CHAPTER 2

Why You Need a Firewall

Introduction

With the rapid growth of interest in the Internet and the Windows NT operating system, network security has become a major concern to companies throughout the world. The fact that the information and tools needed to penetrate the security of corporate networks are widely available has only 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 are constantly trying to keep abreast of the wide number of security issues confronting you in today's world. The next section describes many of the security issues that arise when connecting a private network to the Internet.

Note Although this discussion focuses on those connections made between a private internal network and the Internet, the security issues described and the ways in which Cisco Centri Firewall addresses these issues also applies to private intranetwork connections. The solutions provided by Cisco Centri Firewall are discussed in Chapter 5, "Inside the Cisco Centri Firewall."

Security Issues When Connecting to the Internet

When you connect your private network to the Internet, you are physically connecting your network to well over 50,000 unknown networks and all of their users. While 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 of your network?

The following sections describe the security issues and types of attacks focused around these two questions.

Note When people access information that they should not be accessing, or when they attempt to do something undesirable to a network or its resources, we refer to such attempts as attacks. An *attack* is some action, or attempted action, that you do not want to happen on your network. The person who performs such an action is called an *attacker.*

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 list introduces 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 preventing the theft, destruction, corruption, and introduction of information. These results can cause irreparable damage to sensitive and confidential information. Below, we describe these common methods of attack and provide 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*. Currently, Windows NT distributes network packets in "clear text;" the information sent across the network is not encrypted. (Encryption is the transformation, or "scrambling," of a message into an unreadable format 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 of the network and process them.

Note The Windows NT Remote Access Service (RAS) does provide encryption methods for protecting the packets that are sent across modem connections. The Point-to-Point Tunneling Protocol (PPTP) provides encryption between Windows NT clients and Windows NT servers over RAS, but this solution is limited. It does not include client-to-client encryption or support for non-Windows NT-based computers. Neither of these encryption techniques are standard for all TCP/IP-based communications. Third-party products are available that provide encryption for all TCP/IP-based communications.

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 Windows NT distributes 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. The more 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 work during regular employee hours (such as those in the Payroll Department), 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.

Note A *trusted computer* is a computer that you have administrative control over or one that you consciously make a decision to "trust" to allow access to 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 bi-directional 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 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 of 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 your network.

Password Attacks

Password attacks can be implemented using several different methods, including brute force attacks, Trojan horse programs (discussed later in the *Application Layer Attacks* section), 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 into a shared resource, such as a server. When an attacker successfully gains access to a resource, he 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 your network security policy. While such an attack may not seem obvious to you, the majority of computer break-ins that organizations suffer are at the hands of a disgruntled present or former employee (Miller, Stewart S., *Secure Your Data: Web Site Attacks On The Rise!*, Inter@ctive Week, January 29, 1996.). At the core of these security breaches is the distribution of sensitive information to competitors or others who will use it to your disadvantage. While an outside intruder can use password and IP spoofing attacks to copy information, an internal user can easily place sensitive information on an external computer or share a drive on the network with other users.

As an example, an internal user could place a file on an external FTP server without ever leaving his desk. He 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 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 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

While protecting your information may be your highest priority, protecting the integrity of your network is critical in your ability to protect the information that 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. In this section, we describe 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
- 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 of these concerns are important in maintaining a productive network environment. Below, we describe the previously mentioned attacks and provide examples of how they can be used to compromise your network's integrity.

Network Packet Sniffers

As we mentioned earlier, network packet sniffers can yield critical system information, such as user accounts and passwords. Once an attacker obtains the correct account information, that attacker 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 anytime 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, etc., 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 hard to detect and can be very costly to correct.

IP Spoofing

While IP spoofing can yield access to user accounts and passwords, these attacks also can be used in other ways. One way is where an attacker emulates one of your internal users in ways that prove embarrassing for your organization. For example, 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 where an attacker modifies the routing tables for your network. By doing so, an attacker ensures that all network packets are routed to him before they are transmitted to their final destination. In such a case, an attacker can monitor all of your 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 HTTP or FTP server, these attacks can focus on acquiring and keeping open all of the available connections supported by that server, effectively locking out those valid users of the server or service. Denial-of-service attacks can also be implemented using common Internet protocols, such as TCP and 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 networks by flooding the network with undesired, and often useless, network packets and by providing false information about the status of network resources.

The New York-based Internet service provider (ISP), Panix Public Access Network Corporation, recently brought to light just how vulnerable a network can be to denial-of-service attacks. Panix was subjected to an extended attack that crippled access to all of its TCP-based services, including the web sites and e-mail services that it hosted for its corporate clients. In this example, the attacker sent up to 150 connection requests per second to Panix's hosts—a number of requests that quickly filled up the hosts' crucial memory buffers (TCP pending connection buffers) with pending connection attempts, making the hosts unreachable by legitimate clients.

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 of the functionality that the normal program provides, but they 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 on 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 has incorrectly entered his 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 HTML specification, web browser functionality, and the HTTP protocol. 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 hardware platform and operating systems because of the portability of the programming languages involved.

The next chapter describes the concepts and terminology used to describe today's common firewall architectures, such as security perimeters, packet filters, and proxy services. These terms and concepts are important to understanding the unique architecture and network security advancements provided by Cisco Centri Firewall.