chapter 13

Securing the Internet

"Billions of dollars fly through the air . . . Watch your head."

—IBM E-commerce commercial

In this chapter, you will learn how to:

- Define the security threats and attacks that hackers use to gain unauthorized access to network services and resources.
- List the Internet security safeguards that protect networks by detecting intrusions and defeating attacks.
- Define the methods for digitally signing and encrypting network transmissions to thwart masquerading and deciphering by unauthorized parties.
- Describe how to publish a Web securely with the SFTP protocol.

IMAGINE what would happen if the Net stopped working. The world got a warning when Slammer attacked on January 25, 2003. For several hours that Saturday morning, while network administrators scurried to patch hundreds of thousands of database servers infected by the attack, bank networks were out of service. People could not withdraw money or buy anything online. Luckily, Slammer attacked on a weekend, when the stock markets were closed. Imagine the pandemonium of bringing down the markets.

This chapter is about securing the Internet. By defining the security threats that can bring down a network, you will know the techniques crackers use to attack the Internet. You will learn best-practice methods of detecting intrusions and defeating attacks. The vulnerability Slammer exploited, for example, was a known issue. Server administrators who follow best practices had already installed the patch. Slammer succeeded because too many networks had neglected to install the patch. The extent of the outage demonstrated how important it is for all server administrators to follow the best security practices this chapter presents.

The Internet exists because of the wonderfully creative minds of its inventors. Hackers also have creative minds. Malicious hackers, who are more correctly called crackers, continue to scheme and dream of new ways to attack the Internet. To fool you into installing malicious code, crackers masquerading as legitimate vendors try to trick you by sending messages disguised as security update bulletins. On first glance, the messages appear legitimate, because the crackers emulate the look and feel of the vendor's official update mechanism. This chapter teaches you the digital signing technology that legitimate vendors use to foil the masquerade. Learning how to use the Internet's public-key infrastructure will enable you to encrypt messages and prevent deciphering by unauthorized users. The chapter concludes by teaching you how to publish a Web site securely and set permissions that determine which users have different kinds of access to your site.

Identifying Internet Security Issues

Internet security risks fall into three general categories: (1) unauthorized access, (2) data manipulation, and (3) service interruption. Unauthorized access happens when someone who is not permitted manages to obtain

information from the site. Prime targets of unauthorized intruders include financial records, credit card numbers, and mailing lists.

Data manipulation happens when an intruder modifies records in the database. This kind of attack occurs less frequently because modifying data leaves tracks that can lead to identifying the intruder.

Service interruption happens when malicious code floods a server with bogus requests. Legitimate users experience long delays or cannot log on at all because the server is busy processing the bogus requests or the server crashed under the onslaught.

User-level Issues

It is generally believed that about two-thirds of all security breaches come from inside an organization's network architecture. Known as **inside attacks**, these breaches happen when employees hack in to obtain unauthorized access to network services or data. Malicious employees can cause a lot of harm, especially if they happen to be IT staff members or system administrators. Because of this threat, companies typically remove an employee's computer access before terminating employment. Even so, knowledgeable employees may still have ways to hack into the network, such as through a back door, which is a loophole a programmer creates to bypass the application security mechanisms.

Physical Access Security

To help guard against inside attacks, you should locate equipment behind locked doors and limit access to authorized personnel. Require employees to log off before walking away from their workstations to prevent someone unauthorized from sitting down and continuing the session. Keep employees from writing down their passwords on slips of paper that an eavesdropper could see. Encourage employees to report suspicious activity, such as equipment being moved or worked on by someone who is not a member of the IT staff.

Network Security Threats

Some security threats are more sophisticated than others. Network security threats that rely on technical measures may take the following forms:

- **Data interception** Packet sniffers and network analyzers can intercept data that moves across the network. If the data transmissions are in clear text, unauthorized personnel can capture and access sensitive data. Later in this chapter, you learn how to encrypt network transmissions to prevent sensitive information from crossing the Internet in clear text.
- **Identity interception** Most sites rely on user names and passwords as the first level of security. If the logon form sends user names and passwords across the network in clear text, intruders with packet sniffers can intercept identities. Hackers also intercept identities by guessing passwords. Because password hacking tools use dictionaries

of common words in many languages, you should require employees to have passwords consisting of a combination of characters and numbers. Hackers can also use search engines to find out information about you; avoid passwords consisting of information that can be searched or guessed, such as family member names, birth dates, hobbies, phone numbers, or license plate numbers.

- **Masquerading** When unauthorized users assume the privileges of an authorized user, they are **masquerading**. This can happen through identity interception and is especially dangerous if a hacker discovers a system administrator's username and password. A more subtle form of masquerading is **IP address spoofing**, which happens when an intruder uses the IP address of a trusted system to gain access rights granted to that system.
- **Replay attacks** If a hacker uses a packet sniffer to record a logon sequence, the hacker can playback the sequence at a later time in an attempt to gain access. In this kind of **replay attack**, the intruder does not need to know the username or password. Even if the logon information is encrypted, the replay attack will succeed if the logon server accepts the same codes.

The technical measures employed in these kinds of attacks are sophisticated, and the hackers who can pull them off are impressive. Nevertheless, you must also be on guard for less technically sophisticated means of compromising network security. The following methods exploit human weaknesses that are no less dangerous than the more technical security threats:

- **Social engineering attack** A **social engineering attack** exploits human weaknesses to gain access to the organization's network. The attacker may pretend to be from the IT department, call new employees on the phone, and ask them to verify their passwords. An attacker may call the IT help desk in the guise of a user who needs help logging on. To thwart this kind of attack, the organization should communicate clearly outlined IT procedures to all staff, especially newcomers, to clarify the manner in which passwords and other kinds of sensitive information will be communicated.
- **Misuse of privileges** A network administrator with a high level of system privileges could use them to poke into different parts of the system and look up information about fellow employees or alter financial information such as the payroll. You must therefore give system privileges only to trusted employees of the highest integrity.

