital signatures, and audit trails can be used to mitigate risks and circumvent vulnerabilities.

2.6.1 Security Policies to Mitigate Risks

Security policies, as indicated above, specify how security controls will be implemented in an organization. Security policies can be used to mitigate risks in the following ways:

- Clear assertions about the consequences of unauthorized access and intrusions mitigate risks. For example, dismissal of employees and prosecution of offenders are good deterrents to many intruders.
- Risks can be transferred by buying insurance to cover financial losses due to intrusion.
- Clarity in roles and responsibilities goes a long way in mitigating risks. It is extremely important to clearly specify who can get an authorized user account, who can modify firewall rules, who can access the system from dial-up, and what type of traffic will be accepted from the Internet.
- Explicit policies for the managers at all levels are essential to avert misuse.
- Extreme measures and controls should be only reserved for extreme situations because extreme policies are very expensive to maintain. For example, overuse of "urgent" messages dilutes the meaning of urgency.

It is not enough to develop and document policies. They need to be enforced and monitored diligently. For example, if the policy states that an employee will be dismissed if caught accessing unauthorized information, then the intrusions employees must be dismissed. Otherwise the policy has no value. Similarly, enforcement of policies should be monitored on a regular basis. Forced vacations and job rotations should be used to determine if the employees are doing their jobs according to the corporate policies. In many instances, frauds and non-compliances are found by replacement workers while the main workers are on vacation or other assignments.

2.6.2 Using Core Security Technologies for Circumventions

Many technologies can be employed to circumvent security vulnerabilities. Examples of the technologies are encryption, message digesting, password protection, digital signatures, and digital certificates.

Encryption has been used for a number of years to mask messages so that interveners cannot see/modify the messages. Due to e-commerce, encryption/decryption has become a major area of active work. In the simplest case, data is transformed by a key into an encrypted message. The encrypted message is then transmitted and decrypted on the other side by using the same key. Encryption/decryption can be performed by hardware and/or software. Modern computing systems have the ability to implement very sophisticated encryption/decryption techniques. The same encryption can be used on all data in a system, or encryption keys can be more "personalized." For example, instead of using the same encryption/decryption key on all data from all stations in a network, each station or user can use its own encryption/decryption key. A user can have his or her own encryption card which is inserted into a workstation before the user logs on. This card encrypts the data before sending it across the network. The encrypted data can be read only by those users or programs with access to the same encryption key. Encryption is generally discussed in two different formats:

- **Secret Key**: In a secret key (also known as symmetric or private) encryption scheme, the same key is used by the sender to encrypt the message, and the receiver

to decrypt it. While secret key encryption is usually very fast and efficient, the problem is with key management. In other words, since the sender and receiver have to agree on the same key, sending the key from one side to the other might compromise it.

- **Public key**: In a public key (asymmetric) system, the encryption key E and the decryption key D are different – hence the name "asymmetric." Each user has a pair of keys, a private key D that he keeps secret and a public key E that he publishes. When sender Bob needs to send a message to a user Joe, he encrypts the message with Joe's public key E(J). This encrypted message can only be decrypted with Joe's private key D(J). Therefore if this encrypted message is delivered to user Pat who does not have key D(J), then Pat cannot decrypt it. Thus when Joe receives the message, he can decrypt the message by using D(J) and read the message. Notice that in this key system the decryption key is private and not transmitted over the network. While public key systems solve the problem of key management, they are usually significantly slower than private key systems. The RSA (Rivest, Shamir and Adleman) algorithm, developed in 1976, is by far the most wildly used public key encryption algorithm.

Message Digesting is used to make sure that a certain message was not changed along the way between the sender and the receiver. A message digest algorithm produces a fingerprint of the message, by applying a hashing function to it. The receiver can check for the integrity of the message by reapplying the hash function and comparing the original fingerprint. The hash functions used in these schemes are such that the fingerprint changes dramatically if a single bit of the message changes.

User logon and password is one of the oldest and still most commonly used technologies. In this case, a system keeps track of who can access that system. This technology enables the use of existing systems with minimal disruption to existing infrastructure and applications.

Access control lists (ACLs) show what resources can be accessed by whom and for what purpose. For example, a database ACL may indicate that Joe can read this database, Sam can read and update the database, and Pat can delete the database (Pat is a database administrator).

