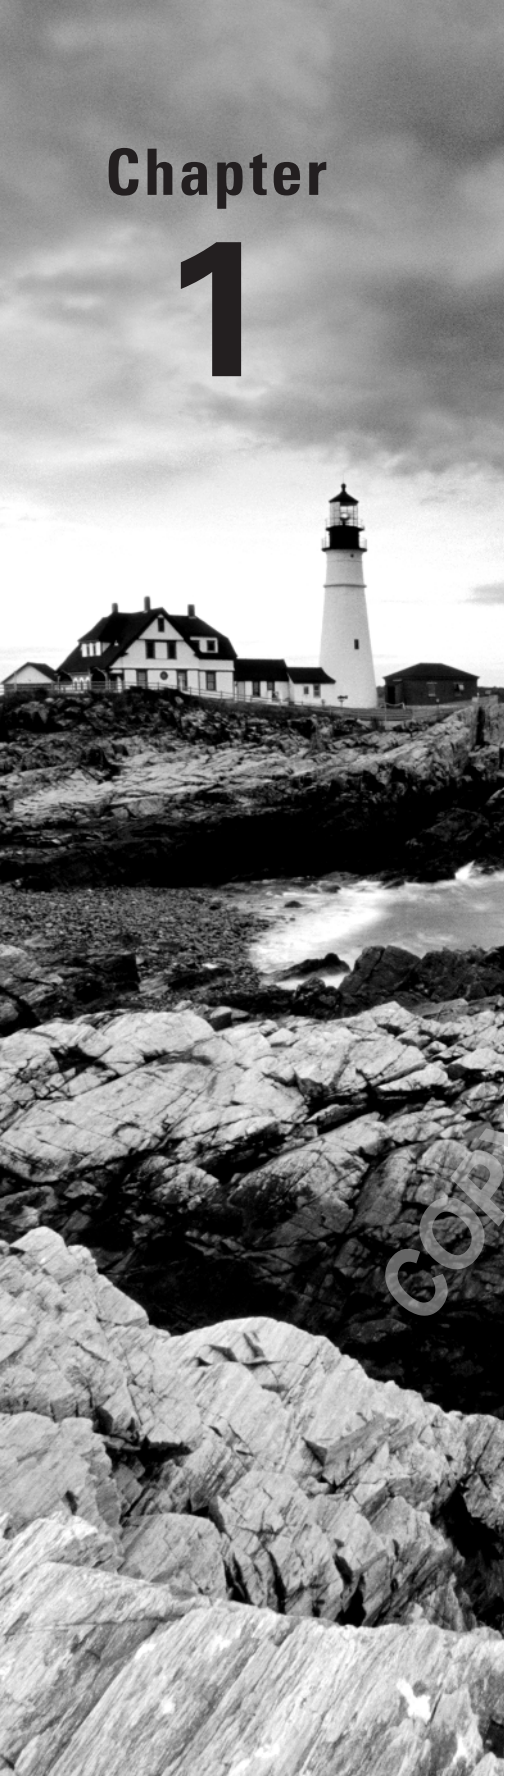

Chapter 1

General Security Concepts

THE FOLLOWING COMPTIA SECURITY+ EXAM OBJECTIVES ARE COVERED IN THIS CHAPTER:

- ✓ **1.6 Explain the purpose and application of virtualization technology.**
- ✓ **2.2 Distinguish between network design elements and components.**
  - DMZ
  - VLAN
  - NAT
  - Network interconnections
  - NAC
  - Subnetting
  - Telephony
- ✓ **3.7 Deploy various authentication models and identify the components of each.**
  - Biometric reader
  - Kerberos
  - CHAP
  - PAP
  - Mutual
- ✓ **3.8 Explain the difference between identification and authentication (identity proofing).**

Security is unlike any other topic in computing. To begin with, the word is so encompassing that it is impossible to know what you mean just by using it. When you talk about security, do you mean physical security of servers and workstations and protecting them from those who might try to steal them or from damage that might occur if the side of the building collapses? Or do you mean the security of data and protecting it from viruses and worms or from hackers and miscreants who have suddenly targeted you and have no other purpose in life than to keep you up at night? Or maybe security to you is the comfort that comes in knowing that you can restore files if a user accidentally deletes them.

The first problem with security is that it is next to impossible for everyone to agree on what it means because it can include all of these items. The next problem with security is that we don't really mean that we want things to be completely secured. If you wanted the customer list file to truly be secure, you would never put it on the server and make it available. It is on the server because you need to access it and so do 30 other people. In this sense, security means that only 30 people can get to it and not anyone outside of the select 30.

The next problem is that while everyone wants security, no one wants to be inconvenienced by it. To use an analogy, few are the travelers who do not feel safer by watching airport personnel frisk and pat down all who head to the terminal—they just don't want it to happen to them. This is true in computing as well; we all want to make sure data is accessed only by those who truly should be working with it, but we don't want to have to enter 12-digit passwords and submit to retinal scans.

As a computer security professional, you have to understand all of these concerns. You have to know that a great deal is expected of you but few users want to be hassled or inconvenienced by the measures you must put in place. You have a primary responsibility to protect and safeguard the information your organization uses. Many times that means educating your users and making certain they understand the "why" behind what is being implemented.

Security is a high-growth area in the computer industry, and it has been for several years now. The need for qualified people is increasing rapidly, as a search of job boards will quickly illustrate. Your pursuit of the Security+ certificate is a good first step in this process. Security+ is not the only security certification on the market, and it is not even the only entry-level certification available to you. It is, however, the only one to truly focus on the topics that most think of when security comes to mind. To pass it, you must have a broad knowledge of all the different types of security mentioned in the first paragraph.

In this chapter, I'll discuss the various aspects of computer security as they relate to your job. I will introduce the basics of computer security and provide several models you can use to understand the risks your organization faces. Not stopping there, I will also present steps you *must* take in order to minimize those risks.

Understanding Information Security

Information security narrows down the definition of *security*. The term *information security* covers a wide array of activities in an organization. It includes not only the products, but also the processes used to prevent unauthorized access to, modification of, and deletion of information. This area also involves protecting resources by preventing them from being disrupted by situations or attacks that may be largely beyond the control of the person responsible for information security.

From the perspective of a computer professional, you're dealing with issues that are much bigger than protecting computer systems from viruses. You're also protecting an organization's most valuable assets from people who are highly motivated to misuse those assets. Fortunately, most of them are outsiders who are trying to break in, but some of these people may already be inside your organization and discontented in their present situation. Not only do you have to keep outsiders out, but you have to be prepared for the accountant who has legitimate access to files and wants to strike out because he did not get as good a performance review as he thought he should.

Needless to say, this job isn't getting any easier. Weaknesses and vulnerabilities in most commercial systems are well known and documented, and more become known each day. Your adversaries can use search engines to find vulnerabilities on virtually any product or operating system. To learn how to exploit the most likely weaknesses that exist in a system, they can buy books on computer hacking, join newsgroups on the Internet, and access websites that offer explicit details. Some are doing it for profit or pleasure, but many are doing it just for the sheer thrill of it. There have been many glamorized characters on television and in movies who break into computer systems and do things they should not. When was the last time you saw a glamorized security administrator on such a show? If you make things look fun and exciting, there is some part of the audience that will attempt it.

Compounding matters, in many situations you'll find yourself constantly dealing with inherent weaknesses in the products you use and depend on. You can't count on the security within an application to be flawless from the moment it is released until the next version comes out three years later. The following sections discuss in detail the aspects you must consider in order to have a reasonable chance of securing your information, networks, and computers. Make sure you understand that I'm always talking about *reasonable*.

One of the first things you must develop as a security administrator is a bit of paranoia. It's important to remember that you're dealing with both system vulnerabilities and human vulnerabilities—although they aren't the same, they both affect the organization significantly. You must assume that you're under attack right now, even as you read this book.

Information security includes a number of topics of primary focus, each addressing different parts of computer security. An effective computer security plan and process must evaluate the risks and create strategies and methods to address them. The following sections focus on three such areas:

- Physical security
- Operational security
- Management and policies

Each of these areas is vital to ensure security in an organization. You can think of information security as a three-legged stool: If any one of the legs of your stool breaks, you'll fall down and hurt yourself. You must look at the overall business and address all the issues that business faces concerning computer security. Figure 1.1 shows how these three components of computer security interact to provide a reasonably secure environment.

FIGURE 1.1 The security triad

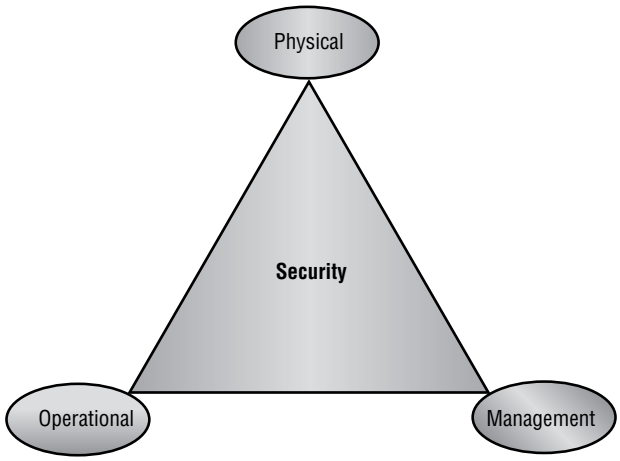

Part of your job is to make recommendations to management about needs and deficiencies; to take action to minimize the risks and exposure of your information and systems; and to establish, enforce, and maintain the security of the systems with which you work. This is not a small task, and you must do each and every one of these tasks well in order to have a reasonable chance of maintaining security in your organization.

Securing the Physical Environment

Physical security, as the name implies, involves protecting your assets and information from physical access by unauthorized persons. In other words, you're trying to protect items that can be seen, touched, and stolen. Threats often present themselves as service technicians, janitors, customers, vendors, or even employees. They can steal your equipment, damage it, or take documents from offices, garbage cans, or filing cabinets. Their motivation may be retribution for some perceived misgiving, a desire to steal your trade secrets to sell to a competitor as an act of vengeance, or just greed. They might steal $1,000 worth of hardware that they can sell to a friend for a fraction of that and have no concept of the value of the data stored on the hardware.

Physical security is relatively easy to accomplish. You can secure facilities by controlling access to the office, shredding unneeded documents, installing security systems, and limiting access to sensitive areas of the business. Most office buildings provide perimeter and corridor security during unoccupied hours, and it isn't difficult to implement commonsense

measures during occupied hours as well. Sometimes just having a person present—even if it's a guard who spends most of their time sleeping—can be all the deterrent needed to prevent petty thefts.

Many office complexes also offer roving security patrols, multiple lock access control methods, and electronic or password access. Typically, the facility managers handle these arrangements. They won't generally deal with internal security as it relates to your records, computer systems, and papers; that is your responsibility in most situations.

The first component of physical security involves making a physical location less tempting as a target. If the office or building you're in is open all the time, gaining entry into a business in the building is easy. You must prevent people from seeing your organization as a tempting target. Locking doors and installing surveillance or alarm systems can make a physical location a less desirable target. You can also add controls to elevators, requiring keys or badges in order to reach upper floors. Plenty of wide-open targets are available, involving less risk on the part of the people involved. Try to make your office not worth the trouble.

The second component of physical security involves detecting a *penetration* or theft. You want to know what was broken into, what is missing, and how the loss occurred. Passive videotape systems are one good way to obtain this information. Most retail environments routinely tape key areas of the business to identify how thefts occur and who was involved. These tapes are admissible as evidence in most courts. Law enforcement should be involved as soon as a penetration or theft occurs. More important from a deterrent standpoint, you should make it well known that you'll prosecute anyone caught in the act of theft to the fullest extent of the law. Making the video cameras as conspicuous as possible will deter many would-be criminals.