Identifying Assets

Depending on the nature of the organization, some security threats will be more of an issue than others. Because it is impossible to prepare for all possible attacks, identifying assets helps prioritize the need for different kinds of defenses.

At a public service Web site that has the goal of making information freely available, for example, the greatest threat is denial of service. In a

bank, on the other hand, security administrators must guard against fraud, which insiders as well as outsiders can commit. Critical assets may reside on one of the following resource groups:

- **Data tier information resources** Any organization that conducts transactions has a back office database that you do not want hacked. Customer information, payroll, student grades, credit card numbers, and mailing lists are prime targets of crackers.
- **Server resources** All of the servers covered in Chapter 11, including Web servers, mail servers, FTP servers, and directory servers, may contain assets that you wouldn't want hacked.
- **Network resources** You need to protect network resources from tapping and spoofing to prevent crackers from gaining unauthorized access to the network.
- **Local workstations** End-user workstations are especially prone to viruses that can cause significant productivity losses, destroy data, and spread to other computers on the network. Every organization must guard against virus attacks.

Viruses and the Hacker Process

Hackers are constantly trying to devise new ways of breaking through network security mechanisms. Some hackers obtain satisfaction from the act of breaking in. For them, cracking security is a kind of game that proves their prowess in hacking. Malicious hackers, on the other hand, can disrupt business, resulting in financial losses that can be considerable.

Defending a network is a two-stage process that requires (1) a proactive pre-attack strategy and (2) a reactive post-attack strategy. In formulating a pre-attack strategy, you list the threats and identify the IT staff that will be responsible for the defenses. The IT staff must be proactive, meaning that they keep up with the latest intelligence on the hacking process, anticipate things that could go wrong, and implement defensive measures in advance of an attack.

Attackers will sometimes succeed, however, especially if your network gets hit by a new virus before security patches become available. You must therefore have a post-attack strategy for mobilizing the appropriate staff to take corrective actions. When my organization got hit by the Slammer worm, for example, dozens of servers were affected. Fast-acting IT staff minimized damage by blocking traffic on port 1434, which runs the Microsoft SQL resolution service over which Slammer waged its attack. This bought precious time for server administrators to apply the security patch before reopening port 1434 to network traffic.

Viruses top the list of attacker's tools. Viruses propagate by attaching themselves to software programs or other kinds of files. When a user executes or opens a file thus infected, the user's computer catches the virus. Types of viruses include:

- **Boot record virus** When a computer boots, it runs code from the boot record of the startup device. **Boot record viruses** spread

through malicious code that runs when the computer boots. You need to warn all employees to be leery of removable media such as USB drives and floppy disks, which can transmit a boot record virus on startup if left in the drive.

- **File infector virus** Malicious code can attach to individual files, which propagate primarily via e-mail attachments. When the user opens the attachment, the file infector virus executes. Every workstation should have a virus scanner configured to scan all incoming files for these kinds of viruses.
- **Macro virus** A macro is a command that executes a set of instructions in a computer application such as a word processor or a spreadsheet to save time users would otherwise spend typing the instructions individually. Macros can carry viruses. One of the most common transmission methods is attaching to an e-mail message a document or spreadsheet containing the **macro virus**. You must therefore institute policies and take protective measures to prevent employees from opening documents from non-trusted sources.
- **Trojan horse** A **Trojan horse** is malicious code that masquerades as a desirable program. Users fooled into running the program execute the malicious code. Hackers have succeeded in delivering Trojans in the guise of security updates attached to e-mail messages made to look as though they are coming from legitimate sources. You need to make sure everyone in your organization knows that no legitimate vendor would ever send a security patch as an attachment to an unsolicited e-mail message.
- **Embedded code** Crackers can embed malicious executable code in Web pages via Java applets or ActiveX controls. You must warn users about the danger of following links to non-trusted sites.

Another kind of malicious code is the **worm**, which can propagate across the Internet and infect other computers without attaching itself to other programs. Worms can wreak havoc by replicating themselves on an infected system to tie up all its resources and then following that system's network connections to attack other computers.

Try This!

Detecting Unauthorized Services

Many users are unaware that a clever attacker can install services that run on the end-user's computer. An attacker who gains access to your computer can install an FTP service, for example, and use your computer to distribute files to other computers, including files propagating the virus that installed the unauthorized process on your workstation. To discover whether any unknown services are running on your computer, you can download and install a freeware program called Process Explorer. Follow these steps:

Try This!
continued

1. The free download version of Process Explorer is a zip file. If you do not have a zip file extractor on your computer, go to www.pkware.com and install PKZip for Windows. There is an evaluation version in the free download section at www.pkware.com/products/free_eval.html.
2. Go to www.sysinternals.com and look for a link to Process Explorer. When this book went to press, the link was www.sysinternals.com/ntw2k/freeware/procexp.shtml. Click the Process Explorer link.
3. On the Process Explorer page, scroll down to find the download link for the version of Process Explorer appropriate for your operating system. You will see links to download versions of Process Explorer for Win9x/Me, WinNT/2K/XP, and 64-bit versions of Windows.
4. Click the download link appropriate for your operating system. When prompted, choose the option to open the file.
5. After the zip file downloads and opens, double-click the file named *procexp.exe*, which runs the Process Explorer.
6. If you get a warning that symbols are not configured, click OK to dismiss the warning. You will not need symbols for this exercise.
7. Figure 13-1 shows how the Process Explorer displays the processes running on your computer. Peruse this display to determine whether any unauthorized services are running on your computer.

