Security Technologies

With the rapid growth of interest in the Internet, network security has become a major concern for companies throughout the world. The fact that the information and tools needed to penetrate the security of corporate networks are widely available has increased that concern.

Because of this increased focus on network security, network administrators often spend more effort protecting their networks than on actual network setup and administration. New tools that probe for system vulnerabilities, such as the Security Administrator Tool for Analyzing Networks (SATAN), assist in these efforts, but these tools only point out areas of weakness instead of providing a means to protect networks. Thus, as a network administrator, you must constantly try to keep abreast of the large number of security issues confronting you in today's world. This chapter describes many of the security issues that arise when connecting a private network to the Internet.

Security Issues When Connecting to the Internet

When you connect your private network to the Internet, you are physically connecting your network to more than 50,000 unknown networks and all their users. Although such connections open the door to many useful applications and provide great opportunities for information sharing, most private networks contain some information that should not be shared with outside users on the Internet. In addition, not all Internet users are involved in lawful activities. These two statements foreshadow the key questions behind most security issues on the Internet:

- How do you protect confidential information from those who do not explicitly need to access it?
- How do you protect your network and its resources from malicious users and accidents that originate outside your network?

Protecting Confidential Information

Confidential information can reside in two states on a network. It can reside on physical storage media, such as a hard drive or memory, or it can reside in transit across the physical network wire in the form of packets. These two information states present multiple opportunities for attacks from users on your internal network, as well as those users on the Internet. We are primarily concerned with the second state, which involves network security issues. The following are five common methods of attack that present opportunities to compromise the information on your network:

- Network packet sniffers
- IP spoofing
- Password attacks
- Distribution of sensitive internal information to external sources

- Man-in-the-middle attacks

When protecting your information from these attacks, your concern is to prevent the theft, destruction, corruption, and introduction of information that can cause irreparable damage to sensitive and confidential information. This section describes these common methods of attack and provides examples of how your information can be compromised.

Network Packet Sniffers

Because networked computers communicate serially (one information piece is sent after another), large information pieces are broken into smaller pieces. (The information stream would be broken into smaller pieces even if networks communicated in parallel. The overriding reason for breaking streams into network packets is that computers have limited intermediate buffers.) These smaller pieces are called *network packets*. Several network applications distribute network packets in *clear text*; that is, the information sent across the network is not encrypted. (Encryption is the transformation, or scrambling, of a message into an unreadable format by using a mathematical algorithm.) Because the network packets are not encrypted, they can be processed and understood by any application that can pick them up off the network and process them.

A network protocol specifies how packets are identified and labeled, which enables a computer to determine whether a packet is intended for it. Because the specifications for network protocols, such as TCP/IP, are widely published, a third party can easily interpret the network packets and develop a packet sniffer. (The real threat today results from the numerous freeware and shareware packet sniffers that are available, which do not require the user to understand anything about the underlying protocols.) A *packet sniffer* is a software application that uses a network adapter card in promiscuous mode (a mode in which the network adapter card sends all packets received on the physical network wire to an application for processing) to capture all network packets that are sent across a local-area network.

Because several network applications distribute network packets in clear text, a packet sniffer can provide its user with meaningful and often sensitive information, such as user account names and passwords. If you use networked databases, a packet sniffer can provide an attacker with information that is queried from the database, as well as the user account names and passwords used to access the database. One serious problem with acquiring user account names and passwords is that users often reuse their login names and passwords across multiple applications.

In addition, many network administrators use packet sniffers to diagnose and fix network-related problems. Because in the course of their usual and necessary duties these network administrators (such as those in the Payroll Department) work during regular employee hours, they can potentially examine sensitive information distributed across the network.

Many users employ a single password for access to all accounts and applications. If an application is run in client/server mode and authentication information is sent across the network in clear text, then it is likely that this same authentication information can be used to gain access to other corporate resources. Because attackers know and use human characteristics (attack methods known collectively as *social engineering attacks*), such as using a single password for multiple accounts, they are often successful in gaining access to sensitive information.