The third component of physical security involves recovering from a theft or loss of critical information or systems. How will the organization recover from the loss and get on with normal business? If a vandal destroyed your server room with a fire or flood, how long would it take your organization to get back into operation and return to full productivity?

Recovery involves a great deal of planning, thought, and testing. What would happen if the files containing all your bank accounts, purchase orders, and customer information became a pile of ashes in the middle of the smoldering ruins that used to be your office? Ideally, critical copies of records and inventories should be stored off-site in a secure facility.

Examining Operational Security

Operational security focuses on how your organization does that which it does. This includes computers, networks, and communications systems as well as the management of information. Operational security encompasses a large area, and as a security professional, you'll be primarily involved here more than any other area.

Operational security issues include network access control (NAC), authentication, and security topologies after the network installation is complete. Issues include the daily operations of the network, connections to other networks, backup plans, and recovery plans. In short, operational security encompasses everything that isn't related to design or physical security in your network. Instead of focusing on the physical components where the data is stored, such as the server, the focus is now on the topology and connections.

Real World Scenario

Survey Your Physical Environment

As a security administrator, you need to put yourself in the position of an intruder. For this exercise, think of yourself as an outsider who wants to gain access to the company server and damage it. Don't think of trying to steal data but rather of trying to pour water into the server. See if you can answer these questions:

1. How would you gain access to the building? Is a key or code required? Is there any security—a guard, a receptionist, or cameras? Are they highly visible, or does someone have to look to even know they are there?
2. How would you gain access to the floor the server is on? Is the elevator keyed, or can anyone use it? Do the doorways to the stairs only open outward, or can anyone walk up and enter?
3. How would you find the server? Is it sitting in the middle of the office, or is it in a separate room? If the latter, is the door to that room secured? How is it secured—by key, badge, punchpad?
4. After you reach the server, would anyone see what you're doing? Does the server room have glass windows? Is there a camera overlooking the server? Is the server viewable from a distance? Would anyone question why you were there?
5. If you do use cameras for surveillance, where are the tape machines? Are they located near the server so someone can steal the evidence of their crime as well?

If you can easily spot flaws in the security using these questions, then there is a risk that someone could harm your operations.

Finally, try to answer similar questions, but instead of imagining that you're an outsider to the company, use the perspective of someone from accounting who didn't get the promotion they thought they should and now wants to hurt the company. They have already gained access to much of the building—what keeps them from carrying out the crime?

Some vendors use the acronym NAC to signify network *admission* control rather than the more commonly accepted network *access* control.

The issues you address in an operational capacity can seem overwhelming at first. Many of the areas you'll address are vulnerabilities in the systems you use or weak or inadequate security policies. For example, if you implement a comprehensive password expiration policy, you can

require users to change their passwords every 30 or 60 days. If the system doesn't require password rotation, though (it allows the same passwords to be reused), you have a vulnerability that you may not be able to eliminate. A user can go through the motions of changing their password only to reenter the same value and keep it in use.

From an operational perspective, the system described has weak password-protection capabilities. There is nothing you can do, short of installing a higher-security logon process or replacing the operating system. Either solution may not be feasible given the costs, conversion times, and possible unwillingness of an organization—or its partners—to make this switch.

Such dependence on a weak system usually stems from the fact that most companies use software that was developed by third parties in order to save costs or meet compatibility requirements. These packages may require the use of a specific operating system. If that operating system has significant security problems or vulnerabilities, your duties will be mammoth because you'll still be responsible for providing security in that environment. For example, when your secure corporate network is connected to the Internet, it becomes subject to many potential vulnerabilities. You can install hardware and software to improve security, but management may decide these measures cost too much to implement. Again, operationally there may be little you can do.

Much of this book discusses the technologies and tools used to help ensure operational security. Figure 1.2 illustrates the various concerns you face from an operational perspective.

FIGURE 1.2 Operational security issues

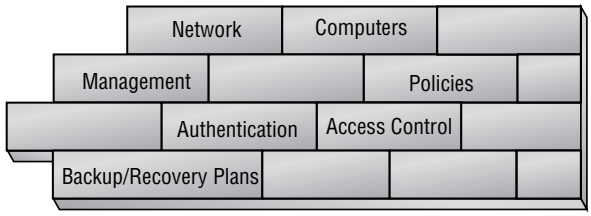

Working with Management and Policies

Management and *policies* provide the guidance, rules, and procedures for implementing a security environment. Policies, to be effective, must have the full and uncompromised support of the organization's management team. Management directions can give security initiatives the teeth they need to be effective. In the absence of support, even the best policies will be doomed to failure.

Information security professionals can recommend policies, but they need the support of management to implement them. There is nothing more ineffective than a self-proclaimed security "czar" who has no support from management. Not only is their tenure often short-lived, but so too is the security of their network.

Real World Scenario

Survey Your Operational Environment

As a security administrator, you'll need to assess the operational environment of your network by looking for "doors" that an outsider could use to gain access to your data. Securing the network involves far more than simply securing what exists within the four walls of your building. Look for openings that intruders can use to enter your network without walking through the door. Don't think of the safeguards that may currently exist, but rather focus on ways someone not on your network might join it.

See if you can answer these questions:

1. How do users on your network access the Internet? Do any users use dial-up connections within the office? Do they use dial-up access when they take their laptops home with them? Are proxy servers in use? Do you use private or public IP addresses? If you are using private IP addresses, are you using something as simple as Internet Connection Sharing or as complex as Network Address Translation (both perform the same function, but the latter offers more functionality and security)?
2. Are there wireless access points on the network? Can a mobile user with a laptop configure their settings to join the network? What is the range of your access points? Are signals stopped at the perimeter, or can someone sitting in the parking lot access the network?
3. Are dial-in connections allowed? Can users call in from home? Can they call in from hotel rooms? Do you verify the number they are calling from or merely allow anyone in with a correct username/password combination?
4. Do you use Terminal Services? Are thin clients employed/allowed? Are entire sessions on the server run remotely? Is remote administration enabled?
5. Do your users have shares on their laptops that would potentially compromise the laptop's data security?
6. What ports are open on your routers and firewalls (or on a user's personal firewall solution)?

The issues that must be decided at the management and policy level affect the entire company and can greatly impact productivity, morale, and corporate culture. Policies also establish expectations about security-related issues. Security policies should be treated no differently than an organization's vacation, sick leave, or termination policies. Most people can tell you exactly how many days of vacation they get per year; however, many can't tell

you what the company's information usage or security policies are. This can be solved by posting such information on an intranet or including it in a manual issued to all employees (with a note in each employee's personnel file indicating that they've received the manual).

A number of key policies are needed to secure a network. The following list identifies some broad areas that require thought and planning:

- Administrative policies
- Disaster recovery plans
- Information policies
- Security policies
- Software design requirements
- Usage policies
- User management policies

Administrative Policies

Administrative policies lay out guidelines and expectations for upgrades, monitoring, backups, and audits. System administrators and maintenance staff use these policies to conduct business. The policies should clearly outline how often and when upgrades appear, when and how monitoring occurs, and how logs are reviewed. They should also identify—not by name, but by title—who is responsible for making decisions on these matters and how often decisions should be reviewed. Ideally, the policies should also include information about who wrote them, who signed off on them, and at what date they were mandated.

The policies must be specific enough to help the administrative staff keep focused on the business of running the systems and networks. At the same time, they must be flexible enough to allow for emergencies and unforeseen circumstances. This trade-off is common to most policies, and you always want to be careful to avoid leaving a gap too wide, making the policy virtually ineffective or unenforceable.

Disaster Recovery Plans

Disaster recovery plans (DRPs) are one of the biggest headaches that IT professionals face. The DRP is expensive to develop and to test, and it must be kept current.

Many large companies invest huge amounts of money in DRPs, including backup or hot sites. A *hot site* is a facility designed to provide immediate availability in the event of a system or network failure. These sites are expensive to maintain and sometimes hard to justify. The likelihood that an organization will need a hot site is relatively small, and the site may seem unimportant—right up to the point when you don't have one and you need it.

A good DRP takes into consideration virtually every type of occurrence or failure possible. It may be as simple as a single system failing or as complicated as a large multinational company needing to recover from a cataclysmic event. The key to its success is its completeness. For example, if a company is located in the Midwest region of the United States, plans should be in place to address tornadoes, floods, fires, and every conceivable disaster.

Information Policies

Information policies refer to the various aspects of information security, including access, classifications, marking and storage, and the transmission and destruction of sensitive information. If your company records audio communications, that should be addressed as well.

The development of information policies is critical to security. It is not uncommon for such a policy to include a data classification matrix that defines various classification levels. The levels are usually similar to the following examples:

Public For all advertisements and information posted on the Web

Internal For all intranet-type information

Private Personnel records, client data, and so on

Confidential Public Key Infrastructure (PKI) information and other items restricted to all but those who must know them

The terms used for data classification might differ with different organizations—many used *top secret, secret,* and *sensitive,* for example—but the most important concept for the organization is that a matrix of levels exist.

As with all other policies, the key is to be as comprehensive as possible. Little should be left to chance or conjecture when you're writing information policies.

Security Policies

Security policies define the configuration of systems and networks, including the installation of software, hardware, and network connections. Security policies also define computer room and data center security as well as how identification and authentication (I&A) occurs. These policies determine how access control, audits, reports, and network connectivity are handled. Encryption and antivirus software are usually covered. Security policies also establish procedures and methods used for password selection, account expiration, failed logon attempts, and related areas.

Although each security policy is intended for a specific purpose, there may be scope overlap in many of the different policies. It is not uncommon as well to have overlap between information policies and security policies.

Software Design Requirements

Software design requirements outline what the capabilities of the system must be. These requirements are typically part of the initial design and greatly affect the solutions you can use. Many vendors will respond to every bid and assure you that they're secure. You can use the requirements to have vendors explain proposed solutions. A software design policy should be specific about security requirements. If your design doesn't include security as an integral part of the implementation, you can bet that your network has vulnerabilities.

Design requirements should be viewed as a moving target. The requirements that exist today shouldn't be the same in two years when the network environment has been significantly modified.

Usage Policies

Usage policies cover how information and resources are used. You need to explain to users how they can use organizational resources and for what purposes. These policies lay down the law about computer usage. Usage policies include statements about privacy, ownership, and the consequences of improper acts. Your usage policies should clearly explain usage expectations about the Internet, remote access, and e-mail.

They should also address how users should handle incidents—whom they should contact if they suspect something is awry. The policy should spell out the fact that monitoring can take place and that users agree to it. Consequences for account misuse, whether termination or something less severe, should also be stated.

User Management Policies

User management policies identify the various actions that must occur in the normal course of employee activities. These policies must address how new employees are added to the system as well as training, orientation, and equipment installation and configuration.

Employee transfers are a normal occurrence within a company. If an employee transfers to a new job, the privileges and access they had in the old position may be inappropriate for the new position. Establishing new access rights allows the employee to continue working. If you forget to revoke the old privileges, this user may have access to more information than they need. Over time, this can result in a situation called *privilege creep*. The user may acquire administrative privileges to the system by accident.

Terminated employees pose a threat to information security. In some cases, a terminated employee may seek to gain access to customer lists, bank accounts, or other sensitive information. When employees leave the company, it's imperative that their accounts be either disabled or deleted and that their access be turned off. You'd be amazed how often system administrators don't know about personnel changes. Your user management policies should clearly outline who notifies the IT department about employee terminations as well as when and how the notification occurs.

Real World Scenario

Assemble and Examine Your Procedures

It's surprising how many businesses think they have a policy in place when one can't be produced when needed. See if you can answer these questions:

1. Does your company have administrative policies in place? What are they, and where can they be found? Are they easily accessed by, or provided to, new employees? Does each written policy offer some indication of who to contact if there is a question or a breach?