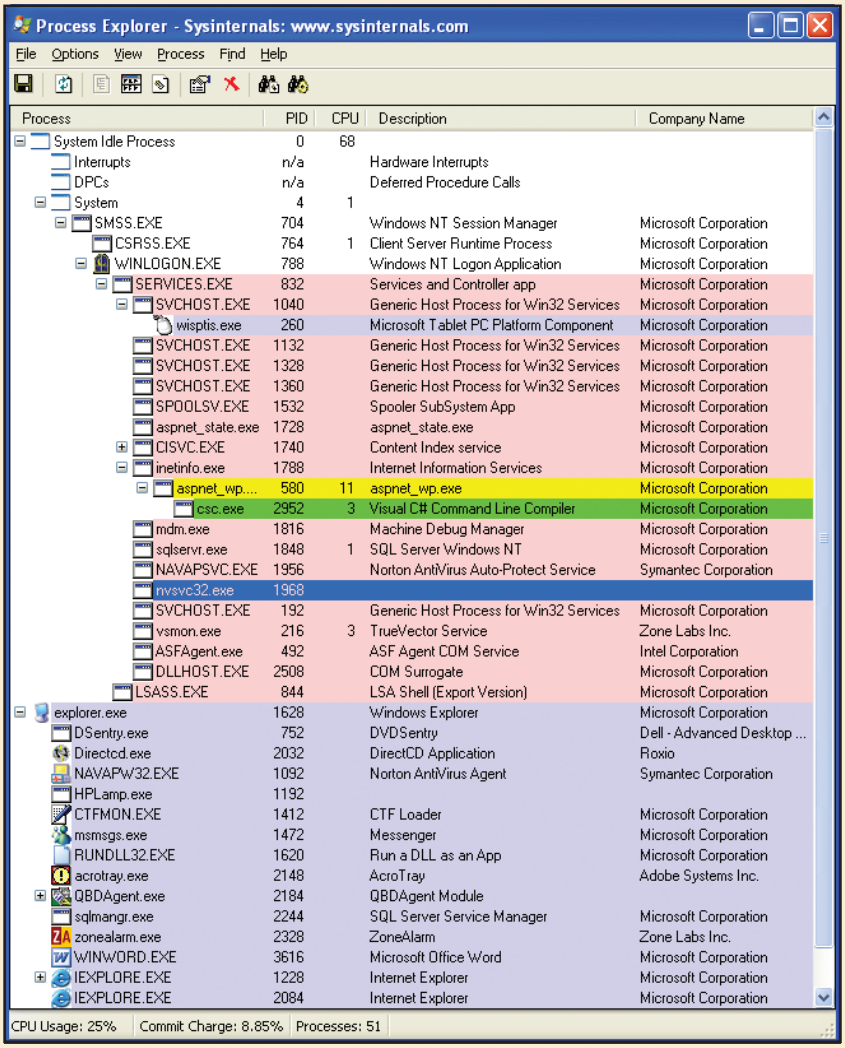

FIGURE 13-1 *The freeware program Process Explorer displays the processes running on your computer. Be leery of processes that have no descriptions or company names. In this example, I suspected nvsvc32.exe because it has neither a description nor a company name. Upon closer inspection, however, this turned out to be the computer's graphics driver.* ■

Applying Internet Security Safeguards

Due to the global importance of Internet security, the computer industry has evolved best practices you can follow to safeguard your network's client workstations and servers. These best practices include (1) subscribing to a security newsletter that keeps you apprised of the latest security issues and threats, (2) using an automatic update service to install the latest security patches, (3) identifying the kinds of attacks to which your network is prone, (4) auditing the network for traces of these attacks, (5) installing software that can automatically detect intrusions, (6) planning how to recover from network disasters, and (7) using firewalls to block non-trusted traffic or processes.

Microsoft Security Newsletters

For millions of end users and server administrators that have the Windows operating system, subscribing to one or more of the Microsoft Security Newsletters is an excellent way to keep up with Microsoft's security strategies and guidance. To subscribe, follow these steps:

1. Go to www.microsoft.com/technet/security/secnews/newsletter.htm and follow the link to subscribe to a newsletter.
2. When the Passport screen appears, log on via your Passport credentials. If you do not have a Passport, follow the onscreen instructions to get one.
3. The Microsoft subscription center appears, as illustrated in Figure 13-2. Follow the link to Newsletter Subscriptions.
4. Figure 13-3 shows how you can subscribe to a wide range of newsletters. Server administrators and IT staff should subscribe

note *Passport is Microsoft's online authentication service.*

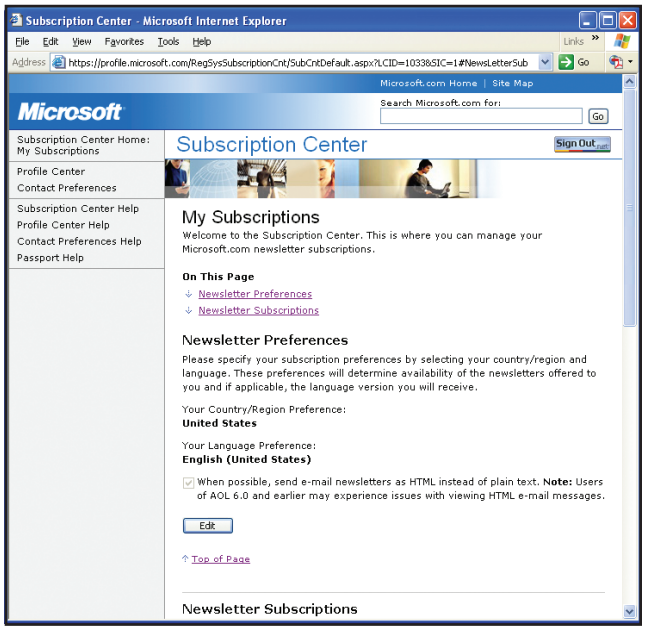

FIGURE 13-2 *The Microsoft Subscription Center lets you set your Newsletter Preferences. Follow the link to Newsletter Subscriptions to bring up the newsletter choices in Figure 13-3.* ■