IP Spoofing

An *IP spoofing attack* occurs when an attacker outside your network pretends to be a trusted computer either by using an IP address that is within the range of IP addresses for your network or by using an authorized external IP address that you trust and to which you wish to provide access to specified resources on your network.

Normally, an IP spoofing attack is limited to the injection of data or commands into an existing stream of data passed between a client and server application or a peer-to-peer network connection. To enable bidirectional communication, the attacker must change all routing tables to point to the spoofed IP address. Another approach the attacker could take is to simply not worry about receiving any response from the applications. If an attacker is attempting to get a system to mail him or her a sensitive file, application responses are unimportant.

However, if an attacker manages to change the routing tables to point to the spoofed IP address, he can receive all the network packets that are addressed to the spoofed address and reply just as any trusted user can. Like packet sniffers, IP spoofing is not restricted to people who are external to the network.

Password Attacks

Password attacks can be implemented using several different methods, including brute-force attacks, Trojan horse programs (discussed later in the chapter), IP spoofing, and packet sniffers. Although packet sniffers and IP spoofing can yield user accounts and passwords, password attacks usually refer to repeated attempts to identify a user account and/or password; these repeated attempts are called brute-force attacks.

Often a brute-force attack is performed using a program that runs across the network and attempts to log in to a shared resource, such as a server. When an attacker successfully gains access to a resource, he or she has the same rights as the user whose account has been compromised to gain access to that resource. If this account has sufficient privileges, the attacker can create a back door for future access, without concern for any status and password changes to the compromised user account.

Distribution of Sensitive Information

Controlling the distribution of sensitive information is at the core of a network security policy. Although such an attack may not seem obvious to you, the majority of computer break-ins that organizations suffer are at the hands of disgruntled present or former employees. At the core of these security breaches is the distribution of sensitive information to competitors or others who will use it to your disadvantage. An outside intruder can use password and IP spoofing attacks to copy information, and an internal user can easily place sensitive information on an external computer or share a drive on the network with other users.

For example, an internal user could place a file on an external FTP server without ever leaving his or her desk. The user could also e-mail an attachment that contains sensitive information to an external user.

Man-in-the-Middle Attacks

A man-in-the-middle attack requires that the attacker have access to network packets that come across the networks. An example of such a configuration could be someone who is working for your Internet service provider (ISP), who can gain access to all network packets transferred between your network and any other network. Such attacks are often implemented using network packet sniffers and routing and transport protocols. The possible uses of such attacks are theft of information, hijacking of an ongoing session to gain access to your internal network resources, traffic analysis to derive information about your network and its users, denial of service, corruption of transmitted data, and introduction of new information into network sessions.

Protecting Your Network: Maintaining Internal Network System Integrity

Although protecting your information may be your highest priority, protecting the integrity of your network is critical in your ability to protect the information it contains. A breach in the integrity of your network can be extremely costly in time and effort, and it can open multiple avenues for continued attacks. This section covers the five methods of attack that are commonly used to compromise the integrity of your network:

- Network packet sniffers
- IP spoofing
- Password attacks
- Denial-of-service attacks
- Application layer attacks

When considering what to protect within your network, you are concerned with maintaining the integrity of the physical network, your network software, any other network resources, and your reputation. This integrity involves the verifiable identity of computers and users, proper operation of the services that your network provides, and optimal network performance; all these concerns are important in maintaining a productive network environment. This section describes the previously mentioned attacks and provide examples of how they can be used to compromise your network's integrity.

Network Packet Sniffers

As mentioned earlier, network packet sniffers can yield critical system information, such as user account information and passwords. When an attacker obtains the correct account information, he or she has the run of your network. In a worst-case scenario, an attacker gains access to a system-level user account, which the attacker uses to create a new account that can be used at any time as a back door to get into your network and its resources. The attacker can modify system-critical files, such as the password for the system administrator account, the list of services and permissions on file servers, and the login information for other computers that contain confidential information.

