

Network Security Basics

Solutions in this chapter:

- Security Overview
- Defining Basic Security Concepts
- Addressing Security Objectives
- Recognizing Network Security Threats
- Designing a Comprehensive Security Plan

Summary

Introduction

Before you can understand firewalls and how ISA Server 2006 works, you need to look at the big picture: what we mean by network security in general – and Internet security in particular – why it's necessary, how we can create a comprehensive security policy to protect our networks from unauthorized access, and where ISA Server fits into that picture.

Network security is a big topic and is growing into a high profile (and often highly paid) Information Technology (IT) specialty area. Security-related websites are tremendously popular with savvy Internet users. The popularity of security-related certifications has expanded. Esoteric security measures like biometric identification and authentication – formerly the province of science fiction writers and perhaps a few ultra-secretive government agencies – have become commonplace in corporate America. Yet, with all this focus on security, many organizations still implement security measures in an almost haphazard way, with no well-thought-out plan for making all the parts fit together. Computer security involves many aspects, from protection of the physical equipment to protection of the electronic bits and bytes that make up the information that resides on the network.

In the next section, we will provide a brief overview of what we mean by “security” and how it applies to your computer network.

NOTE

This chapter focuses on generic computer and Internet security concepts and how to develop a comprehensive security plan for your organization. The rest of this book will discuss how ISA Server fits into that security plan.

Security Overview

The term *computer security* encompasses many related, yet separate, topics. These can be stated as *security objectives*, and include:

- Control of physical accessibility to the computer(s) and/or network
- Prevention of accidental erasure, modification or compromise of data
- Detection and prevention of intentional internal security breaches
- Detection and prevention of unauthorized external intrusions (hacking)

Network security solutions are loosely divided into three categories: *hardware*, *software* and *human*. In this chapter, we will provide an overview of basic security concepts. Then, we will examine the four security objectives and look at each of the three categories of security solutions.

Defining Basic Security Concepts

A generic definition of *security* is “freedom from risk or danger; safety” (The American Heritage Dictionary).

This definition is perhaps a little misleading when it comes to computer and networking security, as it implies a degree of protection that is inherently impossible in the modern connectivity-oriented computing environment.

This is why the same dictionary provides another definition specific to computer science: “The *level to which* a program or device is safe from unauthorized use [emphasis added].” Implicit in this definition is the caveat that the objectives of security and accessibility – the two top priorities on the minds of many network administrators – are, by their very natures, diametrically opposed. The more accessible your data is, the less secure it is. Likewise, the more tightly you secure it, the more you impede accessibility. Any security plan is an attempt to strike the proper balance between the two.

As in any other specialty field, security professionals speak a language all their own and understanding the concepts requires that you learn the jargon. At the end of this section, you will find a list of some common terms that you are likely to encounter in the IT security field.

Knowledge is Power

The above title is a famous hacker’s motto (along with such other gems as “Information wants to be free,” and the simplistic but optimistic, “Hack the world!”). However, it is a truism that applies not only to those attempting to gain access to data they aren’t supposed to see, but also to those who are trying to protect themselves from the intruders. The first step in winning any battle – and network security *is* a battle over the ownership and control of your computer files – is the same as it’s always been: “know thine enemy.”

To protect your network resources from theft, damage, or unwanted exposure, you must understand who initiates these things, why, and how they do it. Knowledge will make *you* powerful, too – and better able to prevent unauthorized intrusions into your network. In the section entitled *Detecting and Preventing Unauthorized External Intrusions*, we will discuss the various motivations that drive different network intruders and the types of people who make a practice of “breaking and entering” networks.

The very best place to learn is from the hackers themselves. Many network administrators and even some security specialists eschew the books and websites that are written to a hacker audience or from the hacker’s point of view. This may be because one fears “guilt by association” or believes that it would be somehow demeaning to hang out with the hackers. This attitude may be based on high moral ground, but strategically, it’s a mistake.

Think Like a Thief

It is well known in law enforcement circles that the best criminal investigators are those who are best able to “get inside the mind” of the lawbreaker. Network intrusion detectives will find that the same is true – to prevent your network from falling prey to hackers, or to catch data thieves when they do get in, requires that you be able to adopt a mindset emulating theirs.


This means learning to anticipate the intruder’s actions. First, you must determine *what* needs to be protected, and to what degree. A wealthy person not only establishes a general security perimeter by building fences around the house and locking doors and windows, but also places the most valuable items in a wall or floor safe. This provides multiple *layers* of protection. The practice of implementing multiple layers of protection is known as *defense in depth*.

ISA Server can be an important layer of protection in your organization’s security plan.

The Intrusion Triangle

Borrowing again from the law enforcement community, crime prevention specialists use a model called the “Crime Triangle” to explain that certain criteria must exist before a crime can occur. We can adapt this same triangle to network security: the same three criteria must exist before a network security breach can take place. The three “legs” or points of the triangle are shown in Figure 1.1.

Figure 1.1 All three legs of the triangle must exist for a network intrusion to occur



Let's look at each point individually:

- **Motive:** An intruder must have a reason to want to breach the security of your network (even if the reason is “just for fun”); otherwise, he/she won't bother.
- **Means:** An intruder must have the ability (either the programming knowledge, or, in the case of “script kiddies,” the intrusion software written by others), or he/she won't be able to breach your security.
- **Opportunity:** An intruder must have the chance to enter the network, either because of flaws in your security plan, holes in a software program that open an avenue of access, or physical proximity to network components; if there is no opportunity to intrude, the would-be hacker will go elsewhere.

If you think about the three-point intrusion criteria for a moment, you'll see that there is really only one leg of the triangle over which you, as the network administrator or security specialist, have any control. It is unlikely that you can do much to remove the intruder's *motive*. The motive is likely to be built into the type of data you have on the network or even the personality of the intruder him/herself. It is also not possible for you to prevent the intruder from having or obtaining the *means* to breach your security. Programming knowledge is freely available, and there are many experienced hackers out there who are more than happy to help out a less-sophisticated ones. The one thing that you *can* affect is the *opportunity* afforded the hacker.

Removing Intrusion Opportunities

Crime prevention officers tell members of the community that the “good guys” probably can’t keep a potential burglar from wanting to steal, and they certainly can’t keep the potential burglar from obtaining burglary tools or learning the “tricks of the trade.” What citizens *can* do is take away, as much as possible, the opportunity for the burglar to target their own homes.

This means putting dead-bolt locks on the doors (and using them), getting a big, loud, unfriendly dog, installing an alarm system, and the like. In other words, as a homeowner, your goal is not to prevent the burglar from burglarizing, but to make your own home a less desirable target. As a network “owner,” your objective is to “harden” your own network so that all those hackers out there who already have the motive and the means will look for an easier victim.

The best and most expensive locks in the world won’t keep intruders out of your house if you don’t use them. And if those locks are difficult to use and result in inconvenience to you in your everyday comings and goings, you probably *won’t* use them – at least, not all the time. A poorly implemented network security system that is difficult to administer or that unduly inconveniences network users may end up similarly unused; eventually, you will throw your hands up in frustration and just turn the darn thing off. And that will leave your network wide open to intruders.

A good network security system will help you to remove the temptations (open ports, exploitable applications) easily and will be as transparent to your users as possible. ISA Server, when properly configured, meets these requirements – and more. We will discuss the characteristics of a good network security system component further in the section entitled “Preventing and Detecting Unauthorized External Intrusions.”

Security Terminology

Every industry has its own “language,” the jargon that describes concepts and procedures peculiar to the field. Computer networking is infamous for the “technotalk” and the proliferation of acronyms that often mystify outsiders. Specialty areas within an industry often have their own brands of jargon, as well, and the computer security sub-field is no exception.

It is not possible to provide a complete glossary of security-related terms within the scope of this chapter, but in this section, we will define some of the more common words and phrases that you may encounter as you begin to explore the fascinating world of computer security:

- **Attack** In the context of computer/network security, an attack is an attempt to access resources on a computer or a network without authorization, or to bypass security measures that are in place.
- **Audit** To track security-related events, such as logging onto the system or network, accessing objects, or exercising user/group rights or privileges.
- **Availability of data** Reliable and timely access to data.
- **Breach** Successfully defeating security measures to gain access to data or resources without authorization, or to make data or resources available to unauthorized persons, or to delete or alter computer files.
- **Brute force attack** Attempt to “crack” passwords by sequentially trying all possible combinations of characters until the right combination works to allow access.

- **Buffer** A holding area for data.
- **Buffer overflow** A way to crash a system by putting more data into a buffer than the buffer is able to hold.
- **CIA triad** Confidentiality, Integrity, and Availability of data. Ensuring the confidentiality, integrity, and availability of data and services are primary security objectives that are often related to each other. See also *availability of data*, *confidentiality of data*, and *integrity of data*.
- **Confidentiality of data** Ensuring that the contents of messages will be kept secret. See also *integrity of data*.
- **Countermeasures** Steps taken to prevent or respond to an attack or malicious code.
- **Cracker** A hacker who specializes in “cracking” or discovering system passwords to gain access to computer systems without authorization. See also *hacker*.
- **Crash** Sudden failure of a computer system, rendering it unusable.
- **Defense-in-depth** The practice of implementing multiple layers of security. Effective defense-in-depth strategies do not limit themselves to focusing on technology, but also focus on operations and people. For example, a firewall can protect against unauthorized intrusion, but training and the implementation of well-considered security policies help to ensure that the firewall is properly configured.
- **Denial of Service attack** A deliberate action that keeps a computer or network from functioning as intended (for example, preventing users from being able to log onto the network).
- **Exposure** A measure of the extent to which a network or individual computer is open to attack, based on its particular vulnerabilities, how well known it is to hackers, and the time duration during which intruders have the opportunity to attack. For example, a computer using a dialup analog connection has less exposure to attack coming over the Internet, because it is connected for a shorter period of time than those using “always-on” connections such as cable, DSL or T-carrier.
- **Hacker** A person who spends time learning the details of computer programming and operating systems, how to test the limits of their capabilities, and where their vulnerabilities lie. See also *cracker*.
- **Integrity of data** Ensuring that data has not been modified or altered, that the data received is identical to the data that was sent.
- **Least privilege** The principle of least privilege requires that users and administrators have only the minimum level of access to perform their job-related duties. In military parlance, the principle of least privilege is referred to as *need to know*.
- **Malicious code** A computer program or script that performs an action that intentionally damages a system or data, that performs another unauthorized purpose, or that provides unauthorized access to the system.
- **Penetration testing** Evaluating a system by attempting to circumvent the computer’s or network’s security measures.

- **Reliability** The probability of a computer system or network continuing to perform in a satisfactory manner for a specific time period under normal operating conditions.
- **Risk** The probability that a specific security threat will be able to exploit a system vulnerability, resulting in damage, loss of data, or other undesired results. That is, a risk is the sum of the threat plus the vulnerability.
- **Risk management** The process of identifying, controlling, and either minimizing or completely eliminating events that pose a threat to system reliability, data integrity, and data confidentiality.
- **Sniffer** A program that captures data as it travels across a network. Also called a *packet sniffer*.
- **Social engineering** Gaining unauthorized access to a system or network by subverting personnel (for example, posing as a member of the IT department to convince users to reveal their passwords).
- **TCSEC** Trusted Computer System Evaluation Criteria. A means of evaluating the level of security of a system.
- **Technical vulnerability** A flaw or bug in the hardware or software components of a system that leaves it vulnerable to security breach.
- **Threat** A potential danger to data or systems. A threat agent can be a virus; a hacker; a natural phenomenon, such as a tornado; a disgruntled employee; a competitor, and other menaces.
- **Trojan horse** A computer program that appears to perform a desirable function but contains hidden code that is intended to allow unauthorized collection, modification or destruction of data.
- **Virus** A program that is introduced onto a system or network for the purpose of performing an unauthorized action (which can vary from popping up a harmless message to destroying all data on the hard disk).
- **Vulnerability** A weakness in the hardware or software or security plan that leaves a system or network open to threat of unauthorized access or damage or destruction of data.
- **Worm** A program that replicates itself, spreading from one machine to another across a network.