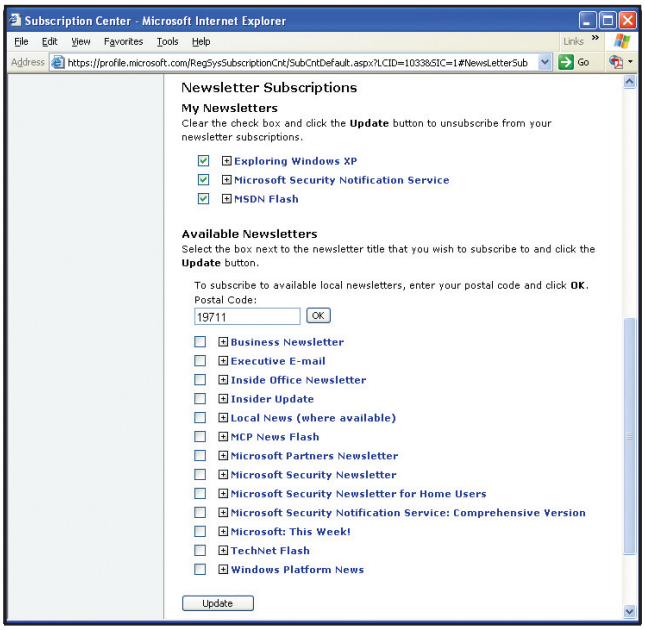

FIGURE 13-3 *Check the boxes alongside the newsletters to which you want to subscribe, and then click the Update button. To find out more about a newsletter, click its plus sign.* ■

to the Microsoft Security Notification Service, which provides detailed technical information on security issues. Other employees should subscribe to the Microsoft Security Newsletter.

5. Check the boxes alongside the newsletters to which you want to subscribe and click the Update button.
6. If you later decide to unsubscribe, repeat these steps but uncheck the boxes alongside the newsletters you want to stop and click the Update button.

Microsoft Windows Update Service

Microsoft runs a **Windows Update Service** that can automatically download the latest security patches to your computer. You can configure the service to install the patches right away or to notify you when the patches are ready to install. To configure the Microsoft Windows Update Service, follow these steps:

1. Log on to your computer as the administrator. Only the administrator can configure the Windows Update Service.
2. Click the Windows Start button, go to the Control Panel and double-click System to bring up your computer's System Properties window.
3. Click the Automatic Updates tab. Figure 13-4 shows how the configuration settings appear onscreen. To enable automatic updating, check the box alongside the option to *Keep my computer up to date*.
4. The settings let you choose one of three service levels. The first level notifies you before downloading any updates. The second level downloads the updates automatically but notifies you before installing them. The third level downloads and installs the updates automatically.
5. Click OK when you finish configuring the settings.

Besides configuring Windows Update to run automatically, you can also force it to run manually at any time. To run Windows Update manually, follow these steps:

1. Click the Windows Start button and choose Help and Support to bring the help center onscreen.
2. Follow the link to Windows Update. When the Windows Update screen appears, follow the link to scan for updates.
3. After scanning for updates, the Windows Update service will display a list of (1) critical updates that are very important for you to install, (2) operating system updates, and (3) driver updates. Follow the

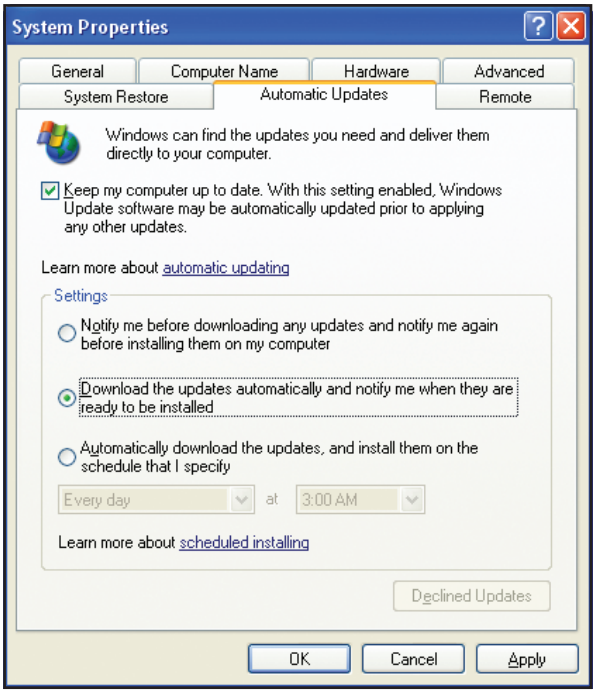

FIGURE 13-4 *You configure the Windows Update settings on the Automatic Updates tab of the System Properties window that you bring onscreen by choosing System from the Windows Control Panel. Shown here are the settings on my home computer, which downloads the updates automatically and notifies me when they are ready to be installed.* ■

onscreen instructions to review the updates and choose the ones you want installed on your computer. By definition, you should always install the critical updates because they patch serious bugs and security holes that make your computer vulnerable to attack.

Defeating Attacks

No matter how careful you are about protecting your network, it will inevitably come under attack. The more you understand about the methods employed by different kinds of attacks, the better your instincts will work in guiding you to take appropriate actions.

The most frequent kind of attack is **Denial of Service (DoS)**, in which the attacker seeks to consume so much of a server's resources that the host cannot respond to legitimate requests. A DoS attacker may send a flood of requests, such as ping commands, in an attempt to overflow the server's input buffer. Due to the overflow, legitimate requests cannot get in. The worm named Slammer, for example, waged its attack by flooding port 1434 with SQL ping requests. When hundreds of thousands of SQL servers began echoing each other's ping requests, large portions of the Internet went out of service for several hours until server administrators applied the patch.

Another mode of DoS attack is the SYN flood, in which an attacker floods a server with the synchronize (SYN) command that is the first part of the TCP handshake. To thwart these kinds of floods, routers can block traffic from an IP address that is obviously sending too many packets. Crackers will continue to find other ways to launch a DoS attack, however. They may send you a flood of megabyte e-mail messages to use up all of your server's memory. To keep you from blocking the attack, crackers may use multiple servers and e-mail addresses from which to launch the attack. It is important for you to subscribe to the security newsletters and apply patches promptly to thwart newly emerging attacks. Most of the recent DoS attacks succeeded because server administrators and end-users were slow to apply the patches.