Packet sniffers provide information about the topology of your network that many attackers find useful. This information, such as what computers run which services, how many computers are on your network, which computers have access to others, and so on can be deduced from the information contained within the network packets that are distributed across your network as part of necessary daily operations.

In addition, a network packet sniffer can be modified to interject new information or change existing information in a network packet. By doing so, the attacker can cause network connections to shut down prematurely, as well as change critical information within the packet. Imagine what could happen if an attacker modified the information being transmitted to your accounting system. The effects of such attacks can be difficult to detect and very costly to correct.

IP Spoofing

IP spoofing can yield access to user accounts and passwords, and it can also be used in other ways. For example, an attacker can emulate one of your internal users in ways that prove embarrassing for your organization; the attacker could send e-mail messages to business partners that appear to have originated from someone within your organization. Such attacks are easier when an attacker has a user account and password, but they are possible by combining simple spoofing attacks with knowledge of messaging protocols.

Password Attacks

Just as with packet sniffers and IP spoofing attacks, a brute-force password attack can provide access to accounts that can be used to modify critical network files and services. An example that compromises your network's integrity is an attacker modifying the routing tables for your network. By doing so, the attacker ensures that all network packets are routed to him or her before they are transmitted to their final destination. In such a case, an attacker can monitor all network traffic, effectively becoming a man in the middle.

Denial-of-Service Attacks

Denial-of-service attacks are different from most other attacks because they are not targeted at gaining access to your network or the information on your network. These attacks focus on making a service unavailable for normal use, which is typically accomplished by exhausting some resource limitation on the network or within an operating system or application.

When involving specific network server applications, such as a Hypertext Transfer Protocol (HTTP) server or a File Transfer Protocol (FTP) server, these attacks can focus on acquiring and keeping open all the available connections supported by that server, effectively locking out valid users of the server or service. Denial-of-service attacks can also be implemented using common Internet protocols, such as TCP and Internet Control Message Protocol (ICMP). Most denial-of-service attacks exploit a weakness in the overall architecture of the system being attacked rather than a software bug or security hole. However, some attacks compromise the performance of your network by flooding the network with undesired, and often useless, network packets and by providing false information about the status of network resources.

Application-Layer Attacks

Application-layer attacks can be implemented using several different methods. One of the most common methods is exploiting well-known weaknesses in software commonly found on servers, such as sendmail, PostScript, and FTP. By exploiting these weaknesses, attackers can gain access to a computer with the permissions of the account running the application, which is usually a privileged system-level account.

Trojan horse program attacks are implemented using programs that an attacker substitutes for common programs. These programs may provide all the functionality that the normal program provides, but also include other features that are known to the attacker, such as monitoring login attempts to capture user account and password information. These programs can capture sensitive information and distribute it back to the attacker. They can also modify application functionality, such as applying a blind carbon copy to all e-mail messages so that the attacker can read all of your organization's e-mail.

One of the oldest forms of application-layer attacks is a Trojan horse program that displays a screen, banner, or prompt that the user believes is the valid login sequence. The program then captures the information that the user types in and stores or e-mails it to the attacker. Next, the program either forwards the information to the normal login process (normally impossible on modern systems) or simply sends an expected error to the user (for example, Bad Username/Password Combination), exits, and starts the normal login sequence. The user, believing that he or she has incorrectly entered the password (a common mistake experienced by everyone), retypes the information and is allowed access.

One of the newest forms of application-layer attacks exploits the openness of several new technologies: the Hypertext Markup Language (HTML) specification, Web browser functionality, and HTTP. These attacks, which include Java applets and ActiveX controls, involve passing harmful programs across the network and loading them through a user's browser.

Users of ActiveX controls may be lulled into a false sense of security by the Authenticode technology promoted by Microsoft. However, attackers have already discovered how to utilize properly signed and bug-free ActiveX controls to make them act as Trojan horses. This technique uses VBScript to direct the controls to perform their dirty work, such as overwriting files and executing other programs.