Once you are comfortable with the terminology, you can begin to address the individual objectives that will assist you in realizing your goal to create a secure network environment.

Addressing Security Objectives

If our security goal is to have complete control over what data comes into and goes out of our networks, we must define objectives that will help us reach that goal. We listed some general security objectives related to computer networks – especially those connected to an outside internetwork such as the Global Internet – as controlling physical access, preventing accidental compromise of data, detecting and

preventing intentional internal security breaches, and detecting and preventing unauthorized external intrusions. In the following sections, we will examine each of these objectives in detail.

Controlling Physical Access

One of the most important, and at the same time most overlooked aspects of a comprehensive network security plan is physical access control. This matter is often left up to facilities managers or plant security departments, or it is outsourced to security guard companies. Network administrators frequently concern themselves with sophisticated software and hardware solutions that prevent intruders from accessing internal computers remotely, while doing nothing to protect the servers, routers, cable, and other physical components of the network from direct access.

Thinking Outside the Box About Security

In far too many supposedly security-conscious organizations, computers are locked away from employees and visitors all day, only to be left open at night to the janitorial staff, which has keys to all offices. It is not at all uncommon for computer espionage experts to pose as members of the cleaning crew to gain physical access to machines that hold sensitive data. This is a favorite ploy for several reasons:

- Cleaning services are often contracted out, and workers in the industry are often transient, so that company employees may not be easily aware of who is or isn't a legitimate employee of the cleaning company.
- Cleaning is usually done late at night, when all or most company employees are gone, making it easier to surreptitiously steal data.
- Cleaning crew members are often paid little or no attention by company employees, who take their presence for granted and think nothing of their being in areas where the presence of others might be questioned.

Physically breaking into the server room and stealing the hard disk on which sensitive data resides may be a crude method; nonetheless, it happens. In some organizations, it may be the easiest way to gain unauthorized access, especially for an intruder who has help “on the inside.”

Physical Access Factors

It is important for you to make physical access control the “outer perimeter” of your security plan. This means:

- Controlling physical access to the servers
- Controlling physical access to networked workstations

- Controlling physical access to network devices
- Controlling physical access to the cable
- Being aware of security considerations with wireless media
- Being aware of security considerations related to portable computers
- Recognizing the security risk of allowing data to be printed out
- Recognizing the security risks involving floppy disks, CDs, tapes, and other removable media

Let's look at why each of these is important and how you can implement a physical security plan that addresses all these factors.

Protecting the Servers

File servers on which sensitive data is stored and infrastructure servers that provide mission critical services such as logon authentication and access control should be placed in a highly secure location. At the minimum, servers should be in a locked room where only those who need to work directly with the servers have access. Keys should be distributed sparingly, and records should be kept of issuance and return.

If security needs are high due to the nature of the business or the nature of the data, access to the server room may be controlled by magnetic card, electronic locks requiring entry of a numerical code, or even biometric access control devices such as fingerprint or retinal scanners. Both ingress and egress should be controlled – ideally with logs, video cameras, and/or other means of recording both who enters and who exits.

Other security measures include monitor detectors or other alarm systems, activated during non-business hours, and security cameras. A security guard or company should monitor these devices.

Keeping Workstations Secure

Many network security plans focus on the servers but ignore the risk posed by workstations with network access to those servers. It is not uncommon for employees to leave their computers unsecured when they leave for lunch or even when they leave for the evening. Often there will be a workstation in the receptionist area that is open to visitors who walk in off the street. If the receptionist must leave briefly, the computer – and the network to which it is connected – is vulnerable unless steps have been taken to ensure that it is secure.

A good security plan includes protection of all unmanned workstations. A secure client operating system such as Windows NT or Windows 2000 requires an interactive logon with a valid account name and password in order to access the operating system (unlike Windows 9x). This allows users to “lock” the workstation when they are going to be away from it so someone else can't just step up and start using the computer.

However, don't depend on access permissions and other software security methods alone to protect your network. If a potential intruder can gain physical access to a networked computer, he/she is that much closer to accessing your valuable data or introducing a virus onto your network.

Ensure all workstation users adhere to a good password policy, as discussed in the section entitled *Planning a Comprehensive Security Plan* later in this chapter.

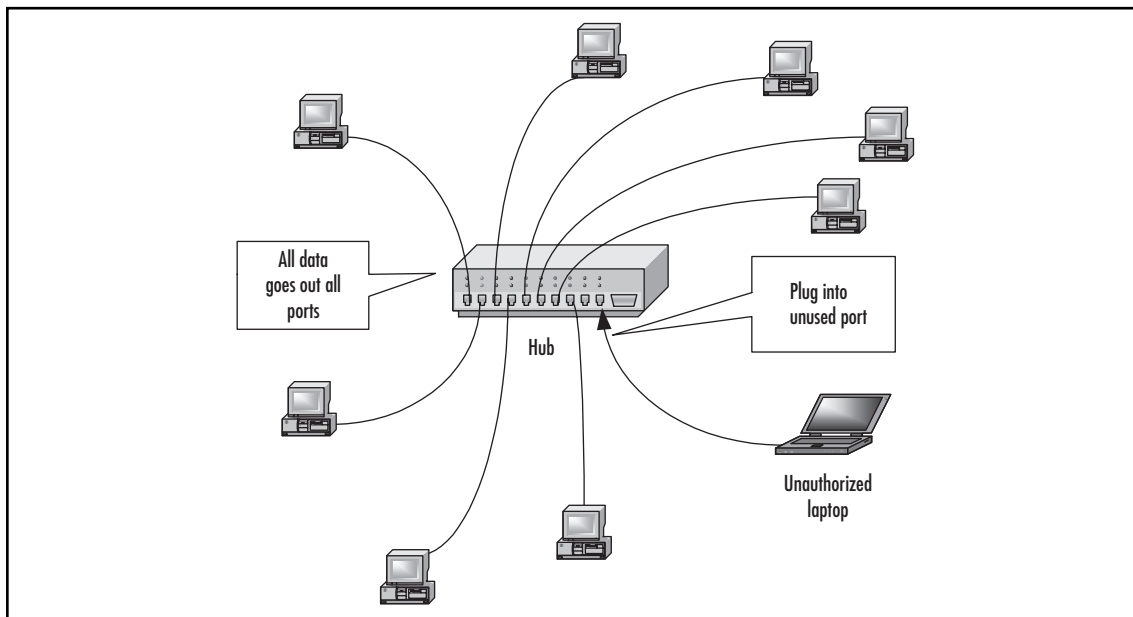
Many modern PC cases come with some type of locking mechanism that will help prevent an unauthorized person from opening the case and stealing the hard disk. Locks are also available to prevent use of the floppy drive, copying data to diskette, and/or rebooting the computer with a floppy.

Protecting Network Devices

Hubs, routers, switches and other network devices should be physically secured from unauthorized access. It is easy to forget that just because a device doesn't have a monitor on which you can *see* data, this does not mean the data can't be captured or destroyed at that access point.

For example, a traditional Ethernet hub sends all data out every port on the hub. An intruder who has access to the hub can plug a packet-sniffing device (or a laptop computer with sniffer software) that operates in "promiscuous mode" into a spare port and capture data sent to any computer on the segment, as shown in Figure 1.2.

Figure 1.2 An intruder who has access to the hub can easily intercept data



Although switches and routers are somewhat more secure, any device through which the data passes is a point of vulnerability. Replacing hubs with switches and routers makes it more difficult for an intruder to "sniff" on your network, but it is still possible to use techniques such as Address Resolution Protocol (ARP) spoofing. This is sometimes called *router redirection*, in which nearby machines are redirected to forward traffic through an intruder's machine by sending ARP packets that contain the router's Internet Protocol (IP) address mapped to the intruder's machine's MAC address. This results in other machines believing the intruder's machine is the router, and so they send their traffic to it. A similar method uses Internet Control Message Protocol (ICMP) router advertisement messages.

It is also possible, with certain switches, to overflow the address tables with multiple false Media Access Control (MAC) addresses or send a continuous flow of random garbage through the switch to

trigger it to change from bridging mode to repeating mode. This means all frames will be broadcast on all ports, giving the intruder the same opportunity to access the data that he would have with a regular hub. This is called *switch jamming*.

Finally, if the switch has a special monitor port designed to be used with a sniffer for legitimate (network troubleshooting) purposes, an intruder who has physical access to the switch can simply plug into this port and capture network data.

Your network devices should be placed in a locked room or closet and protected in the same manner as your servers.

How Packet Sniffers Work

Packet sniffer/protocol analyzer devices and programs are not used solely for nefarious purposes, although intruders use them to capture unencrypted data and clear-text passwords that will allow them to break into systems. Despite the fact that they can be used to “steal” data as it travels across the network, they are also invaluable troubleshooting tools for network administrators.

The sniffer captures individual data packets and allows you to view and analyze the message contents and packet headers. This can be useful in diagnosing network communications problems and uncovering network bottlenecks that are impacting performance. Packet sniffers can also be turned against hackers and crackers and used to discover unauthorized intruders.

The most important part of the sniffer is the capture driver. This is the component that captures the network traffic, filters it (according to criteria set by the user), and stores the data in a buffer. The packets can then be analyzed and decoded to display the contents.

It is often possible to detect an unauthorized packet sniffer on the wire using a device called a Time Domain Reflectometer (TDR), which sends a pulse down the cable and creates a graph of the reflections that are returned. Those who know how to read the graph can tell whether unauthorized devices are attached to the cable and where.

Other ways of detecting unauthorized connections include monitoring hub or switch lights using Simple Network Monitoring Protocol (SNMP) managers that log connections and disconnections or using one of the many tools designed for the specific purpose of detecting sniffers on the network. There are also several techniques using Packet Internet Groper (ping), ARP, and DNS that may help you to catch unauthorized sniffers.

Securing the Cable

The next step in protecting your network data is to secure the cable across which it travels. Twisted pair and coaxial cable are both vulnerable to data capture; an intruder who has access to the cable

can tap into it and eavesdrop on messages being sent across it. A number of companies make “tapping” devices.

Fiber optic cable is more difficult to tap into because it does not produce electrical pulses, but instead, uses pulses of light to represent the 0s and 1s of binary data. It is, however, possible for a sophisticated intruder to use an optical splitter and tap into the signal on fiber optic media.

Compromise of security at the physical level is a special threat when network cables are not contained in one facility but span a distance between buildings. There is even a name for this risk, “manhole manipulation,” referring to the easy access intruders often have to cabling that runs through underground conduits.

Cable taps can sometimes be detected by using a TDR or optical TDR to measure the strength of the signal and determine where the tap is located.