2. When were the software design requirements last checked and/or updated? Are they routinely given to vendors? Who is responsible for reviewing them?
3. When was the last time the disaster recovery plan was checked? Do all administrators know it? Is it in writing and accessible from a remote location should this site become inaccessible?
4. Are informational policies easy to locate? By whom?
5. Are security policies updated frequently? Are they updated with each software change? Do they incorporate the latest patches?
6. Are usage policies part of the employee handbook? Do users sign off that they have seen the policies and are aware of them? How do users receive updates to the policies and signal that they have them and understand them? How do they know when those updates exist?
7. Can the user management policies be located and adhered to in the event that a situation occurs while the chief administrator is at a conference? Is there an escalation procedure in writing indicating who should be notified and when?

Policies not only need to exist, they also must be readily available so they can be referenced by all relevant parties. If this can't be said of the policies we've discussed, then their value is drastically diminished.

Understanding the Goals of Information Security

Like so many things, the goals of information security are straightforward. They create the framework that is used for developing and maintaining a security plan. They're remarkably easy to express but extremely hard to carry out. These goals are as follows:

Prevention *Prevention* refers to preventing computer or information violations from occurring; it is much easier to deal with violations before they occur than after. Security breaches are also referred to as *incidents*. When an incident occurs, it may be the result of a breakdown in security procedures.

Incidents come in all shapes and sizes. Simple incidents include things such as losing a password or leaving a terminal logged on overnight. They can also be quite complex and result in the involvement of local or federal law enforcement personnel. If a group of hackers were to attack and deface your website, you would consider this a major incident. Ideally, your

security procedures and policies would make you invulnerable to an attack; unfortunately, this isn't usually the case. The better your prevention policies, however, the lower the likelihood of a successful attack occurring.

Detection *Detection* refers to identifying events when they occur. Detection is difficult in many situations; an attack on your system may occur over a long period before it's successful. Incident detection involves identifying the assets under attack, how the incident occurred, and who carried it out (or is still doing so). The detection process may involve a variety of complicated tools or a simple examination of the system log files. Detection activities should be ongoing and part of your information security policies and procedures.

Response *Response* refers to developing strategies and techniques to deal with an attack or loss. Developing an appropriate response to an incident involves several factors. If the incident was a probe, the attacker may have done no actual harm but may be gathering intelligence about your network or systems. These types of attacks may be random or targeted, and they usually cause little damage. Occasionally, an attack will be successful. When that happens, it is helpful to have a well-thought-out and tested plan you can use to respond, restore operation, and neutralize the threat. It's always better to have a set of procedures and methods in place to recover from an incident than to try to create those processes on-the-fly.

These goals are an important part of setting benchmarks for an organization. You can't allow these policies or goals to become insignificant. If you do, you and your organization are setting yourselves up for a surprise. Unfortunately, the surprise won't be pleasant, and it may be very costly to deal with.

Comprehending the Security Process

It helps to think of security as a combination of three Ps: *processes*, *procedures*, and *policies*. The security of information involves both human and technical factors. The human factors are addressed by the policies that are enforced in the organization as well as the processes and procedures your organization has in place. The technology components include the tools you install on the systems you work with. There are several parts to this process, and each is described in the following sections.

Appreciating Antivirus Software

Computer *viruses*—applications that carry out malicious actions—are among the most annoying trends happening today. It seems that almost every day someone invents a new virus. Some of these viruses do nothing more than give you a big "gotcha." Others contaminate networks and wreak havoc on computer systems. A virus may act on your data or your operating system, but it's intent on doing harm—and doing so without your consent. Viruses often include replication as a primary objective and try to infect as many machines as they can, as quickly as possible.

The business of providing software to computer users to protect them from viruses has become a huge industry. Several very good and well-established suppliers of antivirus software exist, and new virus-protection methods come on the scene almost as fast as new viruses. Antivirus software scans a computer's memory, disk files, and incoming and outgoing e-mail. The software typically uses a virus definition file that is updated regularly by the manufacturer. If these files are kept up-to-date, the computer system will be relatively secure. Unfortunately, most people don't keep their virus definition files up-to-date. Users will exclaim that a new virus has come out, because they just got it. Upon examination, you'll often discover that their virus definition file is months out-of-date. As you can see, the software part of the system will break down if the definition files aren't updated on a regular basis.

Implementing Access Control

The process of implementing *access control* is critical. Access control defines how users and systems communicate and in what manner. In other words, it limits—or controls—access to system resources, including data, and thus protects information from unauthorized access. Three basic models are used to explain access control.

The Mandatory Access Control Method

The *Mandatory Access Control (MAC)* model is a static model that uses a predefined set of access privileges for files on the system. The system administrators establish these parameters and associate them with an account, files, or resources.

The MAC model can be very restrictive. In a MAC model, administrators establish access. Users can't share resources dynamically unless the static relationship already exists.

The acronym MAC appears in numerous computer-related contexts. One of the most common uses is to represent the Media Access layer in networking. Be careful not to confuse MAC addressing as it relates to network cards with Mandatory Access Control.

MAC uses *labels* to identify the level of sensitivity that applies to objects. When a user attempts to access an object, the label is examined to see if the access should take place or be denied. One key element to remember is that when mandatory control is applied, labels are required and must exist for every object.

The Discretionary Access Control Method

The *Discretionary Access Control (DAC)* model allows the owner of a resource to establish privileges to the information they own. The difference between DAC and MAC is that labels are not mandatory but can be applied as needed.

The DAC model allows a user to share a file or use a file that someone else has shared. It establishes an access control list (ACL) that identifies the users who have authorization to access that information. This allows the owner to grant or revoke access to individuals or

groups of individuals based on the situation. This model is dynamic in nature and allows information to be shared easily between users.

The Role-Based Access Control Method

The *Role-Based Access Control (RBAC)* model allows a user to act in a certain predetermined manner based on the role the user holds in the organization. The roles almost always shadow the organizational structure.

Users can be assigned roles systemwide and can then perform certain functions or duties based on the roles they're assigned. An example might be a role called *salesperson*. The user assigned the salesperson role can access only the information established for that role. Users may be able to access this information from any station in the network, based strictly on their role. A sales manager may have a different role that allows access to all of the individual salespersons' information.

The RBAC model is common in network administrative roles.

Understanding Authentication

Authentication proves that a user or system is actually who they say they are. This is one of the most critical parts of a security system. It's part of a process that is also referred to as *identification and authentication (I&A)*. The identification process starts when a user ID or logon name is typed into a sign-on screen. Authentication is accomplished by challenging the claim about who is accessing the resource. Without authentication, anybody can claim to be anybody.

Authentication systems or methods are based on one or more of these three factors:

- Something you know, such as a password or PIN
- Something you have, such as a smart card or an identification device
- Something physically unique to you, such as your fingerprints or retinal pattern

Systems authenticate each other using similar methods. Frequently, systems pass private information between each other to establish identity. Once authentication has occurred, the two systems can communicate in the manner specified in the design.

Several common methods are used for authentication. Each offers something to security and should be considered when you're evaluating authentication schemes or methods.

Biometrics

Biometric readers use physical characteristics to identify the user. Such devices are becoming more common in the business environment. Biometric readers include hand scanners, retinal scanners, and soon, possibly, DNA scanners. To gain access to resources, you must pass a physical screening process. In the case of a hand scanner, the screening may include fingerprints, scars, and markings on your hand. Retinal scanners compare your eye's retinal pattern to a stored retinal pattern to verify your identity. DNA scanners will examine a unique portion of your DNA structure to verify that you are who you say you are.

Certificates

Certificates are another common form of authentication. A server or *certificate authority (CA)* can issue a certificate that will be accepted by the challenging system. Certificates can be either physical access devices, such as smart cards, or electronic certificates that are used as part of the logon process. A *Certificate Practice Statement (CPS)* outlines the rules used for issuing and managing certificates. A *Certificate Revocation List (CRL)* lists the revocations that must be addressed (often due to expiration) in order to stay current.

This chapter provides only an overview of certificates. Certificates, along with Public Key Infrastructure (PKI) and related topics, are discussed in detail in Chapter 7, "Cryptography Basics, Methods, and Standards."

A simple way to think of certificates is to think of hall passes at school. Figure 1.3 illustrates a certificate being handed from the server to the client after authentication has been established. If you have a hall pass, you can wander the halls of your school. If your pass is invalid, the hallway monitor can send you to the principal's office. Similarly, if you have a certificate, then you can prove to the system that you are who you say you are and are authenticated to work with the resources.

FIGURE 1.3 A certificate being issued after identification has been verified

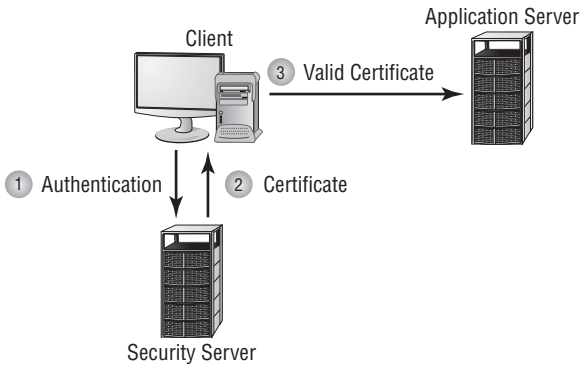

Challenge Handshake Authentication Protocol

Challenge Handshake Authentication Protocol (CHAP) challenges a system to verify identity. CHAP doesn't use a user ID/password mechanism. Instead, the initiator sends a logon request from the client to the server. The server sends a challenge back to the client. The challenge is encrypted and then sent back to the server. The server compares the value from the client and, if the information matches, grants authorization. If the response fails, the session fails, and the request phase starts over. Figure 1.4 illustrates the CHAP procedure. This handshake method involves a number of steps and is usually automatic between systems.

FIGURE 1.4 CHAP authentication

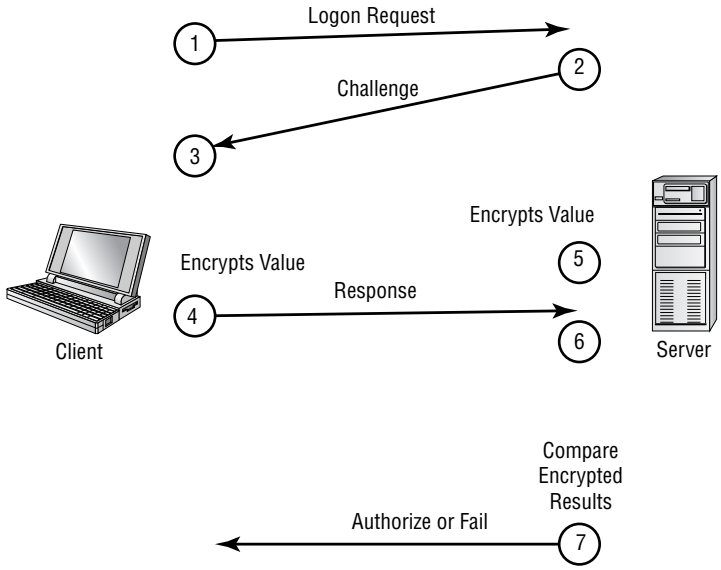

Kerberos

Kerberos is an authentication protocol named after the mythical three-headed dog that stood at the gates of Hades. Originally designed by MIT, Kerberos is becoming very popular as an authentication method. It allows for a single sign-on to a distributed network.

Kerberos authentication uses a *Key Distribution Center (KDC)* to orchestrate the process. The KDC authenticates the *principle* (which can be a user, a program, or a system) and provides it with a ticket. After this ticket is issued, it can be used to authenticate against other principles. This occurs automatically when a request or service is performed by another principle.

Kerberos is quickly becoming a common standard in network environments. Its only significant weakness is that the KDC can be a single point of failure. If the KDC goes down, the authentication process will stop. Figure 1.5 shows the Kerberos authentication process and the ticket being presented to systems that are authorized by the KDC.