Digital signature is used to authenticate the source of a message. It is essentially the same as a public key system except that the order in which the keys are applied is reversed. A sender "signs" the message by applying his private key to it. The sender sends the message and the signature to the receiver. The receiver checks the signature by applying the sender's public key to it. If the receiver gets the original message back, he is sure that the message was signed by the sender's private key, and therefore, was sent by the receiver himself. In essence, a digital signature is a block of data created by applying a cryptographic signing algorithm to some data using the signer's private key. Digital signatures may be used to authenticate the source of the message and to assure message recipients that no one has tampered with a message since the time it was sent by the signer.

A digital certificate binds an entity's identification to its public key and is issued by the Certification Authority. Digital certificates, based on the X.509v3 standard, enable Internet applications and other users to verify the identity of an entity. Unfortunately, certificates produced by one vendor product may not interoperate with other vendor's

because X.509 does not define the formats of the certificate entries and other necessary provisions. PKIX, the X.509 standard by IETF, defines the contents of public key certificates and is intended to resolve these interoperation issues.

Audit trails keep a log of each activity in a system. Most modern systems at present keep detailed logs, with time stamps, for each transaction (time when invoked, invoker ID, resource accessed, actions performed, time when completed, completion code, etc.). These logs serve as confirmation services for non-repudiation and provide digital receipts to indicate that messages were sent and/or received. Timestamps, for example, contain the date and time a document was composed and prove that a document existed at a certain time.

Fragmentation, Replication, and Scattering (FRS) schemes (Deswarte 1988, Deswarte 1991) consist of first cutting all the sensitive information into several fragments. This has to be done such that no significant information is contained in any isolated fragment. Then redundancy might be introduced by copying the fragments. The purpose of making copies is to tolerate accidental or purposeful destruction or alteration of fragments. Finally the fragments along with their copies are to be scattered amongst the different sites of the distributed system. This approach can be used to increase system availability as well as privacy.

We will look at these techniques in later chapters. As expected, different techniques address different aspects of risks and vulnerabilities. In particular, these techniques address different aspects of privacy, integrity, authentication, authorization, accountability, and availability (PIA4) attacks. This is illustrated in Table 2-3.

Table 2-3: Different Technologies Address Different Vulnerabilities

Vulnerabilities	Privacy attacks	Integrity attacks	Authentication attacks	Authorization	Accountability (NR)	Availability
Core Technologies	Encryption (private, public)	Message digests, CRC codes, error correction codes	Password protection, digital signatures, Digital certificates, biological identification	Access Control Lists (ACLs)	Digital Signatures, Audit trails	Fragmentation, replication, and Scattering (FRS)

2.7 Short Case Studies – Protecting the IT Assets from Internal Attacks [2]

The following cases are concerned with securing the environment within a company – i.e., how to secure the systems from damage by the insiders, not the outsiders.

[2] These case studies were collected by Christopher Freiler and Viktoryia Petrashova, students at Fordham Graduate School of Business

2.7.1 Global Retailer

Company: Global retailer X (because the company has a policy of not talking to the press, the CIO agreed to share all the following information on condition of anonymity).

Time of the case: January 2001.

Place: primary headquarters of the company outside New York City.

Description of the security problem: Hundreds of computers were infected with a stealth virus not recognized by the latest antivirus software. Later, the virus was identified as ***Demiurg***, the stealth virus that spreads through Microsoft Excel spreadsheets. When a user opens an infected spreadsheet, the virus infects the Windows Kernel32.dll file, a fundamental part of the operating system. When the computer is rebooted with the infected Kernel32.dll file, the virus spreads to executables and batch files, corrupting so many files that the computer eventually stops working.

What was at risk: the entire company network, in this case the existence of the company. Total estimated damage of lost sales was more than $250,000.

Problem categorization: Management issue. This case illustrates the importance and key role of the CIO in organizing the staff and solving the problem. However, it shows also that little was done before the accident to prevent a virus intrusion into the system. The only precaution was that the company was running the latest version of McAfee's antivirus software. The company did not seem to have a strict policy regarding opening files and e-mails from unknown respondents.

Solution approach: Time, in cases like this one, is the most important factor. Because IT staff became alert when the significant damage was already done, they tried to shut down the network as quickly as possible: through intercom orders, by posting flyers and by physically delivering the news. It took less than two hours to have the network shut down, involving more than 400 PC users. The result was no remote access for mobile users, no connection to offices in other countries, and no communication with stores (which could still ring sales and process credit card transactions but could not look up customer data or inventory at other locations).

The FBI was notified.