Safely Going Wireless

Wireless media is becoming more and more popular as our society becomes more mobile, and many predict it will be next big thing in networking during the first years of the new millennium.

Large companies such as Cisco Systems, Lucent Technologies, Sun Microsystems, and Microsoft have invested large amounts of talent and money into the wireless initiative. Wireless Internet access based on the Wireless Access Protocol (WAP) is common in Europe and beginning to catch on in the U.S. Fixed wireless services are offered by communications giants such as AT&T and Sprint and companies such as Metricom (which offers the Ricochet wireless service).

Wireless networking offers several distinct advantages over traditional cabled networking. Laptop users can easily connect and disconnect as they come and go. Workers out in the field can maintain network communications in areas where there are no cables or phone lines. For professions such as policing, where employees work from a moving vehicle most of the time, wireless is the only way to stay connected to the department LAN. For telecommuters in rural areas where DSL and cable modem access are unavailable, wireless technologies such as satellite provide a broadband alternative to slow analog modems.

There are several different varieties of wireless networking, including:

- Radio (narrow band or spread spectrum)
- Satellite/microwave
- Laser/infrared

The most popular wireless technologies are radio-based and operate according to the IEEE 802.x standards. 802.11b (and increasingly, 802.11g, which is backwardly compatible with b) networks are becoming commonplace as commercial “hot spots” spring up in major cities and businesses and home computer users implement wireless networks because of their convenience. Wireless connectivity is available at hotels, airports, and even coffee shops and restaurants.

Despite the many benefits of these wireless technologies, they also present special problems, especially in the area of network security. Wireless is more vulnerable to inception of data than cabled media. Radio and microwave are known as broadcast media. Because the signals are transmitted across the airwaves, any receiver set to the correct frequency can easily eavesdrop on the communications.

The practice of “war driving” (going out with a wireless NIC-equipped laptop or handheld system and looking for open wireless networks to which they can connect) is a favorite pastime of hackers.

NOTE

Laser signals are not as easy to intercept; however, because laser is a line-of-sight technology, it is more limited in application – and lasers are much more sensitive to environmental factors, such as weather.

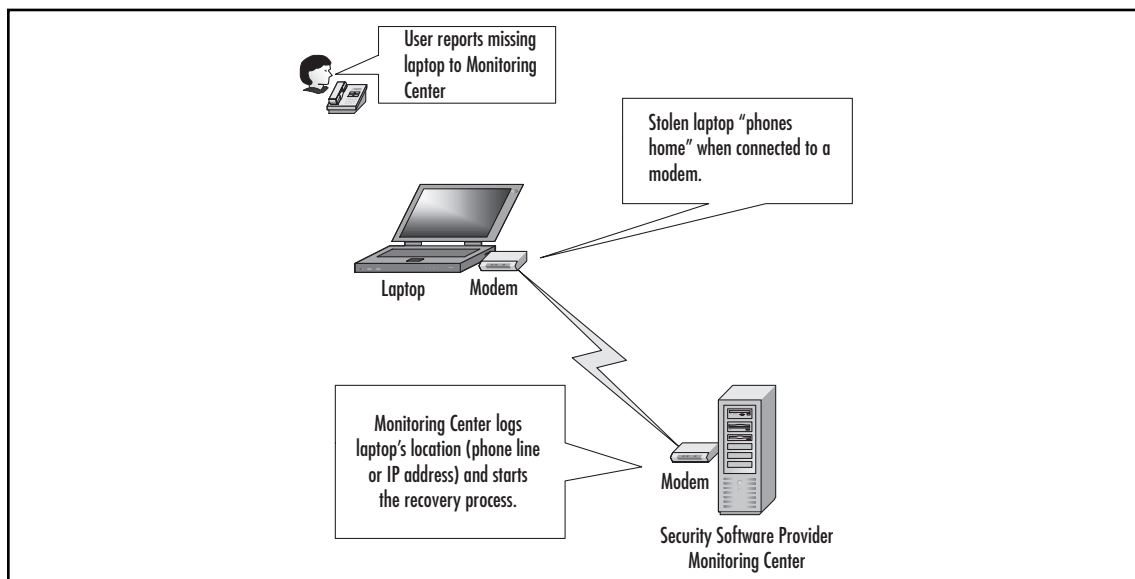
If security is a priority, any data sent via radio or microwave links should be encrypted.

Have Laptop, Will Travel


Portable computers – laptops, notebooks, and new fully functional handheld computers such as the Pocket PC and Palm machines – present their own security problems based on the very features that make them popular– their small size and mobility. Physical security for portable computers is especially important because it is so easy to steal the entire machine, data and all.

Luckily, there are a large number of companies that make theft protection devices and security software for laptops. Locks and alarms are widely available, along with software programs that will disable the laptop’s functionality if it is stolen, or even help track it down by causing the computer to “phone home” the first time the portable computer is attached to a modem (see Figure 1.3).

Figure 1.3 Tracking programs help recover stolen portable computers



Some laptops come with removable hard disks. It is a good idea if you have highly sensitive data that must be accessed with your laptop to store it on a removable disk (PC Card disks and those that plug into the parallel port are widely available) and encrypt it. Separate the disk from the computer when it is not in use.



Theft recovery/tracking software for laptops includes Computrace www.computrace.com from Absolute Software Corporation, Alert PC www.sentryinc.com from Computer Sentry Software. TrackIT www.trackitcorp.com is a hardware anti-theft device for computer cases and other baggage.

The possibility of theft is not the only way in which laptops present a security risk. The threat to your network is that a data thief who is able to enter your premises may be able to plug a laptop into the network, crack passwords (or obtain a password via social engineering), and download data to the portable machine, which can then be easily carried away.

New handheld computers are coming with more security devices built in. For example, the Hewlett-Packard iPAQ 5555 includes biometric (fingerprint recognition) technology to prevent unauthorized users from accessing the data.

The Paper Chase

Network security specialists and administrators tend to concentrate on protecting data in electronic form, but you should recognize that intruders may also steal confidential digital information by printing it out or locating a hard copy that was printed by someone else. It does little good to implement strong password policies and network access controls if employees can print out sensitive material and then leave it lying on desks, stored in unlocked file cabinets, or thrown into an easily accessed trash basket. “Dumpster diving” (searching the trash for company secrets) is a common form of corporate espionage – and one that surprisingly often yields results.

If confidential data must be printed, the paper copy should be kept as physically secure as the digital version. Disposal should require shredding, and in cases of particularly high-security information, the shredded paper can be mixed with water to create a pulp that is impossible to put back together again.

Removable Storage Risks

Yet another potential point of failure in your network security plan involves saving data to removable media. Floppy diskettes, zip and jaz disks, tapes, PC cards, CDs and DVDs containing sensitive data must be kept physically secured at all times.

Don’t make the mistake of thinking that deleting the files on a disk, or even formatting the disk, completely erases the data; it is still there until it has been overwritten and can be retrieved using special software.

NOTE

The residual physical representation of data that has been “erased,” from which that data can be reconstructed, is called *data remanence*. Methods used to prevent this in high-security environments include degaussing, overwriting, and in extreme cases, physical destruction of the media. Degaussing involves use of a device that generates a magnetic field to reduce the magnetic state of the media to zero, which restores it to an unrecorded state. Software (sometimes referred to as “file shredder” software) is available to overwrite all sectors of a disk with random bits in order to prevent recovery of the data.

Although removable media can present a security threat to the network, it can also play a part in your overall security plan. Removable disks (including fully bootable large capacity hard disks installed in mobile “nesting” racks) can be removed from the computer and locked in a safe or removed from the premises to protect the data that is stored there.

Physical Security Summary

Ensuring a physically secure network environment is the first step in controlling access to your network’s important data and system files, but it is only part of a good security plan. This is truer today than in the past, because networks have more “ways in” than they once did. A medium or large network may have multiple dial-in servers, VPN servers, and a dedicated full-time Internet connection. Even a small network is likely to be connected to the Internet part of the time.

Virtual intruders never set foot on your organization’s property and never touch your computers. They can access your network from across the street or from halfway across the world. But they can do as much damage as the thief who breaks into your company headquarters to steal or destroy your data – and they are much harder to catch. In the following sections, we will examine specific network security risks, and how to prevent them.

Preventing Accidental Compromise of Data

The topic of network security may bring to mind a picture of evil corporate rivals determined to steal your company’s most precious trade secrets or malevolent hackers bent on crashing your network and erasing all of your data just for the sheer joy of it. While these risks do exist, often the reality of network data loss is far less glamorous. A large proportion of erased, modified, or disclosed data is the result of the actions of employees or other authorized network personnel. And a large percentage of *that* is the result of *accidental* compromise of the data.

Unintended errors in entering data or accessing network resources or carelessness in use of the computers and network can cause loss of data or crashing of individual computers, the server, and even the network.

Your network security plan should address these unintended compromises, which can be just as disastrous as intentional breaches of security.

Know Your Users

To prevent accidental compromise of data, you should first know your users and their skill levels. Those with few technical skills should be given as little access as possible – allow them the access required to do their jobs, and no more (this philosophy is often referred to as the *principle of least privilege*, or, in government circles, as *need to know*.) Too many network users have, in all innocence, destroyed or changed important files while attempting to clear up space on their hard disks or troubleshoot a computer problem on their own.

Educate Your Users

Educating your users is one of the most important factors in eliminating or reducing such incidents, and an essential component of the multilayered “defense in depth” approach to security. This does not necessarily mean upgrading their technical skills (although it can). Turning all your users into power users may not be cost effective or otherwise desirable. What is essential is to train all of your network users in the proper procedures and rules of usage for the network.

Every person who accesses your company network should be aware of your user policies and should agree to adhere to them. This includes notifying technical support personnel immediately of any hardware or software problems, refraining from installing any unauthorized software on their machines or downloading files from the Internet without authorization, and never dialing up their personal ISPs or other networks or services from company machines without permission.

Control Your Users

In some cases, establishing clear-cut policies and making staffers and other users aware of them will be enough. In other cases, you will find that users are unable or unwilling to follow the rules, and you will have to take steps to enforce them – including locking down desktops with system/group policies and, with software such as ISA Server, implementing access rules and filtering to prevent unauthorized packets from being sent or received over the network.

Fortunately, most users will at least attempt to comply with the rules. A more serious problem is the “insider” who is looking to intentionally breach network security. This may be simply a maverick employee who doesn’t like being told what to do, or it may be someone with a darker motive.

Preventing Intentional Internal Security Breaches

According to most computer security studies, as documented in RFC 2196, *Site Security Handbook*, actual loss (in terms of money, productivity, computer reputation, and other tangible and intangible harm) is greater for internal security breaches than for those from the outside. Internal attackers are more dangerous for several reasons:

- They generally know more about the company, the network, the layout of the building(s), normal operating procedure, and other information that will make it easier for them to gain access without detection.

- They usually have at least some degree of legitimate access and may find it easy to discover passwords and holes in the current security system.
- They know what information is on the network and what actions will cause the most damage.

We discuss common motivations behind intentional security breaches, both internal and external, in the section entitled *Recognizing Network Security Threats*. Preventing such problems begins with the same methods used to prevent unintentional compromises, but goes a step further.

To a large extent, unintended breaches can be prevented through education. The best way to prevent such breaches depends, in part, on the motivations of the employee(s) concerned.