These new forms of attack are different in two respects:

- They are initiated not by the attacker but by the user who selects the HTML page that contains the harmful applet or script stored using the <OBJECT>, <APPLET>, or <SCRIPT> tags.
- Their attacks are no longer restricted to certain hardware platforms and operating systems because of the portability of the programming languages involved.

Trusted, Untrusted, and Unknown Networks

As a network manager creates a network security policy, each network that makes up the topology must be classified as one of three types of networks:

- Trusted
- Untrusted
- Unknown

Trusted Networks

Trusted networks are the networks inside your network security perimeter. These networks are the ones you are trying to protect. Often, you or someone in your organization administers the computers that comprise these networks, and your organization controls their security measures. Usually, trusted networks are within the security perimeter.

When you set up the firewall server, you explicitly identify the type of networks that are attached to the firewall server through network adapter cards. After the initial configuration, the trusted networks include the firewall server and all networks behind it.

One exception to this general rule is the inclusion of virtual private networks (VPNs), which are trusted networks that transmit data across an untrusted network infrastructure. For the purposes of our discussion, the network packets that originate on a VPN are considered to originate from within your internal perimeter network. This origin is logical because of how VPNs are established. For communications that originate on a VPN, security mechanisms must exist by which the firewall server can authenticate the origin, data integrity, and other security principles contained within the network traffic according to the same security principles enforced on your trusted networks.

Untrusted Networks

Untrusted networks are the networks that are known to be outside your security perimeter. They are untrusted because they are outside your control. You have no control over the administration or security policies for these sites. They are the private, shared networks from which you are trying to protect your network. However, you still need and want to communicate with these networks even though they are untrusted.

When you set up the firewall server, you explicitly identify the untrusted networks from which that firewall can accept requests. Untrusted networks are outside the security perimeter and external to the firewall server.

Unknown Networks

Unknown networks are networks that are neither trusted nor untrusted. They are unknown quantities to the firewall because you cannot explicitly tell the firewall server that the network is a trusted or an untrusted network. Unknown networks exist outside your security perimeter. By default, all nontrusted networks are considered unknown networks, and the firewall applies the security policy that is applied to the Internet node in the user interface, which represents all unknown networks. However, you can identify unknown networks below the Internet node and apply more specialized policies to those untrusted networks.

Establishing a Security Perimeter

When you define a network security policy, you must define procedures to safeguard your network and its contents and users against loss and damage. From this perspective, a network security policy plays a role in enforcing the overall security policy defined by an organization.

A critical part of an overall security solution is a network firewall, which monitors traffic crossing network perimeters and imposes restrictions according to security policy. Perimeter routers are found at any network boundary, such as between private networks, intranets, extranets, or the Internet. Firewalls most commonly separate internal (private) and external (public) networks.

A network security policy focuses on controlling the network traffic and usage. It identifies a network's resources and threats, defines network use and responsibilities, and details action plans for when the security policy is violated. When you deploy a network security policy, you want it to be strategically enforced at defensible boundaries within your network. These strategic boundaries are called *perimeter networks*.

Perimeter Networks

To establish your collection of perimeter networks, you must designate the networks of computers that you wish to protect and define the network security mechanisms that protect them. To have a successful network security perimeter, the firewall server must be the gateway for *all* communications between trusted networks and untrusted and unknown networks.

Each network can contain multiple perimeter networks. When describing how perimeter networks are positioned relative to each other, three types of perimeter networks are present: the outermost perimeter, internal perimeters, and the innermost perimeter. Figure 47-1 depicts the relationships among the various perimeters. Note that the multiple internal perimeters are relative to a particular asset, such as the internal perimeter that is just inside the firewall server.

Figure 47-1 There are three types of perimeter networks: outermost, internal, and innermost.