You need to guard against ways in which computers can be programmed to guess usernames and passwords. In a so-called **brute force attack** or **front door attack**, for example, a cracker programs a computer to look up words in a dictionary and generate variants with which the computer keeps trying to log on until it discovers a password that gets in. You can minimize the risk of a front-door attack by limiting the number of logon attempts permitted from an IP address that keeps getting the password wrong.

Do not forget that human beings can create security gaps that require less sophistication to exploit. Dumpster diving is a term used to describe the practice of looking through trash for discarded records that may display in clear text important information such as account numbers, passwords, and social security numbers. Your security policy needs to require that sensitive printed information be shredded before being discarded.

Another security risk created by well-intentioned IT staff is the trapdoor attack, in which crackers find a way into your computer by running diagnostic tools that your IT staff left on the system after troubleshooting some kind of problem. It is important for the IT staff to remove

diagnostic tools and restore security settings to their original state after debugging a problem.

Security Auditing and Intrusion Detection

Security needs to be proactive as well as reactive. **Security auditing** is a proactive process that considers the risks associated with security assets, predicts the methods crackers may use to exploit each risk, and takes protective steps to thwart them. Just as crackers may write computer code to attack your network, so also can you use software to detect attempts to compromise your assets. Microsoft Windows servers, for example, have intrusion detection software for keeping an audit trail that tracks logon attempts and logs changes made to vital assets. The audit trail enables the server administrator to determine whether an attacker is trying to guess passwords. By keeping track of who changes what, the audit trail can identify insiders who may use legitimate logons to make inappropriate use of network assets.

To activate intrusion detection on a Windows server, you set an audit policy on the asset you wish to track. You can specify whether to audit successes, failures, or both. Figure 13-5 shows how audit policies can apply to logons, account management, directory services, objects, policy changes, privileges, processes, and system events. To determine whether an insider may have changed a security setting that opened a trapdoor, you can audit policy changes.

You need to be aware of some pitfalls related to security auditing. In a denial of service attack that sends a flood of requests against the process you are auditing, the attacker may succeed in filling up the log. One of the audit policy settings calls for shutting down the server if security audits cannot be logged. If you have the shut-down setting configured for an asset targeted by a denial of service attack that manages to fill the log, the shutdown will enable the attacker to succeed in denying service to other users.

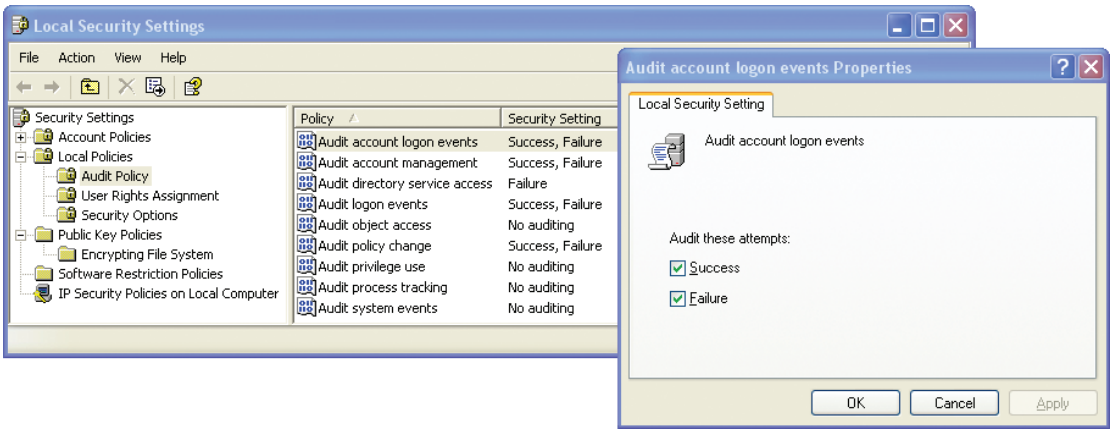

FIGURE 13-5 *To set an audit policy, use Windows control panel | Administrative tools to open the Local Security Policy Settings window shown here. Click to expand the Local Policies tree, and click Audit Policy to display the audit policies. Double-click the policy you want to change. Check the boxes to set whether you want to audit successes, failures, or both.* ■

It is not good practice to log more than you would ever have time to look at. Before you begin setting audit policies, decide what assets are vital, and let this decision be your guide.

Remember that security logs are computer files that crackers could destroy to cover their tracks in a successful attack. Because the logs prior to the attack may track the attacker's attempts to hack your server, studying these logs could be invaluable in identifying the source of the attack. You should back up the security logs regularly and keep them on secondary storage devices, such as optical media that would not get erased in an attack.

Firewall Strategies

The typical end user does not have the degree of sophistication required to guard against many of the different kinds of attacks described in this chapter. Happily, firewall technology enables an organization to take some of the burden off end users. A **firewall** is a combination of hardware, software, and security policies that block certain kinds of traffic from entering or leaving a network, subnet, or individual host computer. There are many ways of configuring firewalls. Depending on the security risks and the importance of the assets being protected, firewalls may make use of one or more of the following strategies:

- **Packet Filtering** In a firewall, a packet filter works at OSI/RM Layers 2 and 3 to inspect the headers of all incoming and outgoing packets and can block transmissions based on source or destination ports or IP addresses. Thus, packet filters can stop attacks waged on specific ports and block access to malicious or forbidden sites. Blocking access to ports that a network does not use protects the network from attackers that may launch attacks on those ports.
- **Proxy Servers and Network Address Translation** In the previous chapter, you learned how a proxy server can apply a technique known as Network Address Translation (NAT) to use different IP addressing for external traffic than the addresses on an internal network. Keeping the internal IP addresses private helps hide them from crackers.
- **Circuit Level Gateway** Figure 13-6 illustrates how a **circuit level gateway** prevents the establishment of end-to-end TCP connections. Instead, the gateway establishes a connection on behalf of an inside host with an outside host, which views the gateway as its destination address. Thus, the circuit level gateway serves as a proxy at the Transport Layer of the Internet.
- **Application Level Gateway** Neither packet filters nor circuit level gateways inspect the packets' contents at the Internet's Application Layer, which can transmit viruses attached to e-mail messages or Trojans embedded in Web