Hiring and Human Resource Policies

A good “defense in depth” security strategy is multifaceted, involving technology, operations, and people. In many cases, the latter is the weakest link in the chain. Thus, prevention starts with good human resources practices. That means management should institute hiring policies aimed at recruiting persons of good character. Background investigations should be conducted, especially for key positions that will have more than normal user access.

The work environment should encourage high employee morale. In many cases, internal security breaches are committed as “revenge” by employees who feel underpaid, under-appreciated, and even mistreated. Employees who are enthusiastic about their jobs and feel valued by the organization will be much more likely to comply with company rules, including network security policies.

Another motivation for internal breaches is money. If the company engages in a highly competitive business, competitors may approach employees with lucrative offers for trade secrets or other confidential data. If you are in a field that is vulnerable to corporate espionage, your security policies should lean toward the “deny all access” model, in which access for a particular network user starts at nothing, and access is added on the basis of the user’s need to know.

NOTE

The “deny all access” policy model is one of two basic starting points in creating a security policy. The other is “allow all access” in which all resources are open to a user unless there are specific reasons to deny access. Neither of these is “right” or “wrong,” although the “deny all access” model is undisputedly more secure, and the “allow all access” model is easier to implement. From which of these starting points you work depends on the security *philosophy* of the organization.

Detecting Internal Breaches

Implementing auditing will help you detect internal breaches of security by recording specified security events. You will be able to track when objects (such as files or folders) are accessed, what user account was used to access them, when users exercise user rights, and when users log onto or off of the computer or network. Modern network operating systems such as Windows 2000 and XP/2003 include built-in auditing functionality.

**WARNING**

You should audit only those events that are necessary to track in keeping with your security policy. Auditing too many events (and access to too many objects) will have a negative impact on your computer's performance and will make relevant events more difficult to find in the security log.

If you choose to audit many events, or often-accessed objects, the security log can grow very large, very quickly. Windows allows you to set the maximum size in kilobytes for the security log by configuring its property sheet in the Event Viewer (right-click **Security Log** and select **Properties**). You can also choose whether to overwrite previous events when the maximum size is reached or to require manual clearing of the log.

Preventing Intentional Internal Breaches

Firewalls are helpful in keeping basically compliant employees from accidentally (or out of ignorance) visiting dangerous websites or sending specific types of packets outside the local network. However, they are of more limited use in preventing intentional internal security breaches. Simply limiting their access to the external network cannot thwart insiders who are determined to destroy, modify, or copy your data. Because they have physical access, they can copy data to removable media, to a portable computer (including tiny handheld machines), or perhaps even print it to paper and remove it from the premises that way. They may change the format of the data to disguise it and upload files to web-based data storage services.

In a high security environment, computers without floppy drives – or even completely diskless workstations – may be warranted. System or group policy can be applied that prevents users from installing software (such as that needed for a desktop computer to communicate with a Pocket PC or Palm Pilot). Cases can be locked, and physical access to serial ports, USB ports, and other connection points can be covered so removable media devices can't be attached. Other internal controls include physical measures such as key cards to limit entry to server rooms and other sensitive resources, as well as software controls such as user and group accounts, encryption, and so forth.

Intentional internal breaches of security constitute a serious problem, and company policies should treat it as such.

Preventing Unauthorized External Intrusions

External intrusions (or “hacking into the system”) from outside the LAN has received a lot of attention in the media and thus is the major concern of many companies when it comes to network security issues. In recent years, there have been a number of high profile cases in which the web servers of prominent organizations (such as Yahoo and Microsoft) have been hacked. Attempts to penetrate sensitive government networks, such as the Pentagon's systems, occur on a regular basis. Distributed Denial of Service (Duos) attacks make front-page news when they crash servers and prevent Internet users from accessing popular sites.

There are psychological factors involved, as well. Internal breaches are usually seen by companies as personnel problems and handled administratively. External breaches may seem more like a “violation” and are more often prosecuted in criminal actions. Because the external intruder could come from anywhere, at any time, the sense of uncertainty and fear of the unknown may cause organizations to react in a much stronger way to this type of threat.

The good news about external intrusions is that the area(s) that must be controlled are much more focused. There are usually only a limited number of points of entry to the network from the outside. This is where a properly configured firewall can be invaluable, allowing authorized traffic into the network while keeping unauthorized traffic out. On the other hand, the popularity of firewalls ensures that dedicated hackers know how they work and spend a great deal of time and effort devising ways to defeat them.

Never depend on the firewall to provide 100 percent protection, even against outside intruders. Remember that in order to be effective, a security plan must be a multifaceted, multilayered one. We hope the firewall will keep intruders out of your network completely – but if they *do* get in, what is your contingency plan? How will you reduce the amount of damage they can do and protect your most sensitive or valuable data?

External Intruders with Internal Access

A special type of “external” intruder is the outsider who *physically* breaks into your facility to gain access to your network. Although not a true “insider,” because he is not authorized to be there and does not have a valid account on the network, he has many of the advantages of those discussed in the section on internal security breaches.

Your security policy should take into account the threats posed by this “hybrid” type of intruder.

Tactical Planning

In dealing with network intruders, you should practice what police officers in defensive tactics training call “if/then thinking.” This means considering every possible outcome of a given situation and then asking yourself, “*If* this happens, *then* what could be done to protect us from the consequences?” The answers to these questions will form the basis of your security policy.

This tactic requires that you be able to plan your responses in detail, which means you must think in specifics rather than generalities. Your security threat must be based in part on understanding the motivations of those initiating the attack and in part on the technical aspects of the type of attack that is initiated. In the next section, we will discuss common intruder motivations and specific types of network attacks.

Recognizing Network Security Threats

In order to effectively protect your network, you must consider the following question: from *whom* or *what* are you protecting it? In this section, we will approach the answer to that question from two perspectives:

- *Who*: types of network intruders and their motivations
- *What*: types of network attackers and how they work

These questions form the basis for performing a *threat analysis*. A comprehensive threat analysis is often the product of collaborative brainstorming among people who are knowledgeable about the business processes, industry, security, and so on. In fact, it is desirable that a threat analysis not be conducted solely by computer security experts, as this group might lack important “big picture” knowledge of the business and industry. The ability to think creatively is a key requirement for members of a threat analysis team.

First, we will look at intruder motivations and classify the different types of people who have the skill and desire to hack into others’ computers and networks.

Understanding Intruder Motivations

There are probably as many different specific motives as there are hackers, but we can break the most common intruder motivations into a few broad categories:

- **Recreation** Those who hack into networks “just for fun” or to prove their technical prowess; often young people or “anti-establishment” types.
- **Remuneration** People who invade the network for personal gain, such as those who attempt to transfer funds to their own bank accounts or erase records of their debts; “hackers for hire” who are paid by others to break into the network; corporate espionage is included in this category.
- **Revenge** Dissatisfied customers, disgruntled former employees, angry competitors, or people who have a personal grudge against someone in the organization.

The scope of damage and extent of the intrusion is often – although by no means always – tied to the intruder’s motivation.

Recreational Hackers

Recreational hackers are often teen hackers who do it primarily for the thrill of accomplishment. In many cases, they do little or no permanent damage, perhaps only leaving “I was here” type messages to “stake their claims” and prove to their peers that they were able to penetrate your network’s security.

There are more malevolent versions of the fun-seeking hacker, however. These are the cyber-vandals, who get their kicks out of destroying as much of your data as possible, or causing your systems to crash.

Profit-motivated Hackers

Those who break into your network for remuneration of some kind – either directly or indirectly – are more dangerous. Because money is at stake, they are more motivated to accomplish their objective. And because many of them are “professionals” of a sort, their hacking techniques may be more sophisticated than the average teenage recreational hacker.

Monetary motivations include:

- Personal financial gain
- Third-party payment
- Corporate espionage

Those motivated by the last are usually the most sophisticated and the most dangerous. There is often *big* money involved in theft of trade secrets. Corporate espionage agents may be employees who have been approached by your competitors and offered money or merchandise, or even threatened with blackmail or physical harm.

In some instances, those working for competitors will go “undercover” and seek a job with your company in order to steal data that they can take back to their own organizations (to add insult to injury, these “stealth spies” are getting paid by your company at the same time they’re working against you to the benefit of your competitor).

There are also “professional” freelance corporate spies. They may be contacted and contracted to obtain your company secrets, or they may do it on their own and auction it off to your competitors.

These corporate espionage agents are often highly skilled. They are technically savvy and intelligent enough to avoid being caught or detected. Fields that are especially vulnerable to the threat of corporate espionage include:

- Oil and energy
- Engineering
- Computer technology
- Research medicine
- Law

Any company that is on the verge of a breakthrough that could result in large monetary rewards or world-wide recognition, especially if the company’s involvement is high profile, should be aware of the possibility of espionage and take steps to guard against it.

Vengeful Hackers

Persons motivated by the desire for revenge are dangerous, as well. Vengeance seeking is usually based on strong emotions, which means these hackers may go all out in their efforts to sabotage your network.

Examples of hackers or security saboteurs acting out of revenge include:

- Former employees who are bitter about being fired or laid off or who quit their jobs under unpleasant circumstances
- Current employees who feel mistreated by the company, especially those who may be planning to leave soon
- Current employees who aim to sabotage the work of other employees due to internal political battles, rivalry over promotions, and the like
- Outsiders who have grudges against the company, such as those at competing companies who want to harm or embarrass the company or dissatisfied customers
- Outsiders who have personal grudges against someone who works for the company, such as former girlfriend/boyfriends, spouses going through a divorce, and other relationship-related problems

Luckily, the intruders in this category are generally less technically talented than those in the other two groups, and their emotional involvement may cause them to be careless and take outrageous chances, which makes them easier to catch.

Hybrid Hackers

Of course, the three categories can overlap in some cases. A recreational hacker who perceives himself to have been mistreated by an employer or in a personal relationship may use his otherwise benign hacking skills to impose “justice” for the wrongs done to him, or a vengeful ex-employee or ex-spouse might pay someone else to do the hacking for him.

It is beneficial to understand the common motivations of network intruders because, although we may not be able to predict which type of hacker will decide to attack our networks, we can recognize how each operates and take steps to protect our networks from all of them.

Even more important in planning our security strategy than the type of *hacker*, however, is the type of *attack*. In the next section, we will examine specific types of network attacks and how you can protect against them.

Classifying Specific Types of Attacks

The *attack type* refers to *how* an intruder gains entry to your computer or network and *what he does* once he has gained entry. In this section, we will discuss some of the more common types of hack attacks, including:

- Social engineering attacks
- Denial of Service (DOS) attacks
- Scanning and Spoofing
- Source routing and other protocol exploits
- Software and system exploits
- Trojans, viruses and worms

When you have a basic understanding of how each type of attack works, you will be better armed to guard against them.

Social engineering attacks

Unlike the other attack types, *social engineering* does not refer to a technological manipulation of computer hardware or software vulnerabilities and does not require much in the way of technical skills. Instead, this type of attack exploits *human* weaknesses – such as carelessness or the desire to be cooperative – to gain access to legitimate network credentials. The talents that are most useful to the intruder who relies on this technique are the so-called “people skills,” such as a charming or persuasive personality or a commanding, authoritative presence.

What is social engineering?