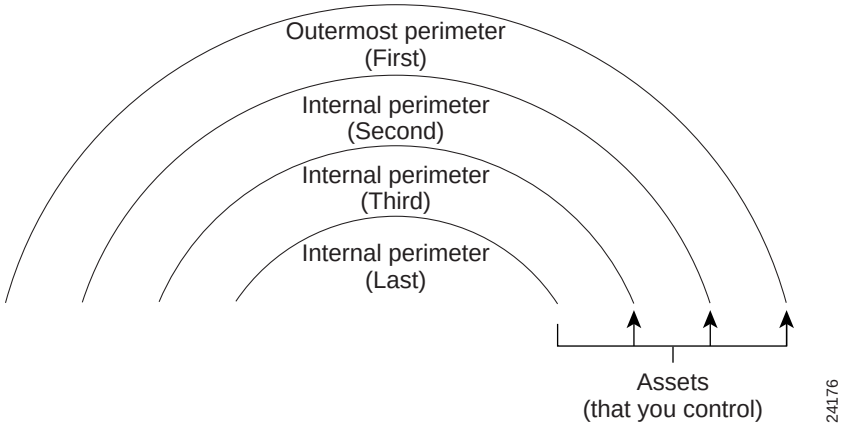

The outermost perimeter network identifies the separation point between the assets that you control and the assets that you do not control—usually, this point is the router that you use to separate your network from your ISP's network. Internal perimeter networks represent additional boundaries where you have other security mechanisms in place, such as intranet firewalls and filtering routers.

Figure 47-2 depicts two perimeter networks (an outermost perimeter network and an internal perimeter network) defined by the placement of the internal and external routers and the firewall server.

Figure 47-2 The diagram is an example of a two-perimeter network security design.

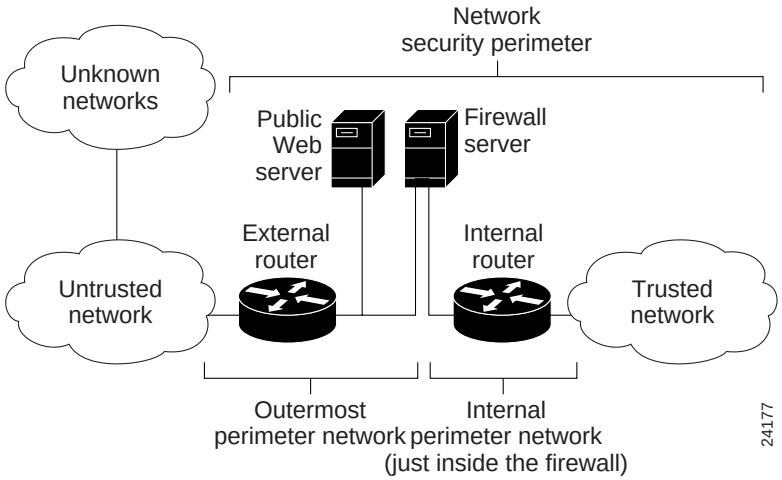

Positioning your firewall between an internal and external router provides little additional protection from attacks on either side, but it greatly reduces the amount of traffic the firewall server must evaluate, which can increase the firewall's performance. From the perspective of users on an external network, the firewall server represents all accessible computers on the trusted network. It defines the point of focus, or choke point, through which all communications between the two networks must pass.

The outermost perimeter network is the most insecure area of your network infrastructure. Normally, this area is reserved for routers, firewall servers, and public Internet servers, such as HTTP, FTP, and Gopher servers. This area of the network is the easiest area to gain access to, and therefore, it is the

most frequently attacked, usually in an attempt to gain access to the internal networks. Sensitive company information that is for internal use only should *not* be placed on the outermost perimeter network. Following this precaution helps avoid having your sensitive information stolen or damaged.

Developing Your Security Design

The design of the perimeter network and security policies requires the following subjects to be addressed.