Multi-Factor Authentication

When two or more access methods are included as part of the authentication process, you're implementing a *multi-factor* system. A system that uses smart cards and passwords is referred to as a *two-factor authentication* system. Two-factor authentication is shown in Figure 1.6. This example requires both a smart card and a logon password process.

FIGURE 1.5 Kerberos authentication process

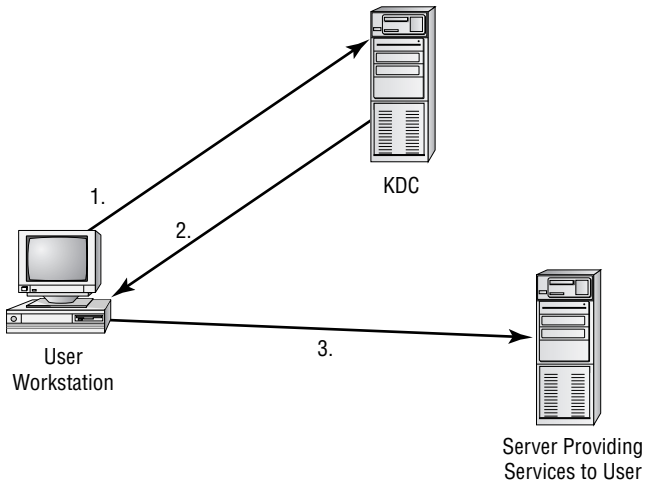

1. User requests access to service running on a different server.
2. KDC authenticates user and sends a ticket to be used between the user and the service on the server.
3. User's workstation sends a ticket to the service.

FIGURE 1.6 Two-factor authentication

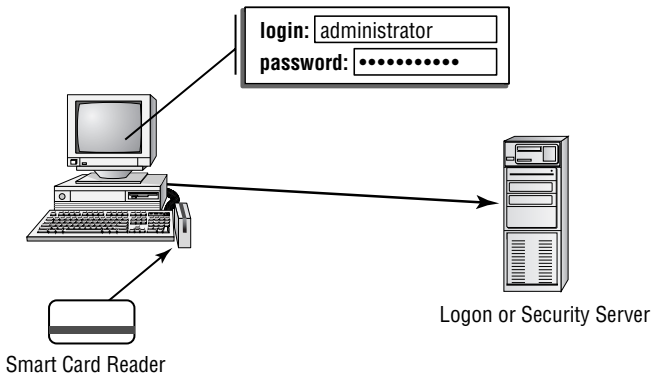

Both factors must be valid:
•User ID and Password
•Smart Card

Mutual Authentication

Whenever two or more parties authenticate each other, this is known as *mutual authentication*. A client may authenticate to a server, and a server authenticate to a client when there is a need to establish a secure session between the two and employ encryption. Mutual authentication ensures that the client is not unwittingly connecting and giving its credentials to a rogue server; which can then turn around and steal the data from the real server.

Commonly, mutual authentication will be implemented when the data to be sent during the session is of a critical nature—such as financial or medical records.

Password Authentication Protocol

Password Authentication Protocol (PAP) offers no true security, but it's one of the simplest forms of authentication. The username and password values are both sent to the server as clear text and checked for a match. If they match, the user is granted access; if they don't match, the user is denied access. In most modern implementations, PAP is shunned in favor of other, more secure authentication methods.

Security Tokens

Security tokens are similar to certificates. They contain the rights and access privileges of the token bearer as part of the token. Think of a token as a small piece of data that holds a sliver of information about the user.

Many operating systems generate a token that is applied to every action taken on the computer system. If your token doesn't grant you access to certain information, then either that information won't be displayed or your access will be denied. The authentication system creates a token every time a user connects or a session begins. At the completion of a session, the token is destroyed. Figure 1.7 shows the security token process.

FIGURE 1.7 Security token authentication

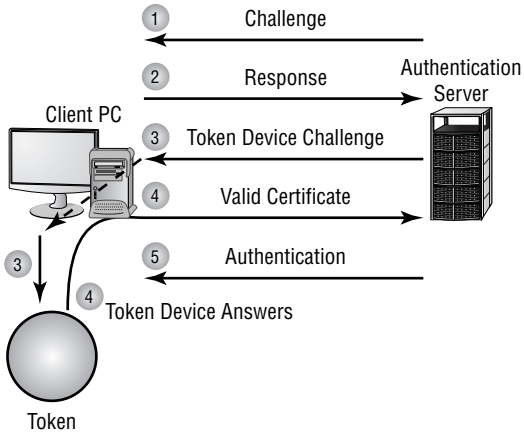

Smart Cards

A *smart card* is a type of badge or card that gives you access to resources, including buildings, parking lots, and computers. It contains information about your identity and access privileges. Each area or computer has a card scanner or a reader in which you insert your card. Smart Cards often also require the use of a small password called a PIN (personal identification number); which further secures the smart card if lost by the true card holder, so that it cannot be used by someone else to gain access to data and resources.

Figure 1.8 depicts a user inserting a smart card into a reader to verify identity. The reader is connected to the workstation and validates against the security system. This increases the security of the authentication process because you must be in physical possession of the smart card to use the resources. Of course, if the card is lost or stolen, the person who finds the card can access the resources allowed by the smart card.

FIGURE 1.8 The smart card authentication process

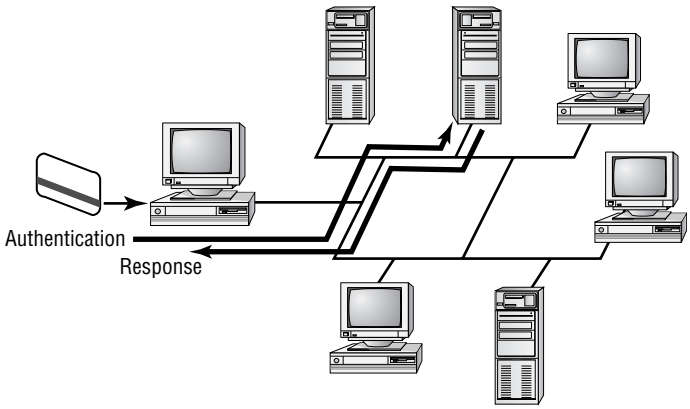

Username/Password

A *username* and *password* are unique identifiers for a logon process. When users sit down in front of a computer system, the first thing a security system requires is that they establish who they are. Identification is typically confirmed through a logon process. Most operating systems use a user ID and password to accomplish this. These values can be sent across the connection as plain text or can be encrypted.

The logon process identifies to the operating system, and possibly the network, that you are who you say you are. Figure 1.9 illustrates this logon and password process. Notice that the operating system compares this information to the stored information from the security processor and either accepts or denies the logon attempt. The operating system might establish privileges or permissions based on stored data about that particular ID.

FIGURE 1.9 A logon process occurring on a workstation

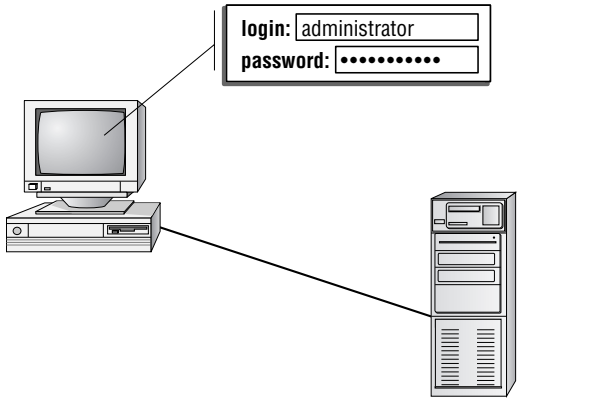

Authentication Issues to Consider

You can set up many different parameters and standards to force the people in your organization to conform. In establishing these parameters, it's important that you consider the capabilities of the people who will be working with these policies. If you're working in an environment where people aren't computer savvy, you may spend a lot of time helping them remember and recover passwords. Many organizations have had to reevaluate their security guidelines after they've already invested great time and expense to implement high-security systems.

Setting authentication security, especially in supporting users, can become a high-maintenance activity for network administrators. On one hand, you want people to be able to authenticate themselves easily; on the other hand, you want to establish security that protects your company's resources.

Be wary of popular names or current trends that make certain passwords predictable. For example, during the first release of *Star Wars*, two of the most popular passwords used on college campuses were C3PO and R2D2. This created a security problem for campus computer centers.

Whenever an issue arises between identification and authentication, *identify proofing* is often called upon. As mentioned earlier in this chapter, the identification process starts when a user ID or logon name is typed into a sign-on screen. Authentication is accomplished by challenging the claim about who is accessing the resource.

Identification proofing is invoked when a person claims they are the user, but cannot be authenticated—such as when they lose their password. Since they can't provide the password, they are typically asked to provide another value—such as mother's maiden name—to prove their identity.

Real World Scenario

Multi-Factor Authentication and Security

The owner of your company is becoming increasingly concerned about computer security and the laxness of users. She reports that users are regularly leaving the office at the end of the day without signing out of their accounts. The company is attempting to win a contract that involves working with the government and that will require additional security measures. What would you suggest to the owner?

The best suggestion is to consider implementing a multi-factor authentication system. This system could consist of a smart card and a logon/password process. Most smart card readers can be configured to require that the card remain inserted in the reader while the user is logged on. If the smart card is removed, say at the end of the day, the workstation will automatically log the user out. By requiring a logon/password process, you can still provide security if the smart card is stolen.

This solution provides reasonable security, and it doesn't significantly increase security costs. The government will probably require additional access control, such as perimeter alarms and physical access control to sensitive areas. However, these measures won't force users to log out when they leave their workstations.

An inherent problem with many identify proofing implementations is that they ask questions which someone other than the user could easily guess or learn the value of (what color are your eyes). To increase the difficulty of someone fraudulently proofing, you should only use questions that are more difficult to guess, or implement biometrics such as voice identification. Under no circumstance should the person proofing be allowed access immediately—instead their access information should be sent to their email account of record.

Distinguishing between Security Topologies

The *security topology* of your network defines the network design and implementation from a security perspective. Unlike a network topology, here we're concerned with access methods, security, and technologies used. Security topology covers four primary areas of concern:

- Design goals
- Security zones
- Technologies
- Business requirements

Setting Design Goals

When setting design goals for a security topology, you must deal with issues of confidentiality, integrity, availability, and accountability, all four of which are discussed continually throughout this book as they apply to various topics. Addressing these four issues as an initial part of your network design will help you ensure tighter security. You'll often see confidentiality, integrity, and availability referred to as the *CIA* of network security, but the accountability component is equally important—design goals must identify who is responsible for the various aspects of computer security. The next few sections introduce these four security components.

Confidentiality

Meeting the goal of *confidentiality* is to prevent or minimize unauthorized access to and disclosure of data and information. In many instances, laws and regulations require specific information confidentiality. For example, Social Security records, payroll and employee records, medical records, and corporate information are high-value assets. This information could create liability issues or embarrassment if it fell into the wrong hands. Over the last few years, there have been a number of cases in which bank account and credit card numbers were published on the Internet. The costs of these types of breaches of confidentiality far exceed the actual losses from the misuse of this information.

Confidentiality entails ensuring that data expected to remain private is seen only by those who should see it. Confidentiality is implemented through authentication and access controls.

If you address confidentiality issues early in the design phase, the steps that must be taken to minimize this exposure will become clear.

Integrity

Meeting the goal of *integrity* involves making sure that the data being worked with is the correct data. Information integrity is critical to a secure topology. Organizations work with and make decisions using the data they have available. If this information isn't accurate or is tampered with by an unauthorized person, the consequences can be devastating.

Take the case of a school district that lost all the payroll and employment records for the employees in the district. When the problem was discovered, the school district had no choice but to send out applications and forms to all the employees, asking them how long they had worked in the school district and how much they were paid. Integrity was jeopardized because the data was vulnerable and then lost.