Social engineering is defined as *obtaining confidential information by means of human interaction* (Business Wire, August 4, 1998). You can think of social engineering attackers as specialized con artists. They

gain the trust of users (or even better, administrators) and then take advantage of the relationship to find out the user's account name and password, or have the unsuspecting users log them onto the system. Because it is based on convincing a valid network user to “open the door,” social engineering can successfully get an intruder into a network that is protected by high-security measures such as biometric scanners.

Social engineering is, in many cases, the easiest way to gain unauthorized access to a computer network. The Social Engineering Competition at a Defcon annual hackers' convention in Las Vegas attracted hundreds of attendants eager to practice their manipulative techniques. Even hackers who are famous for their technical abilities know that *people* make up the biggest security vulnerability on most networks. Kevin Mitnick, convicted computer crimes felon and celebrity hacker extraordinaire, tells in his lectures how he used social engineering to gain access to systems during his hacking career.

These “engineers” often pose as technical support personnel – either in-house, or pretending to work for outside entities such as the telephone company, the Internet Service provider, the network's hardware vendor, or even the government. They often contact their victims by phone, and they will usually spin a complex and plausible tale of why they need the users to divulge their passwords or other information (such as the IP address of the user's machine or the computer name of the network's authentication server).

Protecting your network against social engineers

It is especially challenging to protect against social engineering attacks. Adopting strongly worded policies that prohibit divulging passwords and other network information to anyone over the telephone and educating your users about the phenomenon are obvious steps you can take to reduce the likelihood of this type of security breach. Human nature being what it is, however, there will always be some users on every network who are vulnerable to the social engineer's con game. A talented social engineer is a master at making users doubt their own doubts about his legitimacy.

The “wannabe” intruder may regale the user with woeful stories of the extra cost the company will incur if he spends extra time verifying his identity. He may pose as a member of the company's top management and take a stern approach, threatening the employee with disciplinary action or even loss of job if he doesn't get the user's cooperation. Or he may try to make the employee feel guilty by pretending to be a low-level employee who is just trying to do his job and who will be fired if he doesn't get access to the network and get the problem taken care of right away. A really good social engineer is patient and thorough. He will do his homework, and will know enough about your company, or the organization he claims to represent, to be convincing.

Because social engineering is a human problem, not a technical problem, prevention must come primarily through education rather than technological solutions.

NOTE

For more information about social engineering and how to tell when someone is attempting to pull a social engineering scam, see the preview chapter entitled *Everything You Wanted to Know about Social Engineering – but were Afraid to Ask* at the “Happy Hacker” website, located at www.happyhacker.org/uberhacker/se.shtml.

Denial of Service (DOS) Attacks

Denial of Service (DOS) attacks are one of the most popular choices of Internet hackers who want to disrupt a network's operations. Although they do not destroy or steal data as some other types of attacks do, the objective of the DOS attacker is to bring down the network, denying service to its legitimate users. DOS attacks are easy to initiate; software is readily available from hacker websites and warez newsgroups that will allow anyone to launch a DOS attack with little or no technical expertise.

NOTE

Warez is a term used by hackers and crackers to describe bootlegged software that has been "cracked" to remove copy protections and made available by software pirates on the Internet, or in its broader definition, to describe any illegally distributed software.

In February of 2000, massive DOS attacks brought down several of the biggest websites, including Yahoo.com and Buy.com.

The purpose of a DOS attack is to render a network inaccessible by generating a type or amount of network traffic that will crash the servers, overwhelm the routers or otherwise prevent the network's devices from functioning properly. Denial of service can be accomplished by tying up the server's resources, for example, by overwhelming the CPU and memory resources. In other cases, a particular user/machine can be the target of denial of service attacks that hang up the client machine and require it to be rebooted.

NOTE

Denial of service attacks are sometimes referred to in the security community as "nuke attacks."

Distributed Denial of Service attacks

Distributed DOS (DDOS) attacks use intermediary computers called *agents* on which programs called *zombies* have previously been surreptitiously installed. The hacker activates these zombie programs remotely, causing the intermediary computers (which can number in the hundreds or even thousands) to simultaneously launch the actual attack. Because the attack comes from the computers running the

zombie programs, which may be on networks anywhere in the world, the hacker is able to conceal the true origin of the attack.

Examples of DDOS tools used by hackers are TFN (Tribe FloodNet), TFN2K, Trinoo, and Stacheldraht (German for “barbed wire”). While early versions of DDOS tools targeted UNIX and Solaris systems, TFN2K can run on both UNIX and Windows systems.

It is important to note that DDOS attacks pose a two-layer threat. Not only could your network be the target of a DOS attack that crashes your servers and prevents incoming and outgoing traffic, but your computers could be used as the “innocent middle men” to launch a DOS attack against another network or site.

DNS DOS attack

The Domain Name System (DNS) DOS attack exploits the difference in size between a DNS query and a DNS response, in which all of the network’s bandwidth is tied up by bogus DNS queries. The attacker uses the DNS servers as “amplifiers” to multiply the DNS traffic.

The attacker begins by sending small DNS queries to each DNS server, which contain the spoofed IP address (see *IP Spoofing* later in this chapter) of the intended victim. The responses returned to the small queries are much larger in size, so that if there are a large number of responses returned at the same time, the link will become congested and denial of service will take place.

One solution to this problem is for administrators to configure DNS servers to respond with a “refused” response, which is much smaller in size than a name resolution response, when they received DNS queries from suspicious or unexpected sources.

SYN attack/LAND attack

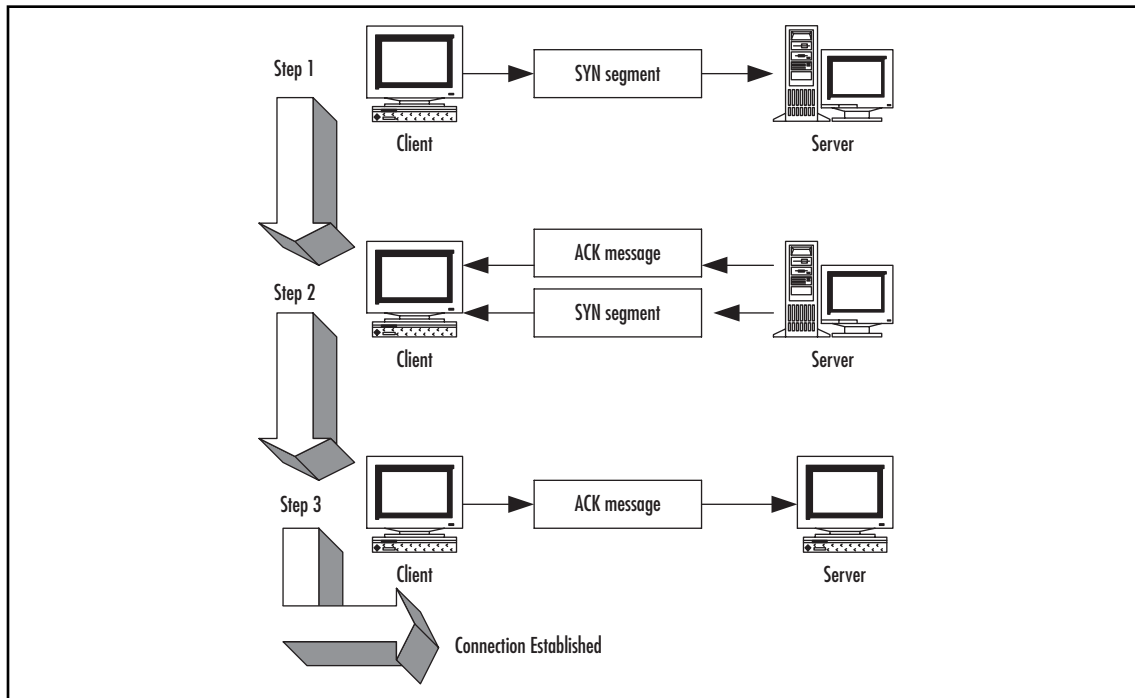
Synchronization request (SYN) attacks exploit the Transmission Control Protocol (TCP) “three-way handshake,” the process by which a communications session is established between two computers. Because TCP, unlike User Datagram Protocol (UDP), is connection-oriented, a *session*, or direct one-to-one communication link, must be created before sending data. The client computer initiates communication with the server (the computer whose resources it wants to access).

The “handshake” includes the following steps:

1. The client machine sends a SYN segment.
2. The server sends an acknowledgement (ACK) message and a SYN, which acknowledges the client machine’s request that was sent in step 1 and sends the client a synchronization request of its own. The client and server machines must synchronize each other’s sequence numbers.
3. The client sends an ACK back to the server, acknowledging the server’s request for synchronization. When both machines have acknowledged each other’s requests, the handshake has been successfully completed and a connection is established between the two computers.

Figure 1.4 illustrates how the client/server connection works.

Figure 1.4 TCP uses a “three-way handshake” to establish a connection between client and server



This is how the process normally works. A SYN attack uses this process to flood the system targeted with multiple SYN packets that have bad source IP addresses, which causes the system to respond with SYN/ACK messages. The problem comes when the system, waiting for the ACK message, puts the waiting SYN/ACK messages into a queue. The queue is limited in the number of messages it can handle, and when it is full, all subsequent incoming SYN packets will be ignored. In order for a SYN/ACK to be removed from the queue, an ACK must be returned from the client, or the interval timer must run out and terminate the three-way handshake process.

Because the source IP addresses for the SYN packets sent by the attacker are no good, the ACKs that the server is waiting for never come. The queue stays full, and there is no room for valid SYN requests to be processed. Thus service is denied to legitimate clients attempting to establish communications with the server.

The LAND attack is a variation on the SYN attack. In the LAND attack, instead of sending SYN packets with IP addresses that do not exist, the flood of SYN packets all have the same spoof IP address – that of the targeted computer.

The LAND attack can be prevented by filtering out incoming packets whose source IP addresses appear to be from computers on the internal network. ISA Server has preset intrusion detection functionality that allows you to detect attempted LAND attacks, and you can configure Alerts to notify you when such an attack is detected.

Ping of Death

Another type of DOS attack that ISA Server can be set to specifically detect is the so-called “Ping of Death” (also known as the “large packet ping”). The Ping of Death attack is launched by creating an IP packet larger than 65,536 bytes, which is the maximum allowed by the IP specification (this is sometimes referred to as a “killer packet”). This can cause the target system to crash, hang or reboot.

Although newer operating systems are generally not vulnerable to this type of attack, many companies still have older operating systems deployed against which the Ping of Death can be used.

ISA allows you to specifically enable detection of Ping of Death attacks.

Teardrop

The teardrop attack works a little differently from the Ping of Death, but with similar results. The teardrop program creates IP fragments, which are pieces of an IP packet into which an original packet can be divided as it travels through the Internet. The problem is that the offset fields on these fragments, which are supposed to indicate the portion (in bytes) of the original packet that is contained in the fragment, overlap.

For example, normally two fragments’ offset fields might appear as shown below:

```
Fragment 1: (offset) 100 - 300
```

```
Fragment 2: (offset) 301 - 600
```

This indicates that the first fragment contains bytes 100 through 300 of the original packet, and the second fragment contains bytes 301 through 600.

Overlapping offset fields would appear something like this:

```
Fragment 1: (offset) 100 - 300
```

```
Fragment 2: (offset) 200 - 400
```

When the destination computer tries to reassemble these packets, it is unable to do so and may crash, hang or reboot.