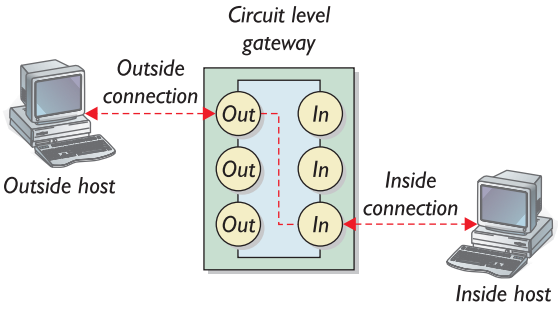

FIGURE 13-6 *When an inside host requests a connection with an outside host, a circuit level gateway sets up two TCP connections, one with the outside network and the other with the inside. Thus, the circuit level gateway works as a proxy that does not permit end-to-end TCP connections.* ■

pages. To protect an internal network from these kinds of attacks, Figure 13-7 shows how you can install an **application level gateway**, a type of firewall that can scan packets for malicious content spread through SMTP (mail), HTTP (Web pages), FTP (file transfers), DNS (attacks on name servers), or Telnet (remote logon). Depending on the network's security policy, an application level gateway can detect and block viruses, perform lexical analysis on message content, block access to offensive content, and report suspicious activity. This kind of analysis can be time-consuming, however; network administrators need to consider the impact of this degree of screening on overall network performance.

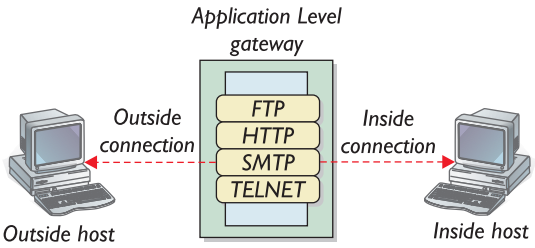

FIGURE 13-7 *An application level gateway acts as a relay that has full packet awareness. It can scan incoming mail messages for viruses, and it can perform lexical analysis to block access to offensive content and report suspicious activity. This kind of application-level screening adds overhead, however, so you should use it only when you really need it.* ■

- **Stateful Inspection** Many applications open temporary ports at addresses above the Internet's Well Known Ports from 0 through 1023. To thwart attacks on temporary ports, a firewall technique known as **stateful inspection** can keep track of when a port opens, what session is using it, and how long the port stays open. If an attacker tries to hijack a session, the firewall can detect the attack and drop the session.

Firewall Topologies

After considering the risks against the security assets, the network administrator implements the appropriate strategies by configuring the hardware, software, and security policies that the firewall comprises. Depending on the configuration, the resulting firewall has one of four topologies, which the following sections define and illustrate the following:

- Packet filtering firewall
- Single-homed bastion host firewall
- Dual-homed bastion host firewall
- Screened subnet firewall with DMZ

Packet Filtering Firewall

The simplest kind of firewall uses a **packet filter**, which monitors the headers of all incoming or outgoing packets and can block transmissions based on source or destination ports or IP addresses. Figure 13-8 illustrates how a packet filtering router can protect an entire network. Because the packet filter operates at Layers 2 and 3 of the OSI reference model, it can quickly inspect the TCP/UDP ports and IP addresses in each packet. Network administrators configure the filter to block certain kinds of traffic that the organization's security policies forbid.

Locating the packet filter at the entrance to the network provides a systemic way for a network administrator to act quickly in thwarting an attack on a specific incoming port. Packet filters enabled network

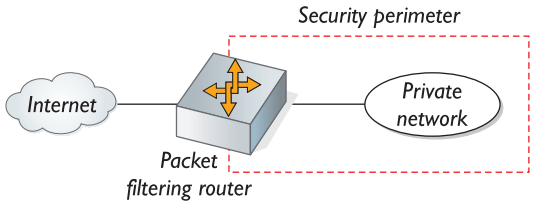

FIGURE 13-8 *Packet filtering firewalls work at OSI/RM Layers 2 and 3 by monitoring the TCP/UDP port numbers and IP addresses in each packet. Routers with packet filtering are relatively inexpensive and easy to install. Locating the packet filter at the entrance to the network provides a systemic way for a network administrator to act quickly in thwarting an attack on a specific incoming port, such as the Slammer attack that exploited the SQL ping vulnerability on port 1434.* ■

administrators to stop the Slammer attack, for example, by blocking traffic on port 1434, which Slammer used to wage its attack. Unprotected networks without packet filtering firewalls were not so lucky.

Single-Homed Bastion Firewall

Although packet filtering firewalls perform an important and essential function by monitoring headers at OSI/RM Layers 2 and 3, certain kinds of attacks happen higher up the protocol stack. Viruses that come in via e-mail, for example, attack at the application level. To detect e-mail viruses, you need to scan the mail messages and attachments. To accomplish this kind of application scanning, Figure 13-9 shows how you can install a bastion host to screen traffic inside the perimeter of your network. A **bastion host** is a computer that sits on the perimeter of a local network and serves as an application-level gateway between the external network and the internal client workstations.