Knowing Your Enemy

Knowing your enemy refers to knowing attackers or intruders. Consider who might want to circumvent your security measures and identify their motivations. Determine what they might want to do and the damage that they could cause to your network.

Security measures can never make it impossible for a user to perform unauthorized tasks with a computer system; they can only make it harder. The goal is to make sure the network security controls are beyond the attacker's ability or motivation.

Counting the Cost

Security measures almost always reduce convenience, especially for sophisticated users. Security can delay work and create expensive administrative and educational overhead. Security can use significant computing resources and require dedicated hardware.

When you design your security measures, understand their costs and weigh those costs against the potential benefits. To do that, you must understand the costs of the measures themselves and the costs and likelihoods of security breaches. If you incur security costs out of proportion to the actual dangers, you have done yourself a disservice.

Identifying Your Assumptions

Every security system has underlying assumptions. For example, you might assume that your network is not tapped, or that attackers know less than you do, that they are using standard software, or that a locked room is safe. Be sure to examine and justify your assumptions. Any hidden assumption is a potential security hole.

Controlling Your Secrets

Most security is based on secrets. Passwords and encryption keys, for example, are secrets. Too often, though, the secrets are not really all that secret. The most important part of keeping secrets is knowing the areas you need to protect. What knowledge would enable someone to circumvent your system? You should jealously guard that knowledge and assume that everything else is known to your adversaries. The more secrets you have, the harder it will be to keep them all. Security systems should be designed so that only a limited number of secrets need to be kept.

Remembering Human Factors

Many security procedures fail because their designers do not consider how users will react to them. For example, because they can be difficult to remember, automatically generated "nonsense" passwords are often found written on the undersides of keyboards. For convenience, a "secure" door that leads to the system's only tape drive is sometimes propped open. For expediency, unauthorized modems are often connected to a network to avoid onerous dial-in security measures.

If your security measures interfere with essential use of the system, those measures will be resisted and perhaps circumvented. To get compliance, you must make sure that users can get their work done, and you must sell your security measures to users. Users must understand and accept the need for security.

Any user can compromise system security, at least to some degree. Passwords, for instance, can often be found simply by calling legitimate users on the telephone, claiming to be a system administrator, and asking for them. If your users understand security issues, and if they understand the reasons for your security measures, they are far less likely to make an intruder's life easier.

At a minimum, users should be taught never to release passwords or other secrets over unsecured telephone lines (especially cellular telephones) or e-mail. Users should be wary of people who call them on the telephone and ask questions. Some companies have implemented formalized network security training for their employees so that employees are not allowed access to the Internet until they have completed a formal training program.

Knowing Your Weaknesses

Every security system has vulnerabilities. You should understand your system's weak points and know how they could be exploited. You should also know the areas that present the greatest danger and prevent access to them immediately. Understanding the weak points is the first step toward turning them into secure areas.

Limiting the Scope of Access

You should create appropriate barriers in your system so that if intruders access one part of the system, they do not automatically have access to the rest of the system. The security of a system is only as good as the weakest security level of any single host in the system.

Understanding Your Environment

Understanding how your system normally functions, knowing what is expected and what is unexpected, and being familiar with how devices are usually used will help you detect security problems. Noticing unusual events can help you catch intruders before they can damage the system. Auditing tools can help you detect those unusual events.

Limiting Your Trust

You should know exactly which software you rely on, and your security system should not have to rely on the assumption that all software is bug-free.

Remembering Physical Security

Physical access to a computer (or a router) usually gives a sufficiently sophisticated user total control over that computer. Physical access to a network link usually allows a person to tap that link, jam it, or inject traffic into it. It makes no sense to install complicated software security measures when access to the hardware is not controlled.

Making Security Pervasive

Almost any change you make in your system may have security effects. This is especially true when new services are created. Administrators, programmers, and users should consider the security implications of every change they make. Understanding the security implications of a change takes practice; it requires lateral thinking and a willingness to explore every way a service could potentially be manipulated.