You can think of integrity as the level of confidence you have that the data is what it's supposed to be—untampered with and unchanged. *Authentic, complete,* and *trustworthy* are often used to describe integrity in terms of data.

Availability

To meet the goal of *availability*, you must protect data and prevent its loss. Data that can't be accessed is of little value. If a mishap or attack brings down a key server or database, that information won't be available to the people who need it. This can cause havoc in an organization. Your job is to provide maximum availability to your users while ensuring integrity and confidentiality. The hardest part of this process is determining the balance you must maintain between these three aspects to provide acceptable security for the organization's information and resources.

Real World Scenario

Compute Availability

Availability is often expressed in terms of *uptime*. High availability strives for 99.9999% uptime over the course of the year (24 hours a day, 7 days a week, 365 days a year). For this exercise, compute how long data wouldn't be available over the course of the year with the following availability percentages. For example, with 98% uptime, there is a 2% downtime of 525,600 minutes in a year. That means the data would be down for 10,512 minutes, or 7⅓ days! Try your math on the following:

1. 99%
2. 99.9%
3. 99.99%
4. 99.999%
5. 99.9999%

The increments may seem small, but over the course of a year, they represent a significant difference in the amount of time data is and isn't available. Answers: (1.) 5,256 minutes, which is over 87 hours and 3.5 days; (2.) 525 minutes, or a little less than 9 hours; (3.) 52.56 minutes; (4.) 5.25 minutes; (5.) About half a minute.

The key to availability is that the data must be available when it's needed and accessible by those who need it.

Accountability

The final and often overlooked goal of design concerns *accountability*. Many of the resources used by an organization are shared between departments and individuals. If

an error or incident occurs, who is responsible for fixing it? Who determines whether information is correct?

It's a good idea to be clear about who owns the data or is responsible for making sure that it's accurate. You should also be able to track and monitor data changes to detect and repair the data in the event of loss or damage. Most systems will track and store logs on system activities and data manipulation, and they will also provide reports on problems.

Real World Scenario

Accountability Is More than a Catchphrase

Accountability, like common sense, applies to every aspect of information technology. Several years ago, a company that relied on data that could never be re-created wrote shell scripts to do backups early in the morning when the hosts were less busy. Operators at those machines were told to insert a tape in the drive around midnight and check back at 3:00 a.m. to make certain that a piece of paper had been printed on the printer, signaling the end of the job. If the paper was there, they were to remove the tapes and put them in storage; if the paper was not there, they were to call for support.

The inevitable hard drive crash occurred on one of the hosts one morning, and an IT "specialist" was dispatched to swap it out. The technician changed the hard drive and then asked for the most recent backup tape. To his dismay, the data on the tape was two years old. The machine crash occurred before the backup operation ran, he reasoned, but the odds of rotating two years' worth of tapes was pretty amazing. Undaunted, he asked for the tape from the day before, and found that the data on it was also two years old.

Beginning to sweat, he found the late shift operator for that host and asked her if she was making backups. She assured him that she was and that she was rotating the tapes and putting them away as soon as the paper printed out. Questioning her further on how the data could be so old, she said she could verify her story because she also kept the pieces of paper that appeared on the printer each day. She brought out the stack and handed them to him. They all reported the same thing—*tape in drive is write protected*.

Where did the accountability lie in this true story? The operator was faithfully following the procedures given to her. She thought the fact that the tape was protected represented a good thing. It turned out that all the hosts had been printing the same message, and none of them had been backed up for a long while.

The problem lay not with the operator, but with the training she was given. Had she been shown what correct and incorrect backup completion reports looked like, the data would never have been lost.

Creating Security Zones

Over time, networks can become complex beasts. What may have started as a handful of computers sharing resources can quickly grow to something resembling an electrician's nightmare. The networks may even appear to have lives of their own. It's common for a network to have connections among departments, companies, countries, and public access using private communication paths and through the Internet.

Not everyone in a network needs access to all the assets in the network. The term *security zone* describes design methods that isolate systems from other systems or networks. You can isolate networks from each other using hardware and software. A router is a good example of a hardware solution: You can configure some machines on the network to be in a certain address range and others to be in a different address range. This separation makes the two networks invisible to each other unless a router connects them. Some of the newer data switches also allow you to partition networks into smaller networks or private zones.

When discussing security zones in a network, it's helpful to think of them as rooms. You may have some rooms in your house or office that anyone can enter. For other rooms, access is limited to specific individuals for specific purposes. Establishing security zones is a similar process in a network: Security zones allow you to isolate systems from unauthorized users. Here are the four most common security zones you'll encounter:

- Internet
- Intranet
- Extranet
- Demilitarized zone (DMZ)

The next few sections identify the topologies used to create security zones to provide security. The Internet has become a boon to individuals and to businesses, but it creates a challenge for security. By implementing intranets, extranets, and DMZs, you can create a reasonably secure environment for your organization.

The Internet

The *Internet* is a global network that connects computers and individual networks together. It can be used by anybody who has access to an Internet portal or an Internet service provider (ISP). In this environment, you should have a low level of trust in the people who use the Internet. You must always assume that the people visiting your website may have bad intentions; they may want to buy your product, hire your firm, or bring your servers to a screaming halt. Externally, you have no way of knowing until you monitor their actions. Because the Internet involves such a high level of anonymity, you must always safeguard your data with the utmost precautions.

Figure 1.10 illustrates an Internet network and its connections.

Sometimes the data leaving a network can be as much a sign of trouble as the data entering it. Examining data leaving the network for signs of malicious traffic is a fairly new field of computer security and is known as *extrusion*.

FIGURE 1.10 A typical LAN connection to the Internet

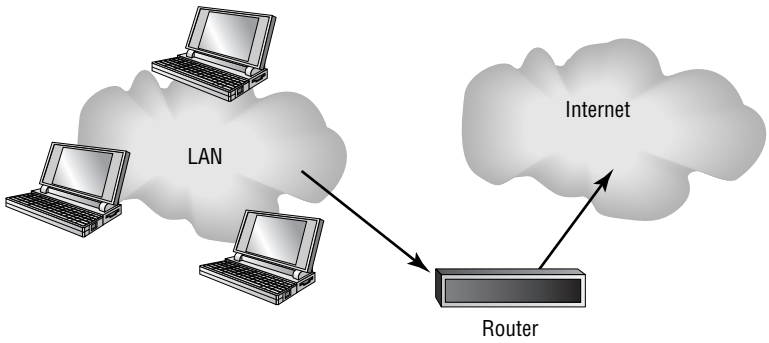

Intranets

Intranets are private networks implemented and maintained by an individual company or organization. You can think of an intranet as an Internet that doesn't leave your company; it's internal to the company, and access is limited to systems within the intranet. Intranets use the same technologies used by the Internet. They can be connected to the Internet but can't be accessed by users who aren't authorized to be part of them; the anonymous user of the Internet is instead an authorized user of the intranet. Access to the intranet is granted to trusted users inside the corporate network or to users in remote locations.

Figure 1.11 displays an intranet network.

FIGURE 1.11 An intranet network

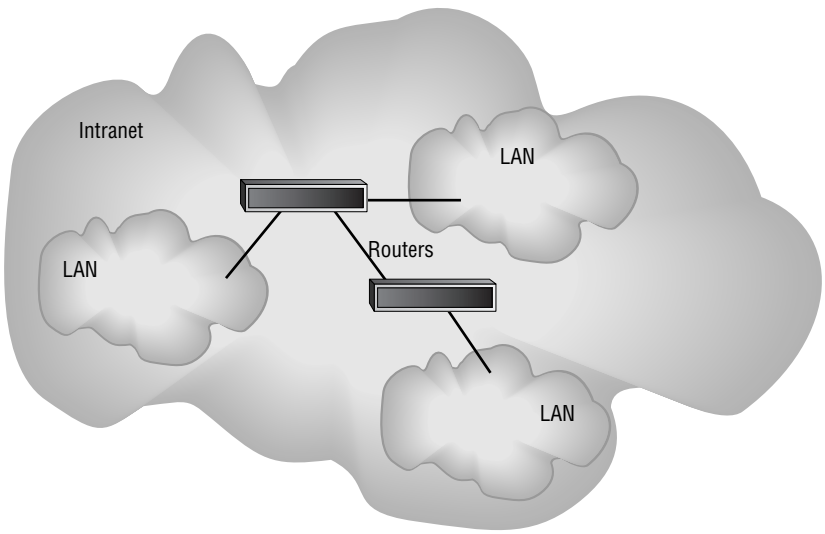

Extranets

Extranets extend intranets to include outside connections to partners. The partners can be vendors, suppliers, or similar parties who need access to your data for legitimate reasons. An extranet allows you to connect to a partner via a private network or a connection using a secure communications channel across the Internet. Extranet connections involve connections between trustworthy organizations.

An extranet is illustrated in Figure 1.12. Note that this network provides a connection between the two organizations. The connection may be through the Internet; if so, these networks would use a tunneling protocol to accomplish a secure connection.

FIGURE 1.12 A typical extranet between two organizations

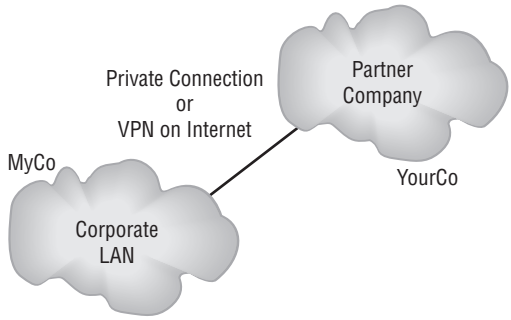

Demilitarized Zone (DMZ)

A *demilitarized zone (DMZ)* is an area where you can place a public server for access by people you might not trust otherwise. By isolating a server in a DMZ, you can hide or remove access to other areas of your network. You can still access the server using your network, but others aren't able to access further network resources. This can be accomplished using firewalls to isolate your network.

When establishing a DMZ, you assume that the person accessing the resource isn't necessarily someone you would trust with other information. Figure 1.13 shows a server placed in a DMZ. Notice that the rest of the network isn't visible to external users. This lowers the threat of intrusion in the internal network.

Anytime you want to separate public information from private information, a DMZ is an acceptable option.

The easiest way to create a DMZ is to use a firewall that can transmit in three directions: to the internal network, to the external world (Internet), and to the public information you're sharing (the DMZ). From there, you can decide what traffic goes where; for example, HTTP traffic would be sent to the DMZ, and e-mail would go to the internal network.

FIGURE 1.13 A typical DMZ

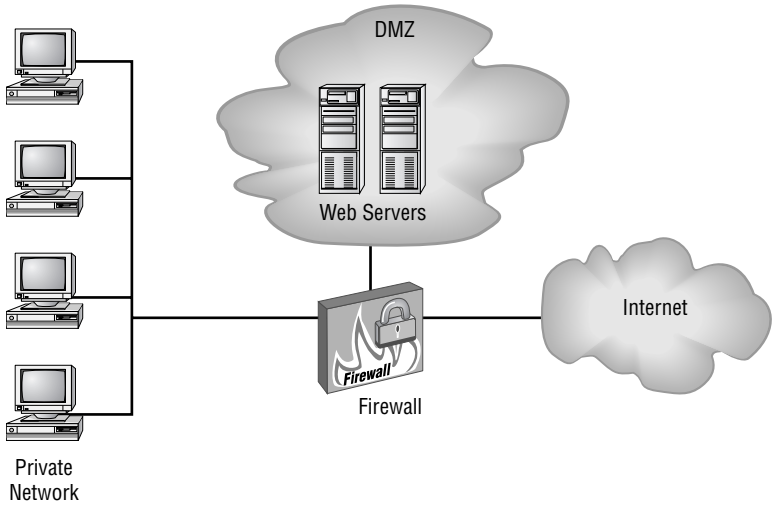

Designing Security Zones

Security zone design is an important aspect of computer security. You can use many different approaches to accomplish a good solid design. Some of the design trade-offs involve risk and money. You can create layers of security to protect systems from less-secure connections, and you can use Network Address Translation (NAT) (discussed later) to hide resources. New methods and tools to design secure networks are being introduced on a regular basis. It's important to remember that after you have a good security design, you should revisit it on a regular basis based on what you learn about your security risks.