Variations include:

- NewTear
- Teardrop2
- SynDrop
- Boink

All of these programs generate some sort of fragment overlap.

Ping Flood (ICMP flood)

The ping flood or ICMP flood is a means of tying up a specific client machine. It is caused by an attacker sending a large number of ping packets (ICMP echo request packets) to the Winsock or dialer software. This prevents it from responding to server ping activity requests, which causes the server to eventually timeout the connection. A symptom of a ping flood is a huge amount of modem activity, as indicated by the modem lights. This is also referred to as a *ping storm*.

The *fraggle attack* is related to the ping storm. Using a spoofed IP address (which is the address of the targeted victim), an attacker sends ping packets to a subnet, causing all computers on the subnet to respond to the spoofed address and flood it with echo reply messages.

NOTE

During the Kosovo crisis, the fraggle attack was frequently used by pro-Serbian hackers against U.S. and NATO sites to overload them and bring them down.

You can use programs such as NetXray or other IP tracing software to record and display a log of the flood packets. Firewalls can be configured to block ping packets to prevent these attacks.

SMURF attack

The Smurf attack is a form of “brute force” attack that uses the same method as the ping flood, but directs the flood of ICMP echo request packets at the network’s router. The destination address of the ping packets is the broadcast address of the network, which causes the router to broadcast the packet to every computer on the network or segment. This can result in a very large amount of network traffic if there are many host computers, which can create congestion that causes a denial of service to legitimate users.

NOTE

The broadcast address is normally represented by all 1s in the host ID. This means, for example, that on class C network 192.168.1.0, the broadcast address would be 192.168.1.255 (255 in decimal represents 11111111 in binary), and in a class C network, the last or z octet represents the host ID. A message sent to the broadcast address is sent simultaneously to all hosts on the network.

In its most insidious form, the Smurf attacker spoofs the source IP address of a ping packet. Then both the network to which the packets are sent *and* the network of the spoofed source IP address will be overwhelmed with traffic. The network to which the spoofed source address belongs will be deluged with responses to the ping when all the hosts to which the ping was sent answer the echo request with an echo reply.

Smurf attacks can generally do more damage than other forms of DoS, such as SYN floods. The SYN flood affects only the ability of other computers to establish a TCP connection to the flooded server, but a Smurf attack can bring an entire ISP down for minutes or hours. This is because a single attacker can easily send 40–50 ping packets per second, even using a slow modem connection. Because each is broadcast to every computer on the destination network, that means the number of responses per second is 40–50 times the number of computers on the network – which could be hundreds or thousands. This is enough data to congest even a T-1 link.

One way to prevent a Smurf attack from using your network as the broadcast target is to turn off the capability to transmit broadcast traffic on the router. Most routers allow you to do this. To prevent

your network from being the victim of the spoofed IP address, you will need to configure your firewall to filter out incoming ping packets.

UDP bomb or UDP flood

An attacker can use the UDP and one of several services that echo packets upon receipt to create service-denying network congestion by generating a flood of UDP packets between two target systems. For example, the UDP chargen service on the first computer, which is a testing tool that generates a series of characters for every packet that it receives, sends packets to another system's UDP echo service, which echoes every character it receives. By exploiting these testing tools, an endless flow of echos go back and forth between the two systems, congesting the network. This is sometimes called a *UDP packet storm*.

In addition to port 7, the echo port, an attacker can use port 17, the quote of the day service (quotd) or the daytime service on port 13. These services will also echo packets they receive. UDP chargen is on port 19.

Disabling unnecessary UDP services on each computer (especially those mentioned above) or using a firewall to filter those ports/services, will protect you from this type of attack.

UDP Snork attack

The snork attack is similar to the UDP bomb. It uses a UDP frame that has a source port of either 7 (echo) or 9 (chargen), with a destination port of 135 (Microsoft location service). The result is the same as the UDP bomb – a flood of unnecessary transmissions that can slow performance or crash the systems that are involved.

WinNuke (Windows out-of-band attack)

The out-of-band (OOB) attack is one that exploits a vulnerability in Microsoft networks, which is sometimes called the *Windows OOB bug*. The WinNuke program (and variations such as Sinnerz and Muerte) creates an out-of-band data transmission that crashes the machine to which it is sent. It works like this: a TCP/IP connection is established with the target IP address, using port 139 (the NetBIOS port). Then the program sends data using a flag called MSG_OOB (or Urgent) in the packet header. This flag instructs the computer's Winsock to send data called out-of-band data. Upon receipt, the targeted Windows server expects a pointer to the position in the packet where the Urgent data ends, with normal data following, but the OOB pointer in the packet created by WinNuke points to the end of the frame with no data following.

The Windows machine does not know how to handle this situation and will cease communicating on the network, and service will be denied to any users who subsequently attempt to communicate with it. A WinNuke attack usually requires a reboot of the affected system to reestablish network communications.

Windows 95 and NT 3.51 and 4.0 are vulnerable to the WinNuke exploit, unless the fixes provided by Microsoft have been installed. Windows 98/ME and Windows 2000 are not vulnerable to WinNuke, but ISA server allows you to enable detection of attempted OOB attacks.

Mail bomb attack

A mail bomb is a means of overwhelming a mail server, causing it to stop functioning and thus denying service to users. A mail is a relatively simple form of attack, accomplished by sending a

massive quantity of email to a specific user or system. There are programs available on hacking sites on the Internet that allow a user to easily launch a mail bomb attack, automatically sending floods of email to a specified address while protecting the attacker's identity.

A variation on the mail bomb program automatically subscribes a targeted user to hundreds or thousands of high volume Internet mailing lists, which will fill the user's mailbox and/or the mail server. Bombers call this *list linking*. Examples of these mail bomb programs include Unabomber, extreme Mail, Avalanche, and Kaboom.

The solution to repeated mail bomb attacks is to block traffic from the originating network using packet filters. Unfortunately, this does not work with list linking because the originator's address is obscured; the deluge of traffic comes from the mailing lists to which the victim has been subscribed.

Scanning and Spoofing

The term *scanner*, in the context of network security, refers to a software program that is used by hackers to remotely determine what TCP/UDP ports are open on a given system, and thus vulnerable to attack. Administrators also use scanners to detect and correct vulnerabilities in their own systems before an intruder finds them. Network diagnostic tools such as the famous Security Administrator's Tool for Analyzing Networks (SATAN), a UNIX utility, include sophisticated port scanning capabilities.

A good scanning program can locate a target computer on the Internet (one that is vulnerable to attack), determine what TCP/IP services are running on the machine, and probe those services for security weaknesses.

NOTE

A common saying among hackers is: *a good port scanner is worth a thousand passwords.*

Many scanning programs are available as freeware on the Internet.

Port scan

Port scanning refers to a means of locating “listening” TCP or UDP ports on a computer or router and obtaining as much information as possible about the device from the listening ports. TCP and UDP services and applications use a number of *well-known ports*, which are widely published. The hacker uses his knowledge of these commonly used ports to extrapolate information.

For example, Telnet normally uses port 23. If the hacker finds that port open and listening, he knows that Telnet is probably enabled on the machine. He can then try to infiltrate the system, for example by guessing the appropriate password in a brute force attack.

Back to Basics: TCP/UDP Well Known Ports

The official well-known port assignments are documented in RFC 1700, available on the web at www.freesoft.org/CIE/RFC/1700/index.htm. The port assignments are made by the Internet Assigned Numbers Authority (IANA). In general, a service will use the same port number with UDP as with TCP, although there are some exceptions. The assigned ports were originally those from 0–255, but the number was later expanded to 0–1023.

Some of the most used well-known ports include:

- TCP/UDP port 20: FTP (data)
- TCP/UDP port 21: FTP (control)
- TCP/UDP port 23: Telnet
- TCP/UDP port 25: SMTP
- TCP/UDP port 53: DNS
- TCP/UDP port 67: BOOTP server
- TCP/UDP port 68: BOOTP client
- TCP/UDP port 69: TFTP
- TCP/UDP port 80: HTTP
- TCP/UDP port 88: Kerberos
- TCP/UDP port 110: POP3
- TCP/UDP port 119: NNTP
- TCP/UDP port 137: NetBIOS name service
- TCP/UDP port 138: NetBIOS datagram service
- TCP/UDP port 139: NetBIOS session service
- TCP/UDP port 194: IRC
- TCP/UDP port 220: IMAPv3
- TCP/UDP port 389: LDAP

Ports 1024–65,535 are called *registered ports*; these numbers are not controlled by IANA and can be used by user processes or applications.

There are a total of 65,535 TCP ports (and the same number of UDP ports) used for various services and applications. If a port is open, it will respond when another computer attempts to contact it over the network. Port scanning programs such as *Nmap* are used to determine which ports

are open on a particular machine. The program sends packets for a wide variety of protocols and, by examining which messages receive responses and which don't, creates a map of the computer's listening ports.

Port scanning in itself does no harm to your network or system, but it provides hackers with information they can use to penetrate the network.

IP half scan attack

“Half scans” (also called “half open scans” or FIN scans) attempt to avoid detection by sending only initial or final packets, rather than establishing a connection. A half scan starts the SYN/ACK process with a targeted computer, but does not complete it. Software that conducts half scans, such as Jaka!, is called a *stealth scanner*.

Many port scanning detectors are unable to detect half scans; however, ISA Server provides IP half scan as part of its intrusion detection.

IP Spoofing

IP spoofing involves changing the packet headers of a message to indicate that it came from an IP address other than the true source. The spoofed address is normally a trusted port, which allows a hacker to get a message through a firewall or router that would otherwise be filtered out. Modern firewalls protect against IP spoofing.

Spoofing is used whenever it is beneficial for one machine to impersonate another. It is often used in combination with one of the other types of attacks. For example, a spoofed address is used in the SYN flood attack to create a “half open” connection, in which the client never responds to the SYN/ACK message because the spoofed address is that of a computer that is down or doesn't exist. Spoofing is also used to hide the true IP address of the attacker in Ping of Death, Teardrop and other attacks.

IP spoofing can be prevented by using Source Address Verification on your router, if it is supported.

Source Routing attack

TCP/IP supports *source routing*, a means that permits the sender of network data to route packets through a specific point on the network. There are two types of source routing:

- *Strict source routing*: the sender of the data can specify the exact route (rarely used).
- *Loose source record route (LSRR)*: the sender can specify certain routers (hops) through which the packet must pass.

The source route is an option in the IP header that allows a sender to override routing decisions normally made by routers between the source and destination machines. Source routing is used by network administrators to map the network, or for troubleshooting routing and communications problems. It can also be used to force traffic through the route that will provide the best performance. Unfortunately, source routing can be exploited by hackers.

If the system allows source routing, an intruder can use it to reach private internal addresses on the LAN that normally would not be reachable from the Internet, by routing the traffic through another machine that is reachable from both the Internet and the internal machine.

Source routing can be disabled on most routers to prevent this type of attack.

Other protocol exploits

The attacks we have discussed so far involve exploiting some feature or weakness of the TCP/IP protocols. Hackers can also exploit vulnerabilities of other common protocols, such as Hypertext Transfer Protocol (HTTP), Domain Name System (DNS), Common Gateway Interface (CGI), and other commonly used protocols.

Active-X controls, Java script, and VBscript can be used to add animations or applets to web sites, but hackers can exploit these to write controls or scripts that allow them to remotely plant viruses, access data, or change or delete files on the hard disk of unaware users who visit the page and run the script. Many e-mail client programs have similar vulnerabilities.