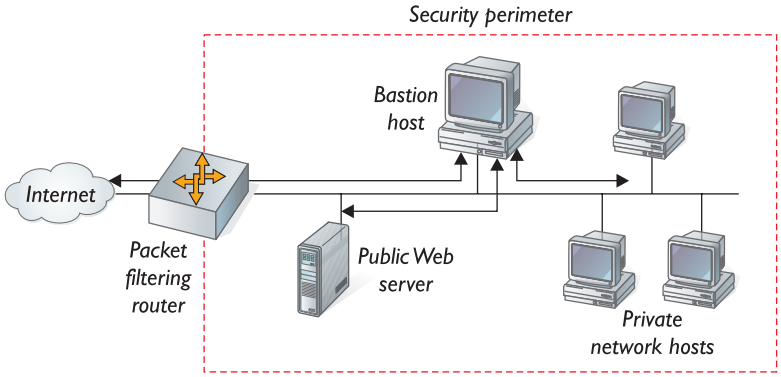

FIGURE 13-9 *The single-homed bastion firewall topology uses a packet filter to establish the security perimeter, inside of which a gateway computer containing a single NIC (hence the term single-homed) serves as a bastion host through which all local network traffic passes. The bastion host can inspect messages at the Application Layer and, depending on the organization's security policies, block transmissions containing viruses or offensive content.* ■

Because the bastion host shown in Figure 13-9 has a single Network Interface Card (NIC), this kind of firewall configuration is called a **single-homed bastion**, which can monitor all of the network's incoming and outgoing traffic at the Application Layer. A single-homed bastion firewall could scan all mail messages for viruses, for example, and quarantine infected messages instead of passing the viruses on to the network. The bastion host could also prevent downloading from the Internet, block access to certain Web sites, or alert administrators when messages contain certain keywords, such as obscenities or sexually explicit content. Before installing a bastion host, however, you need to weigh the potential benefits against the overhead of the added processing time. The more monitoring you do at the application level, the longer end-users need to wait because of increased network delay. When used in combination with a packet filter, on the other hand, the bastion introduces a second layer of security that makes it harder for a cracker to break in.

Dual-Homed Bastion Firewall

Figure 13-10 illustrates how the **dual-homed bastion** firewall uses two NICs (hence the term dual-homed) on which IP forwarding is disabled, thereby creating a complete physical break between the internal and external networks. This adds a second layer of defense by making it impossible for a cracker to sneak packets past the perimeter without being screened

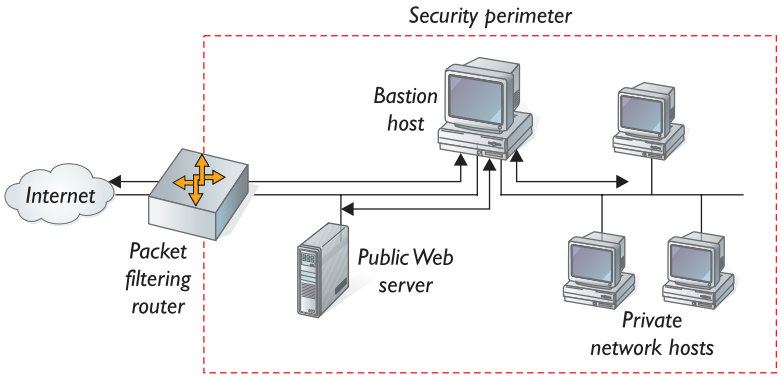

FIGURE 13-1 *The dual-homed bastion host uses two NICs with IP packet forwarding disabled, thereby completely isolating the internal network from the external Internet. This topology makes it impossible for a packet to sneak past the bastion, which can monitor, filter, block, quarantine, and log any kind of information passing through the network.* ■

by the bastion. Because the dual-homed bastion firewall screens each incoming and outgoing packet, it is also known as a screened host firewall.

DMZ Screened-Subnet Firewall

Figure 13-11 shows how the **screened-subnet firewall** establishes a **demilitarized zone** (DMZ) by placing packet filtering routers on both the Internet side and the private network side of the bastion host. This makes it impossible for insiders to communicate directly over the Internet, because the inside router's address does not appear in the Internet's routing tables, which see only the DMZ's outside router. Thus, insiders cannot bypass the private network's security measures. External crackers must now fool three devices, namely, the bastion host and the two packet filtering routers. Furthermore, the DMZ provides a secure location for the network's modem pool and the organization's public Web and FTP servers.

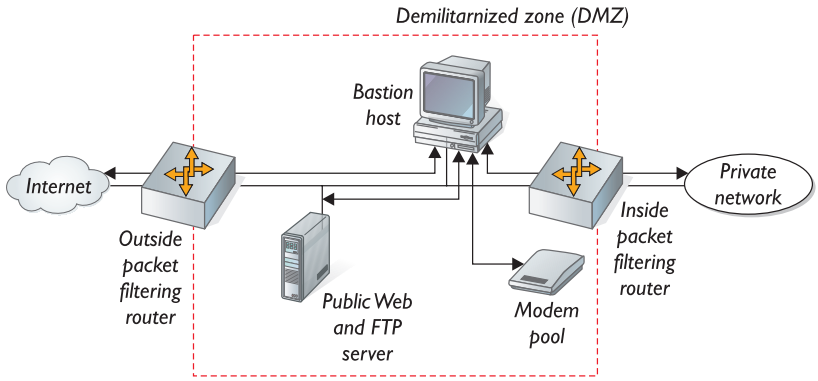

FIGURE 13-2 *A screened subnet firewall uses two packet filtering routers to create a demilitarized zone (DMZ), in which you locate the bastion host, the public Web and FTP servers, and the organization's modem pool. Because the Internet's routing tables do not know the address of the inside router, the DMZ blocks direct traffic across the screened subnet. This makes the internal network invisible to the public Internet.* ■

note *At first glance, you may not notice the subtle difference between the single-homed bastion host in Figure 13-9 and the dual-homed host in Figure 13-10. The only difference is the number of NICs in the bastion host. In Figure 13-9, the network traffic flows in and out of the bastion via the single NIC. In Figure 13-10, on the other hand, there are two NICs with IP packet forwarding disabled, which completely isolates the internal network from the external Internet.*

For more on firewalls, go to Microsoft's firewall page at www.microsoft.com/technet/security/guidance/secmod155.mspx, which explains the security policies you can configure with different kinds of firewalls.