Working with Newer Technologies

One of the nice things about technology is that it's always changing. One of the bad things about technology is that it's always changing. Several relatively new technologies have become available to help you create a less-vulnerable system. The four technologies this section will focus on are virtualization, virtual local area networks (VLANs), Network Address Translation, and tunneling. These technologies allow you to improve security in your network at little additional cost.

Virtualization Technology

Virtualization is easily the technology du jour, with VMWare, one of the largest vendors of such technology, counting 100% of the Fortune 100 as part of their customer base. In addition to proprietary solutions, there are also open source solutions as well, with Xen being the most well-known example.

Virtualization technology allows you to take any single physical device and hide its characteristics from users—in essence allowing you to run multiple items on one device and make them appear as if they are standalone entities. For example, workstations can only run one operating system at a time. Using virtualization, it is possible for a workstation running Windows XP to also be running Fedora, Red Hat, Windows Server 2003, and any number of other operating systems within virtual windows. The developer working on code can move between windows, cutting and pasting if they choose, and do all they need to do on one machine without needing to run four different workstations. Thanks to virtualization, the workstation can run multiple operating systems, multiple versions of the same operating system, multiple applications, and so on.

Just as a workstation can be virtualized, so, too, can a server. A single server can host multiple logical machines. By using one server to do the function of many, cost savings can be immediately gained in terms of hardware, utility, infrastructure, and so on.

As wonderful as virtualization is, from a security standpoint, it can present challenges. A user accessing the system could have access to everything on the system (not just within their logical machine) if they could override the physical layer protection. As of this writing, the threat of that occurring has been far more rumored than performed, but with virtualization growing in popularity, it is a safe bet that virtual machines will become a popular target of miscreants in coming years.

Virtual Local Area Networks

A *virtual local area network (VLAN)* allows you to create groups of users and systems and segment them on the network. This segmentation lets you hide segments of the network from other segments and thereby control access. You can also set up VLANs to control the paths that data takes to get from one point to another. A VLAN is a good way to contain network traffic to a certain area in a network.

Think of a VLAN as a network of hosts that act as if they're connected by a physical wire even though there is no such wire between them.

On a LAN, hosts can communicate with each other through broadcasts, and no forwarding devices, such as routers, are needed. As the LAN grows, so too does the number of broadcasts. Shrinking the size of the LAN by segmenting it into smaller groups (VLANs) reduces the size of the broadcast domains. The advantages of doing this include reducing the scope of the broadcasts, improving performance and manageability, and decreasing dependence on the physical topology. From the standpoint of this exam, however, the key benefit is that VLANs can increase security by allowing users with similar data sensitivity levels to be segmented together.

Figure 1.14 illustrates the creation of three VLANs in a single network.

Network Address Translation

Network Address Translation (NAT) creates a unique opportunity to assist in the security of a network. Originally, NAT extended the number of usable Internet addresses. Now it allows an

organization to present a single address to the Internet for all computer connections. The NAT server provides IP addresses to the hosts or systems in the network and tracks inbound and outbound traffic.

A company that uses NAT presents a single connection to the network. This connection may be through a router or a NAT server. The only information that an intruder will be able to get is that the connection has a single address.

NAT effectively hides your network from the world, making it much harder to determine what systems exist on the other side of the router. The NAT server effectively operates as a firewall for the network. Most new routers support NAT; it provides a simple, inexpensive firewall for small networks.

FIGURE 1.14 A typical segmented VLAN

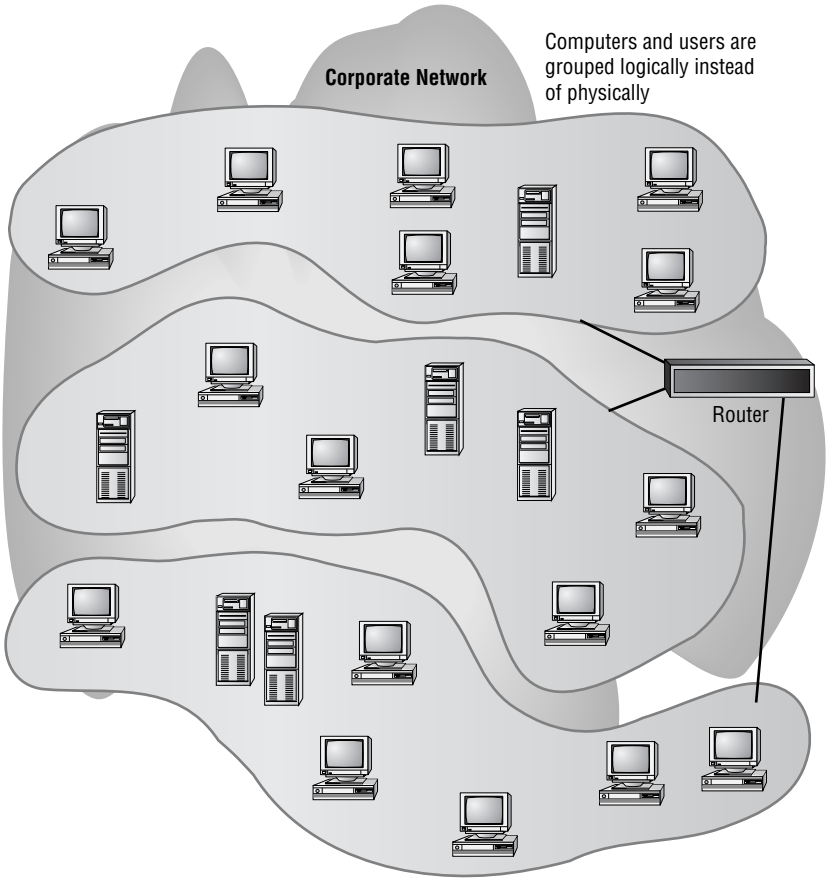

It's important to understand that NAT acts as a proxy between the local area network (which can be using private IP addresses) and the Internet. Not only can NAT save IP addresses, but it can also act as a firewall.

Most NAT implementations assign internal hosts private IP address numbers and use public addresses only for the NAT to translate to and communicate with the outside world. The private address ranges are as follows:

10.0.0.0–10.255.255.255

172.16.0.0–172.31.255.255

192.168.0.0–192.168.255.255

Figure 1.15 shows a router providing NAT services to a network. The router presents a single address for all external connections on the Internet.

FIGURE 1.15 A typical Internet connection to a local network

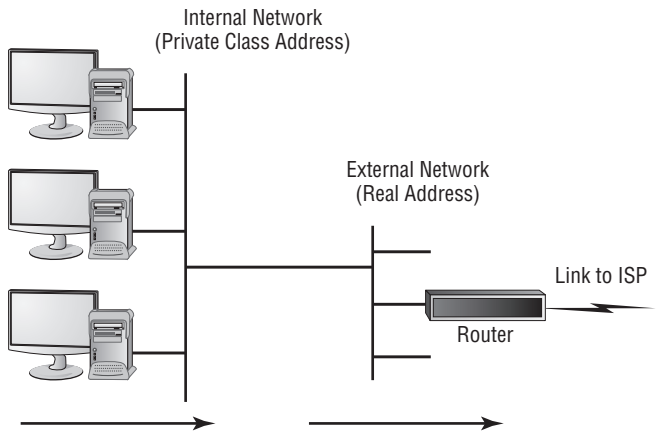

In addition to NAT, Port Address Translation (PAT) is possible. Whereas NAT can use multiple public IP addresses, PAT uses a single one and shares the port with the network. Because it is only using a single port, PAT is much more limited and typically only used on small and home-based networks. Microsoft's Internet Connection Sharing is an example of a PAT implementation.

IP addressing is a subject on the Network+ exam, as opposed to Security+, but CompTIA still expects you to know the basics. In addition to understanding the concept behind NAT, you should know that subnetting is how networks are divided. RFCs 1466 and 1918 detail subnetting and can be found at `http://www.faqs.org/rfcs/`.

Tunneling

Tunneling refers to creating a virtual dedicated connection between two systems or networks. You create the tunnel between the two ends by encapsulating the data in a mutually agreed-upon protocol for transmission. In most tunnels, the data passed through the tunnel appears at the other side as part of the network.

Tunneling protocols usually include data security as well as encryption. Several popular standards have emerged for tunneling, with the most popular being the Layer 2 Tunneling Protocol (L2TP).

Tunneling sends private data across a public network by placing (encapsulating) that data into other packets. Most tunnels are virtual private networks (VPNs).

Figure 1.16 shows a connection being made between two networks across the Internet. To each end of the network, this appears to be a single connection.

FIGURE 1.16 A typical tunnel

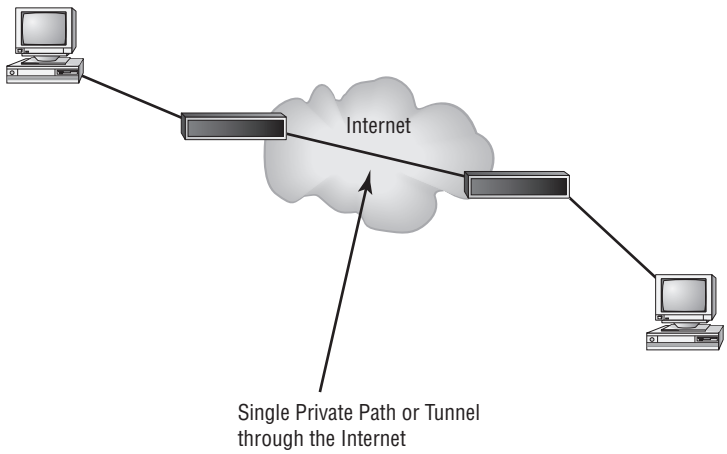

Addressing Business Concerns

An organization or business is well served if it makes a conscious examination of its security situation. This examination includes identifying assets, doing a comprehensive risk assessment, identifying threats, and evaluating vulnerabilities. These four components will help the business principals understand what they're up against and how to cost-effectively address these issues.

The following sections explain the various business requirements you need to address when designing a security topology. The failure to consider any one of these aspects can cause the entire design to be flawed and ineffective.

Real World Scenario

Creating a Corporate Connection to a Business Partner

Your company has just signed an agreement with a large wholesaler to sell your products. The wholesaler has an extensive network that utilizes a great deal of technology, which will benefit you and improve your profitability. You must design a network security topology that will allow you both access to some of each other's systems and information while protecting the confidentiality of your own critical records and information. How might you accomplish this?

A good implementation would connect your network to theirs using a VPN across the Internet. You could use a secure tunneling protocol to ensure that unauthorized parties wouldn't be able to sniff or access information streams between the companies. This approach would create an extranet environment for you and your new business partner.

The challenge lies in creating secure areas in your network that the wholesaler can't access. You can accomplish this by establishing VLANs in your internal network that aren't visible to the extranet. VLANs and network segmentation can be implemented using routers, firewalls, and switches.

Identifying Assets

Every business or organization has valuable assets and resources. These assets must be accounted for, both physically and functionally. *Asset identification* is the process in which a company attempts to place a value on the information and systems it has in place. In some cases, the process may be as simple as counting systems and software licenses. These types of physical asset evaluations are part of the normal accounting procedures a business must perform routinely.

The more difficult part of an asset-identification process is attempting to assign values to information. In some cases, you may only be able to determine what would happen if the information were to become unavailable or lost. If absence of this information would effectively shut down the business, the information is priceless. If you have this type of information, determining which methods and approaches you should take to safeguard it becomes easier.