System and software exploits

System and software exploits are those that take advantage of weaknesses of particular operating systems and applications (often called *bugs*). Like protocol exploits, they are used by intruders to gain unauthorized access to computers or networks or to crash or clog up the systems to deny service to others.

Common “bugs” can be categorized as follows:

- **Buffer overflows** Many common security holes are based on buffer overflow problems. Buffer overflows occur when the number of bytes or characters input exceeds the maximum number allowed by the programmer in writing the program.
- **Unexpected input** Programmers may not take steps to define what happens if invalid input (input that doesn’t match program specifications) is entered. This could cause the program to crash or open up a way into the system.
- **System configuration bugs** These are not really “bugs,” per se, but rather are ways of configuring the operating system or software that leaves it vulnerable to penetration.

Popular software such as Microsoft’s Internet Information Server (IIS), Internet Explorer (MSIE) and Outlook Express (MSOE) are popular targets of hackers looking for software security holes that can be exploited.

Major operating system and software vendors regularly release security patches to fix exploitable bugs. It is very important for network administrators to stay up to date in applying these fixes and/or service packs to ensure that their systems are as secure as possible.

NOTE

Microsoft issues *security bulletins* and makes security patches available as part of TechNet. See the website at www.microsoft.com/technet/security/default.asp.

Trojans, viruses and worms

Intruders who access your systems without authorization or inside attackers with malicious motives may plant various types of programs to cause damage to your network. There are three broad categories of *malicious code*, as follows:

- Trojans
- Viruses
- Worms

We will take a brief look at each of these attack types.

Trojans

The name is short for “Trojan horse,” and refers to a software program that appears to perform a useful function, but in fact, performs actions that the user of the program did not intend or was not aware of. Trojan horses are often written by hackers to circumvent the security of a system. Once installed, the hacker can exploit the security holes created by the Trojan to gain unauthorized access, or the Trojan program may perform some action such as:

- Deleting or modifying files
- Transmitting files across the network to the intruder
- Installing other programs or viruses

Basically, the Trojan can perform any action that the user has privileges and permissions to do on the system. This means a Trojan is especially dangerous if the unsuspecting user who installs it is an administrator and has access to the system files.

Trojans can be very cleverly disguised as innocuous programs, such as utilities or screensavers. A Trojan can also be installed by an executable script (Javascript, a Java applet, Active-X control, and others) on a web site. Accessing the site may initiate the installation of the program if the web browser is configured to allow scripts to run automatically.

Viruses

The most common use of the term “virus” is any program that is installed without the awareness of the user and performs undesired actions (often harmful, although sometimes merely annoying). Viruses may also replicate themselves, infecting other systems by writing themselves to any floppy disk that is used in the computer or sending themselves across the network. Viruses are often distributed as attachments to e-mail, or as macros in word processing documents. Some activate immediately upon installation, and others lie dormant until a specific date/time or a particular system event triggers them.

Viruses come in thousands of different varieties. They can do anything from popping up a message that says “Hi!” to erasing the computer’s entire hard disk. The proliferation of computer viruses has also led to the phenomenon of the *virus hoax*, which is a warning – generally circulated via email or websites – about a virus that does not exist or that does not do what the warning claims it will do.

Viruses, however, present a real threat to your network. Companies such as Symantec and McAfee make anti-virus software that is aimed at detecting and removing virus programs. Because new viruses are being created daily, it is important to download new *virus definition files*, which contain information required to detect each virus type, on a regular basis to ensure that your virus protection stays up to date.

Worms

A worm is a program that can travel across the network from one computer to another. Sometimes different parts of a worm run on different computers. Technically, a worm – unlike a virus – can replicate itself without user interaction; however, much modern documentation makes little distinction between the two, or classifies the worm as a subtype of the virus. Worms make multiple copies of themselves and spread throughout a network. Originally the term *worm* was used to describe code that attacked multiuser systems (networks) while *virus* was used to describe programs that replicated on individual computers.

The primary purpose of the worm is to replicate. These programs were initially used for legitimate purposes in performing network management duties, but their ability to multiply quickly has been exploited by hackers who create malicious worms that replicate wildly, and may also exploit operating system weaknesses and perform other harmful actions.

Designing a Comprehensive Security Plan

Now that you have some understanding of basic security concepts and terminology, general security objectives, common motivation of network intruders, different types of specific attacks and how they are used, and an overview of available hardware and software solutions, you can begin to design a comprehensive security policy for your organization.

A widely accepted method for developing your network security plan is laid out in Request for Comments (RFC) 2196, *Site Security Handbook*, and attributed to Fites, et al (1989). It consists of the following steps:

- Identify what you are trying to protect.
- Determine what you are trying to protect it from.
- Determine how likely the anticipated threats are.
- Implement measures that will protect your assets in a cost-effective manner.
- Review the process continually and make improvements each time a weakness is discovered.

NOTE

The entire text of RFC 2196, which provides many excellent suggestions that focus primarily on the implementation phase, can be found on the web at www.faqs.org/rfcs/rfc2196.html.

It is important to understand that a security *plan* is not the same thing as a security *policy*, although the two words are sometimes used interchangeably. Your security policies (and there are likely to be many of them) grow out of the security plan. Think of policy as “law” or “rules,” while the security plan is procedural; it lays out *how* the rules will be implemented.

Your security plan will generally address three different aspects of protecting your network:

1. *Prevention*: the measures that are implemented to keep your information from being modified, destroyed, or compromised.
2. *Detection*: the measures that are implemented to recognize when a security breach has occurred or has been attempted, and if possible, the origin of the breach.
3. *Reaction*: the measures that are implemented to recover from a security breach, to recover lost or altered data, to restore system or network operations, and to prevent future occurrences.

These can be divided into two types of actions: *proactive* and *reactive*. The first, prevention, is proactive because it takes place *before* any breach has occurred and involves actions that will, if successful, make further actions unnecessary. Unfortunately, our proactive measures don't always work. Reactive measures such as detection and reaction do, however, help us to develop additional proactive measures that will prevent future intrusions.

Regardless of how good your prevention and detection methods may be, it is essential that you have in place a reaction in case attackers do get through and damage your data or disrupt your network operations. As the old folk saying goes: “hope for the best, and plan for the worst.”

Evaluating Security Needs

Before you can develop a security plan and policies for your organization, you must assess the security needs, which will generally be based on the following broad considerations:

- Type of business in which the organization engages
- Type of data that is stored on the network
- Type of connection(s) that the network has to other networks
- Philosophy of the organization's management

Each of these will play a part in determining the level of security that is desirable or necessary for your network.

Assessing the type of business

Certain fields have inherent high-security requirements. An obvious example is the military, or other government agencies that deal with defense or national security issues. Private companies with government defense contracts also fall into this category. Others may be less obvious:

- Law firms are bound by law and ethics to protect client confidentiality.
- Medical offices must protect patient records.
- Law enforcement agencies, courts, and other governmental bodies must secure information.

- Educational institutions store student records.
- Companies that gather information from individuals or organizations guarantee that the data will be kept confidential.

The competitive nature of the business is also a consideration. In a field such as biogenetic research, which is a “hot” market where new developments are being made on a daily basis, any of which could involve huge profits for the company that patents the idea, protecting trade secrets becomes vitally important.

Most businesses will have *some* data of a confidential nature on the network’s computer systems, but the security requirements in some fields are much higher than others. This should be considered as you begin to develop your security plan.

Assessing the type of data

The second question to consider is what type of data is stored on your network, and where. You may find that a higher level of security is needed in one department or division than another. You may, in fact, want to divide the network physically, into separate subnets, to allow you to better control access to different parts of the company network independently.

Generally, payroll and human resource records (such as personnel files and insurance claim documents), company financial records (accounting documents, financial statements, tax documents), and a variety of other common business records will need to be protected. Even in cases where these documents are required to be made public, you will want to take steps to ensure that they can’t be modified or destroyed. Remember that *data integrity*, as well as *data confidentiality and availability*, is protected by a good security plan.

Assessing the network connections

Your exposure to outside intruders is another consideration in planning how to implement security on your network. A LAN that is self-contained and has no Internet connectivity, nor any modems or other outside connections, will not require the degree of protection (other than physical security) that is necessary when there are many avenues “in” that an intruder can take.

Dialup modem connections merit special consideration. While a dialup connection is less open to intrusion than a fulltime dedicated connection – both because it is connected to the outside for a shorter time period, reducing the window of opportunity for intrusion, and because it will usually have a dynamic IP address, making it harder for an intruder to locate it on multiple occasions – allowing workstations on your network to have modems and phone lines can create a huge security risk.

If improperly configured, a computer with a dialup connection to the Internet that is also cabled to the internal network can act as a router, allowing outside intruders to access not just the workstation connected to the modem, but other computers on the LAN.

One reason for allowing modems at individual workstations is to allow users to dialup connections to other private networks. A more secure way to do this is to remove the modems and have the users establish a virtual private networking (VPN) connection with the other private network through the LAN’s Internet connection.

The best security policy is to have as few connections from the internal network to the outside as possible, and control access at those entry points (called the *network perimeter*).

Assessing management philosophy

This last criteria is the most subjective, but can have a tremendous influence on the security level that is appropriate for your organization. Most companies are based on one (or a combination of more than one) management model.

Understanding management models

Some companies institute a highly structured, formal management style. Employees are expected to respect a strict chain of command, and information is generally disseminated on a “need to know” basis. Governmental agencies, especially those that are law-enforcement-related, such as police departments and investigative agencies, often follow this philosophy. This is sometimes referred to as the paramilitary model.

Other companies, particularly those in the IT industry and other fields that are subject to little state regulation, are built on the opposite premise: that all employees should have as much information and input as possible, that managers should function as “team leaders” rather than authoritarian supervisors, and that restrictions on employee actions should be imposed only when necessary for the efficiency and productivity of the organization. This is sometimes called the “one big happy family” model. Creativity is valued more than “going by the book,” and job satisfaction is considered an important aspect of enhancing employee performance and productivity.

In business management circles, these two diametrically-opposed models are called Theory X (traditional paramilitary style) and Theory Y (modern, team-oriented approach). Although there are numerous other management models that have been popularized in recent years, such as Management by Objective (MBO) and Total Quality Management (TQM), each company’s management style will fall somewhere on the continuum between Theory X and Theory Y. The management model is based on the personal philosophies of the company’s top decision-makers regarding the relationship between management and employees.

The management model can have a profound influence on what is or isn’t acceptable in planning security for the network. A “deny all access” security policy that is viewed as appropriate in a Theory X organization may meet with so much resentment and employee dissatisfaction in a Theory Y company that it disrupts business operations. Always consider the company “atmosphere” as part of your security planning. If you have good reasons to implement strict security in a Theory Y atmosphere, realize that you will probably have to justify the restrictions to management and “sell” them to employees, whereas those same restrictions might be accepted without question in a more traditional organization.

Understanding Security Ratings

Security ratings may be of interest as you develop your company’s security policy, although they are not likely to be important unless your organization works under government contract requiring a specified level of security.

The U.S. Government provides specifications for the rating of network security implementations in a publication often referred to as the *orange book*, formally called the *Department of Defense Trusted Computer System Evaluation Criteria*, or TCSEC. The *red book*, or *Trusted Network Interpretation of the TCSEC (TNI)* explains how the TCSEC evaluation criteria are applied to computer networks.