During the following four days the virus was identified. Then, with the help of a simple program, it was rendered inactive. Finally the system was disinfected. McAfee was notified during the process to provide any help in the problem's solution.

2.7.2 FBI

Company: Federal Bureau of Investigation (FBI).

Time of the case: From 1985 to 2000 (approximately).

Place: Alexandria, Virginia.

Description of the security problem: A career FBI agent with significant experience and access to FBI IT systems was charged with spying for Russia since 1985, in what FBI Director Louis Freeh has called the worst case of insider espionage in bureau history. The accused person, an expert in counterintelligence methods at the FBI, was

assigned to the New York Field Office's intelligence division in 1979 to help establish the FBI's automated counterintelligence database in that office. Investigators characterized him as having a "high degree of computer technology expertise." Although the accused was arrested while dropping off classified paper documents for his Russian handlers, he made extensive use of computer media, such as encrypted floppy disks, removable storage devices and a Palm II handheld computer, to communicate with Russian intelligence officers, according to the affidavit. In fact, he provided as many as 26 encrypted floppy disks during the course of his espionage activities.

What was at risk: national defense system.

Problem categorization: network and software security. A person or a group responsible for the security must be able to identify and understand the network and system intrusions. Also, artificial intelligence-enabled security software that uses profiles can tip administrators off to "anomalous activity" on the network.

Solution approach: We can only guess about what exactly has been done by FBI security to find the threat. What is known for sure is that the FBI keeps records on what sites every employee is using and also on his/her everyday activities; when it's needed, that information would lead to a special investigation. Also, in addition to these regulations, the order has been issued that a special panel be formed to review all FBI processes and systems and to study the issue of insider abuse.

2.7.3 Verizon

Company: Verizon Communications (a telecommunications company with subsidiaries that provide local telephone services in the region stretching from Maine to Virginia). At the time of the case, the company's name was GTE Corp., before the merger with Bell Atlantic Corp.

Time of the case: May 2001.

Place: Tampa, Florida.

Description of the security problem: An employee used his ability to gain access to GTE's secure computers at about 3 a.m. on May 15, 2001. Once he had access, he began to erase data contained in the computers and entered a command that prevented anyone from stopping the destruction process. Other IT workers at the GTE facility could not stop the self-destruction of the material once it had started.

What was at risk: customer database and confidential information that would be difficult (and costly) to recover. Total estimated damage: $200,000.

Problem categorization: Procedures. The accused had access to a secure area, and he, according to the record, severely abused those privileges. Therefore, the problem in this case can be identified as the problem of managing people within the security structure. It is also important to assign the right level of responsibilities, and to enforce them so that the lowest possible damage is incurred when they are used against the company. The challenge is to keep everybody happy simultaneously.

Solution approach: Obviously, there should be back up for the database information. Procedures can be implemented to secure information in the sensitive databases from deleting/transporting, by requiring signoffs from multiple people. The deleted

information can also be logged as part of the deletion procedure (most database managers support such logs) and stored somewhere for later use, if needed.

2.7.4 Ethical and Organizational Issues in these Cases

All three cases involve some ethical and organizational issues. In the Global Retailer virus case, the CIO questioned the intent of bringing the virus into the company's network (was it a competitor?). In the FBI case, the following was said by the FBI Director: "At the end of the day, all of our systems probably need to be looked at and maybe improved. But at the end of the day, what we rely upon is honest people." In the Verizon case, the ethical issue is obvious – what do you do when you cannot trust your employees?

The main point is that to develop security systems, companies should decide on the following organizational issues, in addition to the technologies: how to prevent staff from wrongdoing through clear guidelines, how to control the access to a system to minimize possible damage, and how to survive attacks from internal as well as external adversaries.

2.8 Chapter Summary

Security management is an important part of enterprise security. It is needed to protect the corporate IT and physical assets by using policies and employing the latest security technologies to respond to external factors and organizational requirements. This chapter has given a short overview of management issues with a discussion of policies, requirements, and risk assessment.

2.9 Review Questions and Exercises

1) What are the key ingredients of a security management approach? Which ones are absolutely essential and which ones are nice to have?

2) List some generic security requirements that apply to the current breed of e-business applications.

3) Choose a security package, install it, and conduct some very simple tests to understand how the package can be used to do risk analysis.

4) What is the importance of policies in establishing a secure environment and mitigating risks?

5) How can attack trees be used to identify risks, analyze them, and then mitigate them? Give an example.