You wouldn't necessarily assign the same value to the formula for Coca-Cola that you'd assign to your mother's chicken and rice recipe. The Coke formula would be worth a fortune to a person who stole it; they could sell it to competitors and retire. Your mother's recipe would make a nice dinner, but it wouldn't be valuable from a financial perspective.

Real World Scenario

Assign a Value to Data Assets

Think of yourself as a collection of data elements. Some of the data about you, such as your last name, isn't of great value since it's known by almost everyone you come into contact with. Other data, such as your Social Security number, should be closely guarded and is worth more than your name because you stand to lose more if it falls into the wrong hands. See if you can assign a value to each of these items and rank which is worth the most according to what would be most harmful in the hands of a miscreant:

1. Full name
2. Birth date
3. Telephone number
4. Passport number

If this data were spread across a number of databases on a computer system, you would naturally want to assign higher value to the databases containing the most sensitive data and then take more drastic steps to protect them than you would for those containing generic information.

Assessing Risk

There are several ways to perform a *risk assessment* or *risk analysis*. They range from highly scientific formula-based methods to a conversation with the owner. In general, you should attempt to identify the costs of replacing stolen data or systems, the costs of down-time, and virtually any risk factor you can imagine.

You can move to risk assessment only after completing the asset identification. After you know that databases containing information from freely available sources (such as the U.S. Census Bureau) can always be re-created if need be and shouldn't be viewed in the same light as those containing business-specific data, you can start computing costs.

After you've determined the costs, you can then evaluate the likelihood that certain types of events will occur and the most likely outcome if they do occur. If you work in New York City, what is the likelihood of damage to your business from an earthquake? Will your risk assessment place the high probability of an earthquake on your list of primary concerns?

Identifying Threats

Implementing a security policy requires that you evaluate the risks of both internal and external *threats* to the data and network. It does little good to implement a high-security

environment to protect your company from the outside if the threat is mostly internal. If a member of your team brings a disk containing a virus into the office and loads it onto a computer, the virus may spread throughout the entire network and effectively be immune to your external security measures. This is a common problem in schools, libraries, and environments where people regularly use shared resources. If a library offers computers for public use and those computers are in a network, a virus could infect all of the systems throughout the network. External security measures won't prevent potential damage or data loss.

Internal threats also include employee fraud, abuse or alteration of data, and theft of property. Both policies and systems must be put into place to detect and mitigate these possibilities. Investigating and making recommendations to management on procedural changes and policies is a key role for computer security professionals. Figure 1.17 depicts some examples of internal and external threats.

FIGURE 1.17 Internal and external threats to an organization

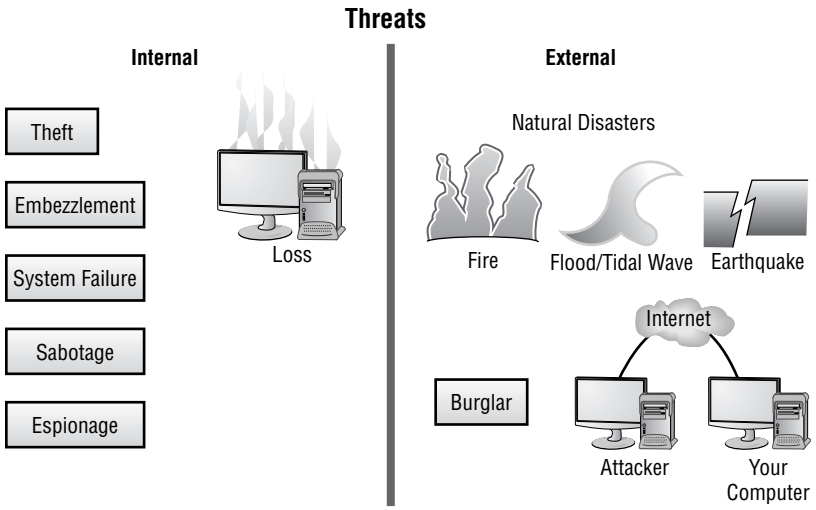

Internal Threats

Most well-publicized *internal threats* involve financial abuses. Some of these abuses are outright fraud or theft. These types of threats, especially in a computer-intensive environment, can be difficult to detect and investigate. They are typically ongoing and involve small transactions over long periods. A recent incident of fraud that occurred in a large software manufacturer involved an accounting professional who generated bogus checks in payment for work that never occurred. Over a few months, this employee was able to make over $100,000 in fraudulent payments to companies that she or relatives had created. It took considerable investigation by computer and financial auditors to determine how this theft occurred. From a computer security perspective, this was an internal threat that was the result of failures

in financial, operational, and computer security controls. These types of incidents probably occur more frequently than anyone wants to admit, and many times more often than anyone becomes aware of.

Another incident involved an employee who was using corporate computer resources to operate a financial accounting service. This employee had been running this business for several years. When the company found out, it immediately fired the employee and confiscated his records. During the investigation, the process used to collect evidence inadvertently tainted it. The chain of custody in this case was broken. When the employee went to court over this situation, his attorney was able to have the evidence thrown out of court. Even though the employee was clearly guilty, the judge dismissed the case due to a lack of admissible evidence. The employee then sued the company for wrongful discharge, harassment, and several other charges. He won those suits, and he got his old job back. In this instance, the internal policies and systems put into place to detect, investigate, and correct the problem broke down. It cost the company a huge amount of money and allowed a known embezzler back in.

We'll discuss chains of custody, incident response, and the proper way to conduct investigations in Chapter 8. For now, it's important to know that finding and dealing with internal threats is a key aspect of the computer security job.

External Threats

Many of the internal threats that a company must deal with involve procedures and methods that are standard across industries. *External threats*, on the other hand, are increasing at an alarming rate. Several years ago, most computer incidents were caused by groups of kids or hobbyists who were primarily in it for fun. Most of the time, these incidents were not intentionally malicious in nature. A few of them did involve alteration or destruction of data and records.

Today, many companies take orders online, process payments, track shipments, manage inventory, use online databases, and administer other key information using complicated systems. These systems are connected to other systems that contain private corporate records, trade secrets, strategic plans, and many other types of valuable information.

Unfortunately, when these systems are compromised, an entire business or industry can be compromised. Incidents have occurred where security breaches remained open for years, and the companies involved had no knowledge that a compromise ever took place. One of a professional criminal's greatest joys is creating and exploiting this type of security breach.

Early methods of cracking systems were primitive and labor intensive. Today, software packages exist that find targets automatically and then systematically attack the targets to find their vulnerabilities. Many of these tools use graphical user interfaces that require little technical expertise on the part of the would-be hacker. Many computer systems are being repeatedly and methodically attacked by the curious or by criminals attempting to commit a crime.

The job of a computer security professional in this situation is to detect the attack, find ways to counter it, and assist law-enforcement personnel in investigating the activity. This type of work is interesting and involves many of the skills you'll learn in this book.

Understanding Vulnerabilities

A computer security specialist's main area of concern will probably revolve around the security capabilities of the software and systems used in the business. Until recently, many operating system manufacturers only paid lip service to security. One popular operating system used a logon and password scheme for security. When the logon prompt occurred, all you had to do was click the Cancel button and the system would provide most of the network capabilities and local access to all resources. If the screensaver was password protected, you could either enter the password to unlock the system or reboot the computer to have the system be unsecure. This was worse than having no security. Many users thought they had a secure computer system, but they didn't—and many thefts of data by coworkers occurred as a result.

The *Transmission Control Protocol/Internet Protocol (TCP/IP)* network protocol used by most corporate networks was designed to allow communications in a trustful environment. This protocol was primarily experimental and was used by schools and governmental agencies for research. Although it's robust in its error handling, by its nature it's unsecured. Many modern network attacks occur through the TCP/IP.

Operating systems and applications programs have long been vulnerable to external and internal attacks. Software companies want to sell software that is easy to use, graphically driven, and easily configured. Users want the same thing. Unfortunately, this creates additional security problems in many networks.

One of the most popular products in use today allows e-mail and attachments to begin executing programs or instructions embedded in a message. This functionality allows e-mail messages to have fancy formatting, but it also lets e-mails carry viruses that can damage networks or spread to other networks. The manufacturer of this software is now releasing security updates, but it seems that every time it introduces a security update, someone comes up with a new way around it.

Many operating system manufacturers are completely rethinking security measures. They've recognized that the products they produce can't protect the companies that use them from data loss or abuse. It has become such a problem for many customers that security support is now made available by most operating system and network software manufacturers. In the past, software manufacturers hid security vulnerabilities; now those vulnerabilities are published, and solutions are provided as soon as a vulnerability is discovered. Of course, this situation helps hackers who know that these changes won't be made on many computer systems for a while.

In the most basic sense, progress is the computer security expert's worst nightmare. As a Security+ certification holder, you're part of the team that must evaluate threats to the systems currently installed and proactively be able to anticipate what should be done to keep your systems secure.

Dealing with Telephony Issues

When telephone technology is married with information technology, the result is known as *telephony*. A breach in your telephony infrastructure is just as devastating as any other violation and can lead to the loss of valuable data.

With the exodus from land lines to Voice over IP (VoIP) in order for companies to save money in full swing, it is imperative that you treat this part of the network the same as you would any other. As an example of some of the information available, SecureLogix markets a voice firewall (`http://www.securelogix.com/ip-telephony-security.html`), and Cisco has published a white paper on IP Telephony Security in Depth (`http://www.cisco.com/warp/public/cc/so/cuso/epso/sqfr/safip_wp.pdf`).

From a security standpoint, the biggest problem with VoIP and data being on the same line is that they are then both vulnerable in the event of a PBX (Private Brance Exchange) attack.

Summary

In this chapter, we covered the key elements that an information security specialist should consider. Every organization has a different set of priorities and a different focus when it comes to security. Your responsibility is to take this information and create or maintain a security-oriented environment to address these priorities and concerns.

You'll encounter the following primary areas of responsibility:

- Physical security
- Operational security
- Management and policies

You should consider actions that you perform in this environment to accomplish one or more of the goals of information security:

- Prevention
- Detection
- Response

Security is a set of processes and products. In order for a security program to be effective, all of the following parts must work and be coordinated by the organization:

- Antivirus software
- Access control
- Authentication

Typically, your network will run many different protocols and services. These protocols allow connections to other networks and products. However, they also create potential vulnerabilities that must be understood. You must work to find ways to minimize the vulnerabilities. Many protocols and services offered by modern operating systems are highly vulnerable to attack. New methods of attacking these systems are developed every day.

Security topologies provide a mechanism to design networks that have multiple ways of implementing security. Design goals for a security topology must address these four areas of security to be effective:

- Confidentiality
- Integrity
- Availability
- Accountability

Your network can be made more secure by considering the impact of security zones and access. Here are the three most common security zones you'll encounter in the workplace:

- Internets
- Intranets
- Extranets

You can improve the likelihood of a successful security implementation if you consider putting externally accessed servers into areas called DMZs.

You can take advantage of several technologies to minimize your network's risk of being compromised:

- VLANs
- NAT
- Tunneling

The final part of this chapter discussed business requirements in a security environment:

- Identifying assets
- Assessing risks
- Identifying threats
- Evaluating vulnerabilities

Exam Essentials

Know the various aspects of information security. Ensuring a secure network involves good design, implementation, and maintenance. The information in your organization is potentially vulnerable to both internal and external threats. Identify these threats and create methods of countering them before they happen.

Identify the potential physical, operational, and management policy decisions that affect your information security efforts. It isn't good enough to have a plan if the plan is unsound or has gaping holes. You must make sure that the plans you develop and the procedures you follow to ensure security make sense for the organization and are effective in addressing the organization's needs.