Other countries have security rating systems that work in a similar way. For example:

- CTPEC (Canada)
- AISEP (Australia)
- ITSEC (Western Europe)

To obtain a government contract in the U.S., companies are often required to obtain a C2 rating. A C2 rating has several requirements:

1. That the operating system in use be capable of tracking access to data, including both who accessed it and when it was accessed (as is done by the auditing function of Windows NT/2000)
2. That users' access to objects be subject to control (access permissions)
3. That users are uniquely identified on the system (user account name and password)
4. That security-related events can be tracked and permanently recorded for auditing (audit log)

If your organization needs a C2 rating for its systems, you should consult the National Computer Security Center (NCSC) publications to ensure that it meets all of the requirements.

Legal Considerations

Another important step in preparing to design your network security plan is to consider legal aspects that may affect your network. It is a good idea to have a member of your company's legal department who specializes in computer law to be involved in the development of your security plan and policies. If this is not possible, the written policies should be submitted for legal review before you put them into practice.

Designating Responsibility for Network Security

In any undertaking as complex as the development and implementation of a comprehensive corporate security plan and accompanying policies, it is vital that areas of responsibility be clearly designated.

Best practices dictate that no one person should have complete authority or control, and in an enterprise-level network, it would be difficult for any single person to handle all facets of developing and implementing the security plan anyway.

Responsibility for Developing the Security Plan and Policies

The initial creation of a good security plan will require a great deal of thought and effort. The policy will impact those at all levels of the organization, and soliciting input from as many representatives of different departments and job descriptions as is practical is desirable. An effective approach is to form a committee consisting of persons from several areas of the organization to be involved in creating and reviewing the security plan and policies.

The Security Planning Committee might include some or all of the following:

1. The network administrator and one or more assistant administrators

2. The site's security administrator
3. Department heads of various company departments or their representatives
4. Representatives of user groups that will be impacted by the security policies (for example, the secretarial staff or the data processing center)
5. A member of the legal department who specializes in computer and technology law
6. A member of the finance or budget department

Responsibility for Implementing and Enforcing the Security Plan and Policies

Security policies will generally be implemented and enforced by network administrators and members of the IT staff. Job descriptions and policies should designate exactly who is responsible for the implementation of which parts of the plan. There should be a clear-cut chain of command that specifies whose decision prevails in case of conflict.

In some cases – such as physical penetration of the network – the company security staff will become involved. There should be written, clearly formulated policies that stipulate which department has responsibility for which tasks in such situations.

The security plan should also address the procedures for reporting security breaches, both internally, and if the police or other outside agencies are to be brought in (as well as who is responsible for or has the authority to call in outside agents).

One of the most important factors in a good security policy is that it must be enforceable, and going a step further, it must be enforced. This is important for legal as well as practical reasons. If your company has policies in place that they routinely fail to enforce, this can be seen as an informal voiding of the policy, leaving the company legally liable for the actions of employees who violate the policy. If the policy can be enforced through technological means, this is preferred. If the policies must be enforced through reprimand or other actions against employees who violate them, there should be clearly worded, universally distributed written documentation of what constitutes a violation and what sanctions will result, as well as who is responsible for imposing such sanctions.

Designing the Corporate Security Policy

Designing a good corporate network security policy will differ, depending on the particular organization. However, there are common elements that should be addressed, including (but not limited to) the following:

- Developing an effective password and authentication policy
- Developing a privacy policy that sets forth reasonable expectations of privacy as to employees' e-mail, monitoring access to Web sites, access to users' directories and files, and so forth
- Developing an accountability policy that defines responsibility concerning security issues, including policies regarding users' obligation to report security violations and the process for doing so
- A network use statement that defines users' responsibilities in regard to accessing network resources, protection of password confidentiality, reporting of problems, and expectations as to availability of network resources

- A disaster protection and recovery policy that specifies policies for fault tolerance, scheduling of data backups and storage of backed-up data, failover plans for critical systems, and other related matters


It is beyond the scope of this chapter to provide detailed examples of all of the above. We will, however, address the first issue: how to go about developing an effective password policy and some of the factors that should be considered. The other policy areas should be addressed in similar depth and detail.

Developing an Effective Password Policy

In the networking world, passwords (in combination with user account names) are normally the “keys to the kingdom” that provide access to network resources and data. It may seem simplistic to say that your comprehensive security plan should include an effective password policy, but it is a basic component that is more difficult to implement than it might appear at first glance.

In order to be effective, your password policy must require users to select passwords that are difficult to “crack” – yet easy for them to remember so they don’t commit the common security breach of writing the password on a sticky note that will end up stuck to the monitor or sitting prominently in the top desk drawer.

A good password policy is the first line of defense in protecting your network from intruders. Careless password practices (choosing common passwords such as “god” or “love” or the user’s spouse’s name; choosing short, all-alpha, one-case passwords, writing passwords down or sending them across the network in plain text) are like leaving your car doors unlocked with the keys in the ignition. Although some intruders may be targeting a specific system, many others are just “browsing” for a network that’s easy to break into. Lack of a good password policy is an open invitation to them.



TIP

Expensive, sophisticated firewalls and other strict security measures (short of biometric scanning devices that recognize fingerprints or retinal images) will not protect you if an intruder has knowledge of a valid user name and password. It is particularly important to use strong passwords for administrative accounts.

Best practices for password creation require that you address the following:

- Password length and complexity
- Who creates the password?
- Forced changing of passwords

Let’s discuss each of these considerations.

Password Length and Complexity

It’s easy to define a “bad” password – it’s one that can be easily guessed by someone other than the authorized user.

One way in which “crackers” (hackers who specialize in defeating passwords to break into systems) do their work is called the *brute force* attack. In this kind of attack, the cracker manually, or more often, using a script or specially written software program, simply tries every possible combination of characters until he finally hits upon the right one. It goes without saying that using this method, it will be easier to guess a short password than a longer one; there are more possible combinations. For this reason, most security experts recommend that passwords have a minimum required length (for example, eight characters). Modern network operating systems such as Windows 2000 allow domain administrators to impose such rules so that if a user attempts to set a password that doesn’t meet the minimum length requirement, the password change will be rejected.

Who creates the password?

Network administrators may be tempted to institute a policy whereby they create all passwords and “issue” them to the users. This has the advantage of ensuring that all passwords will meet the administrator’s criteria in regard to length and complexity. However, it has a few big disadvantages as well:

1. This places a heavy burden on administrators who must handle all password changes and be responsible for letting users know what their passwords are. Of course, you would not want to notify the user of his/her password via e-mail or other insecure channels. In fact, the best way to do so is to personally deliver the password information. In a large organization, this becomes particularly taxing if you have a policy requiring that passwords be changed on a regular basis (as you should; we will discuss this in the next section).
2. Users will have more difficulty remembering passwords that they didn’t choose. This means they are more likely to write the passwords down, resulting in security compromises. Otherwise, they may have to contact the administrator frequently to be reminded of their passwords.
3. If the administrator creates all passwords, this means the administrator *knows* everyone’s password. This may or may not be acceptable under your overall security policy. Some users (including management) may be uncomfortable with the idea that the administrator knows their passwords. Even though an administrator can generally access a user’s account and/or files without knowing the password, it is less obvious to the users, and thus, less of a concern.

Allowing users to create their own passwords, within set parameters (length and complexity requirements) is usually the best option. The user is less likely to forget the password because he can create a complex password that is meaningless to anyone else, but which has meaning to him.

For example, it would be difficult for others to guess the password “Mft2doSmis.” It has 10 characters, combines alpha and numeric characters, and combines upper and lower case in a seemingly random manner. To the user, it would be easy to remember because it means, “My favorite thing to do on Sunday morning is sleep.”

Password Change Policy

Best practices dictate that users change their passwords at regular intervals, and after any suspected security breach. Windows 2000 allows the administrator to set a maximum password age, forcing users to change their passwords at the end of the specified period (in days). Password expiration periods can be set from 1 to 999 days. The default is 42 days.

NOTE

Individual user accounts that need to keep the same passwords can be configured so that their passwords never expire. This overrides the general password expiration setting.

Because it is the nature of most users to make their passwords as easy to remember as possible, you must institute policies to prevent the following practices, all of which can present security risks:

- Changing the password to a variation of the same password (for example, changing from Tag2mB to Tag3mB)
- Changing the password back and forth between two favored passwords each time a change is required (that is, changing from Tag2mB to VERoh9 and back again continuously)
- “Changing” the password to the same password (entering the same password for the new password as was already being used)

Administrators can use operating system features or third party software to prevent most of these practices. For example, in Windows 2000, you can configure the operating system to remember the user’s password history, so that up to a maximum of the last 24 passwords will be recorded, and the user will not be able to change the password to one that has been used during that time.

Summary of Best Password Practices

- Passwords should have a minimum of eight characters.
- Passwords should not be “dictionary” words.
- Passwords should consist of a mixture of alpha, numeric and symbol characters.
- Passwords should be created by their users.
- Passwords should be easy for users to remember.
- Passwords should never be written down.
- Passwords should be changed on a regular basis.
- Passwords should be changed anytime compromise is suspected.
- Password change policies should prevent users from making only slight changes.

Educating Network Users on Security Issues

The best security policies in the world will be ineffective if the network users are unaware of them, or if the policies are so restrictive and place so many inconveniences on users that they go out of their ways to attempt to circumvent them.

The security plan itself should contain a program for educating network users – not just as to what the policies are, but *why* they are important, and how the users benefit from them. Users should also be instructed in the best ways to comply with the policies, and what to do if they are unable to comply or observe a deliberate violation of the policies on the part of other users.

If you involve users in the planning and policy-making stages, you will find it must easier to educate them and gain their support for the policies at the implementation and enforcement stages.

Summary

To get the most out of ISA's features, you must be able to recognize the security threats to which your network is subject and understand a little about the motivations of typical intruders. It is not necessary that you *be* a hacker in order to prevent your network from hacking attempts, but it *will* benefit you to know something about how unscrupulous hackers think and how they do their dirty work.

You must be aware of the different types of attacks with which you could be confronted, and understand how to protect your network from social engineering attacks, DoS attacks, scanning and spoofing, source routing and other protocol exploits, software and system exploits, and Trojans, viruses and worms.

There are a number of hardware-based security solutions available, and even more software-based firewalls on the market. You should have a basic understanding of the capabilities and limitations of each type, and how ISA Server compares – in features and cost – to some of the others. We think you will find that ISA Server offers an excellent value in comparison to competitive products, along with easy configurability and options to integrate third-party programs for even more functionality.

Your comprehensive security plan is integral to protecting your network from both internal and external threats. There is no “one size fits all” when it comes to corporate security plans and policies; yours should be based on the nature of the business in which your organization engages, the nature of the data stored on the network, the number and types of connections your network has to the “outside world,” and the management philosophy regarding organizational structure.

A good security plan is one that meets the needs of IT administration, company management, and network users. The best way to ensure that your security plan meets these criteria is to involve persons from all levels of the organization in the planning process. Once you have a good, comprehensive security plan and corresponding policies worked out, you will be able to use ISA Server as an important element in your security plan, to implement and enforce those policies and provide monitoring, notification, and record-keeping to document the successful functioning of your security plan.

This page intentionally left blank