Explain the relative advantages of the technologies available to you for authentication. You have many tools available to establish authentication processes. Some of these tools start with a password and user ID. Others involve physical devices or the physical characteristics of the person who is requesting authentication. This area is referred to as I&A.

Be able to explain the relative capabilities of the technologies available to you for network security. In most situations, you can create virtual LANs, create connections that are encrypted, and isolate high-risk assets from low-risk assets. You can do so using tunneling, DMZs, and network segmenting.

Identify and describe the goals of information security. The three primary goals of information security are prevention, detection, and response. Your policies and systems must include these three aspects to be effective. Ideally, you want to prevent a security breach. If a breach happens, you should have methods to detect and respond to it as quickly as possible.

Be able to describe the processes and mechanisms that can be used to implement a secure environment. Antivirus software, access control, and authentication are the three primary methods you have to implement a secure environment.

Identify the various access control methods used in systems and networks. Three primary access control methods are used in computer systems today: MAC, DAC, and RBAC. The MAC method establishes all connections and relationships between users statically. The DAC method allows the user to have some control over what information and resources are accessible. The RBAC method sets access levels and permissions based on the role the user plays in a particular situation or job.

Know which services and protocols should be offered and which should not. Many protocols and services in modern operating systems offer little if any security. These protocols and services may also be vulnerable to attack or offer no encryption in the logon process. Services that should be offered include only those that are necessary for legitimate business needs.

Be able to identify the three aspects of design goals of any security topology. The design goals of a security topology must take into consideration the need for confidentiality, integrity, and availability. These three aspects are called the CIA of security topology. Additionally, you must consider issues of accountability. Who owns the data or is responsible for verifying that it is accurate?

Know the characteristics of the three types of commonly used security zones. The three common security zones in place are the Internet, intranets, and extranets. The Internet offers low security. Intranets are considered high security, and extranets may be low to high security. Anytime you connect your network to another network, you increase the vulnerability of your network. One of the primary tools you can use to isolate less-secure resources from more-secure resources is a DMZ.

Be able to identify the differences and characteristics of the technologies available to you. A network can be segmented, and VLANs can be created to improve security. NAT presents only one Internet address to the world, hiding the other elements of the network. Tunneling allows you to make relatively secure connections to other networks using the Internet.

Identify the four business requirements of a network security design. Identifying assets, assessing risks, identifying threats, and evaluating vulnerabilities are the four primary business requirements that must be considered in a security design.

Hands-On Labs

The labs in this chapter are as follows:

Lab 1.1: Update a Linux System

Lab 1.2: Update a Windows-Based System

Lab 1.1: Update a Linux System

It is important to keep your system current and up-to-date. As soon as a weakness in an operating system becomes known, the number of people trying to exploit that weakness grows at an almost exponential rate. In this exercise, you'll apply patches and updates to an SuSE Linux Enterprise Server (SLES).

To apply patches and updates through YaST, the primary administration tool in SLES, follow these steps:

1. Log in as root and start YaST.
2. Choose Software and then select Online Update.
3. Click Next. YaST retrieves information about new updates. If prompted for a username and password, enter these values, and then choose to install any updates that are found.

   You can also choose to install patches from a CD by selecting Software and Patch CD Update. To specify settings on the server, choose Software and then System Update. Choose Change and then Update Options.

   Finally, the YOU (YaST Online Update) server can be configured by choosing Software and then YOU Server Configuration.

Lab 1.2: Update a Windows-Based System

Whether you are running Windows Server 2003 or 2008, you'll use these steps to look for updates to your system and to begin installing them:

1. Log in as administrator and start Microsoft Internet Explorer.
2. Go to `http://v4.windowsupdate.microsoft.com/en/default.asp`.
3. Click Express. The system will be checked, and you can choose to install any updates that are found.

Review Questions

1. Of the following types of security, which would be primarily concerned with someone stealing the server from the premises?
   - **A.** Physical security
   - **B.** Operational security
   - **C.** Management and policy
   - **D.** Authentication

2. Upper management has suddenly become concerned about security. As the senior network administrator, you are asked to suggest changes that should be implemented. Which of the following access methods should you recommend if the method is to be one that is primarily based on preestablished access and can't be changed by users?
   - **A.** MAC
   - **B.** DAC
   - **C.** RBAC
   - **D.** Kerberos

3. Your office administrator is being trained to perform server backups. Which authentication method would be ideal for this situation?
   - **A.** MAC
   - **B.** DAC
   - **C.** RBAC
   - **D.** Security tokens

4. You've been assigned to mentor a junior administrator and bring him up to speed quickly. The topic you're currently explaining is authentication. Which method uses a KDC to accomplish authentication for users, programs, or systems?
   - **A.** CHAP
   - **B.** Kerberos
   - **C.** Biometrics
   - **D.** Smart cards

5. Which authentication method sends a challenge to the client that is encrypted and then sent back to the server?
   - **A.** Kerberos
   - **B.** PAP
   - **C.** DAC
   - **D.** CHAP

6. After a careful risk analysis, the value of your company's data has been increased. Accordingly, you're expected to implement authentication solutions that reflect the increased value of the data. Which of the following authentication methods uses more than one authentication process for a logon?
   - **A.** Multi-factor
   - **B.** Biometrics
   - **C.** Smart card
   - **D.** Kerberos

7. Which of the following IP addresses is within the private address range?
   - **A.** 192.1.1.5
   - **B.** 192.168.0.10
   - **C.** 192.225.5.1
   - **D.** 192.255.255.255

8. After acquiring another company, your organization is in a unique position to create a new—much larger—network from scratch. You want to take advantage of this reorganization to implement the most secure environment that users, and managers, can live with. You've already decided that the only way this will be possible is to implement security zones. Which of the following isn't an example of a type of security zone?
   - **A.** Internet
   - **B.** Intranet
   - **C.** Extranet
   - **D.** NAT

9. Which of the following protocols allows an organization to present a single TCP/IP address to the Internet while utilizing private IP addressing across the LAN?
   - **A.** NAT
   - **B.** VLAN
   - **C.** DMZ
   - **D.** Extranet

10. You're the administrator for Mercury Technical. Due to several expansions, the network has grown exponentially in size within the past two years. Which of the following is a popular method for breaking a network into smaller private networks that can coexist on the same wiring and yet be unaware of each other?
    - **A.** VLAN
    - **B.** NAT
    - **C.** MAC
    - **D.** Security zone

11. Of the following services, which one would be most likely to utilize a retinal scan?

A. Auditing

B. Authentication

C. Access control

D. Data confidentiality

12. One of the vice presidents of the company calls a meeting with information technology after a recent trip to competitors' sites. She reports that many of the companies she visited granted access to their buildings only after fingerprint scans, and she wants similar technology employed at this company. Of the following, which technology relies on a physical attribute of the user for authentication?

A. Smart card

B. Biometrics

C. Mutual authentication

D. Tokens

13. Which technology allows a connection to be made between two networks using a secure protocol?

A. Tunneling

B. VLAN

C. Internet

D. Extranet

14. A new director of information technology has been hired, and you report directly to him. At the first meeting, he assigns you the task of identifying all the company resources that IT is responsible for and assigning a value to each. The process of determining the value of information or equipment in an organization is referred to as which of the following?

A. Asset identification

B. Risk assessment

C. Threat identification

D. Vulnerabilities scan

15. You have been asked to address a management meeting and present the types of threats your organization could face from hackers. Which of the following would best categorize this type of information?

A. Asset identification

B. Risk assessment

C. Threat identification

D. Vulnerabilities

16. Over the years, your company has upgraded its operating systems and networks as it has grown. A recent survey shows that numerous databases on the network haven't been accessed in more than a year. Unfortunately, the survey doesn't identify who created or last accessed those databases. Which aspect of design goals would involve determining who owns a particular database file?
    - **A.** Auditing
    - **B.** Access control
    - **C.** Threat analysis
    - **D.** Accountability

17. A user just complained to you that his system has been infected with a new virus. Which of the following would be a first step to take in addressing and correcting this problem?
    - **A.** Verifying that the most current virus definition file is installed
    - **B.** Reformatting the hard disk
    - **C.** Reinstalling the operating system
    - **D.** Disabling the user's e-mail account

18. You're awakened in the middle of the night by a frantic junior administrator. The caller reports that the guest account—which you have forbidden anyone to use—suddenly logged in and out of the network, and the administrator believes an attack occurred. Which of the following would be the most useful in determining what was accessed during an external attack?
    - **A.** System logs
    - **B.** Antivirus software
    - **C.** Kerberos
    - **D.** Biometrics

19. You want to install a server in the network area that provides web services to Internet clients. You don't want to expose your internal network to additional risks. Which method should you implement to accomplish this?
    - **A.** Install the server in an intranet.
    - **B.** Install the server in a DMZ.
    - **C.** Install the server in a VLAN.
    - **D.** Install the server in an extranet.

20. Your company provides medical data to doctors from a worldwide database. Because of the sensitive nature of the data you work with, it's imperative that authentication be established on each session and be valid only for that session. Which of the following authentication methods provides credentials that are valid only during a single session?
    - **A.** Tokens
    - **B.** Certificate
    - **C.** Smart card
    - **D.** Kerberos

Answers to Review Questions

1. A. Physical security is primarily concerned with the loss or theft of physical assets. This would include theft, fire, and other acts that physically deny a service or information to the organization.

2. A. Mandatory Access Control (MAC) is oriented toward preestablished access. This access is typically established by network administrators and can't be changed by users.

3. C. Role-Based Access Control (RBAC) allows specific people to be assigned to specific roles with specific privileges. A backup operator would need administrative privileges to back up a server. This privilege would be limited to the role and wouldn't be present during the employee's normal job functions.

4. B. Kerberos uses a Key Distribution Center (KDC) to authenticate a principle. The KDC provides a credential that can be used by all Kerberos-enabled servers and applications.

5. D. Challenge Handshake Authentication Protocol (CHAP) sends a challenge to the originating client. This challenge is sent back to the server, and the encryption results are compared. If the challenge is successful, the client is logged on.

6. A. A multi-factor authentication method uses two or more processes for logon. A two-factor method might use smart cards and biometrics for logon.

7. B. The private address range includes IP addresses between 192.168.0.0 and 192.168.255.255.

8. D. Network Address Translation (NAT) is a method of hiding TCP/IP addresses from other networks. The Internet, intranets, and extranets are the three most common security zones in use.

9. A. Network Address Translation (NAT) allows an organization to present a single address to the Internet. Typically, the router or NAT server accomplishes this. The router or NAT server maps all inbound and outbound requests and maintains a table for returned messages.

10. A. Virtual local area networks (VLANs) break a large network into smaller networks. These networks can coexist on the same wiring and be unaware of each other. A router or other routing-type device would be needed to connect these VLANs.

11. B. Authentication is a service that requests the principal user to provide proof of their identity. A retinal scan is a very secure form of evidence used in high-security companies and government agencies.

12. B. Biometric technologies rely on a physical characteristic of the user to verify identity. Biometric devices typically use either a hand pattern or a retinal scan to accomplish this.

13. A. Tunneling allows a network to make a secure connection to another network through the Internet or other network. Tunnels are usually secure and present themselves as extensions of both networks.

14. A. Asset identification is the process of identifying the types and values of assets in an organization.

15. C. A threat assessment examines the potential for internal and external threats to your systems and information.

16. D. Accountability involves identifying who owns or is responsible for the accuracy of certain information in an organization. The department or individual that is accountable for certain information would also be responsible for verifying accuracy in the event of a data-tampering incident.

17. A. Your first step would be to verify that the user's antivirus software is the most current version. This includes checking the virus definition files.

18. A. System logs will frequently tell you what was accessed and in what manner. These logs are usually explicit in describing the events that occurred during a security violation.

19. B. A DMZ is an area in a network that allows access to outside users while not exposing your internal users to additional threats.

20. A. Tokens are created when a user or system successfully authenticates. The token is destroyed when the session is over.