Try This!

Installing ZoneAlarm

ZoneAlarm is a popular firewall product. There is a freeware version you can download and install at no cost on any Windows computer. Millions of users run the free version as their personal firewall. In this exercise, you install the freeware version and explore its abilities to block traffic from computers that are not trusted, restrict specific programs on your computer from accessing or serving other computers, view security alerts and logs, and quarantine VBScript files in e-mail attachments. To download and install the free version of ZoneAlarm, follow these steps:

1. Go to www.zonelabs.com and follow the link to ZoneAlarm (free). When the download screen appears, click the ZoneAlarm download button. Then click the link to download free ZoneAlarm.
2. The free ZoneAlarm download is a self-extracting archive that has the *.exe* filename extension. When prompted, save this file on your computer in the folder of your choice.
3. When the download completes, click the Open button to run the self-extracting archive, and follow the onscreen instructions to install ZoneAlarm. When asked whether you want to take the tutorial to learn how to use ZoneAlarm, say Yes to study the instructions.
4. Figure 13-12 shows that ZoneAlarm presents the settings in five categories named Overview, Firewall, Program Control, Alerts & Logs, and E-mail Protection. You should set up ZoneAlarm to load automatically on startup. To check that setting, click the Preferences tab in the Overview category.
5. Figure 13-13 shows that the ZoneAlarm Firewall category has two zones. The Internet zone contains unknown computers, and the trusted zone is where you define the computers you choose to trust. To put a computer into the trusted zone, click the Firewall category's Zones tab and then click the Add button. A menu prompts you to enter a host by name or IP address. Make your choice and follow the onscreen instructions.
6. When programs on your computer attempt to access the Internet, ZoneAlarm will prompt you to permit, deny, or allow access just this once. Figure 13-14 shows the Program Control setting that makes this happen. You need to be careful, however, in granting programs permission to access the Internet. To review the programs to which you have permitted or denied access, click the Programs tab, which brings up the screen illustrated in Figure 13-15.
7. To see the logs kept by ZoneAlarm, click the Alerts & Logs category and then click the Log Viewer tab. For help on this or any other ZoneAlarm feature, click the Help button that is in the upper-right corner of every ZoneAlarm screen.
8. If you like the free version of ZoneAlarm, you may wish to consider the added security features of ZoneAlarm Pro, which quarantines many kinds of harmful e-mail attachments, protects against spoofing, identifies the physical location of attackers, blocks ads and popups, manages the cache, and controls cookies. Go to www.zonelabs.com and follow the link to learn more about ZoneAlarm Pro.

Try This!
continued

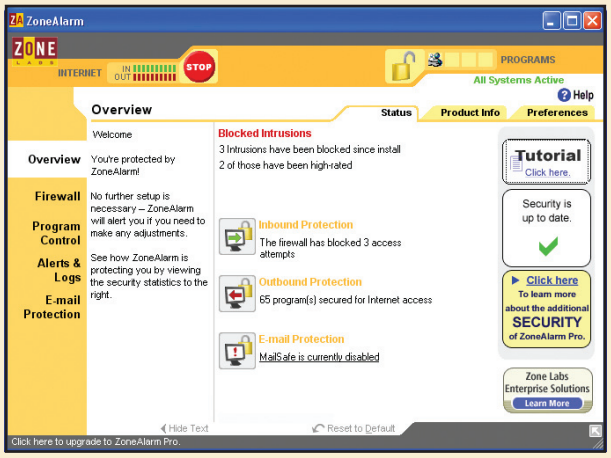

FIGURE 13-12 *To configure ZoneAlarm, you click one of the five general categories in the menu on the left, and then click the tab containing the settings you want. Shown here is the status tab of the Overview category.* ■

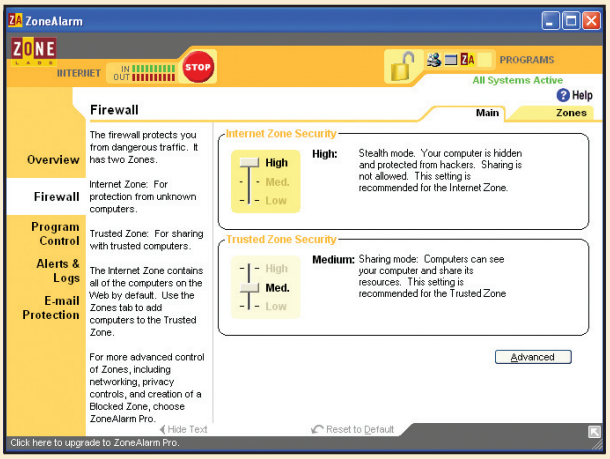

FIGURE 13-13 *ZoneAlerm's Firewall category has two zones: the Internet zone and the trusted zone. The Internet zone contains unknown computers, and the trusted zone is where you define the computers you choose to trust. You should set the Internet zone's security to the high setting, which hides your computer from workstations that are not in your trusted zone.* ■

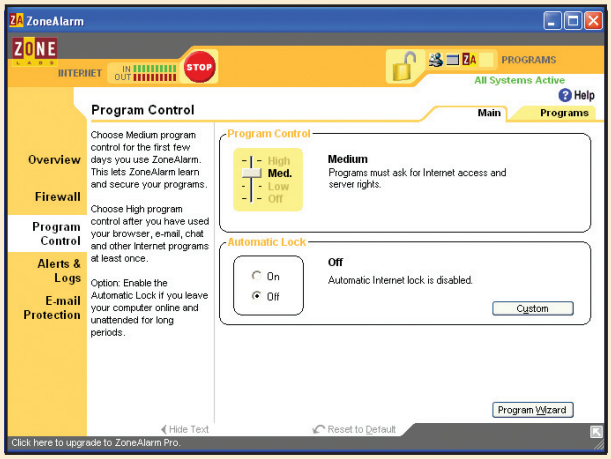

FIGURE 13-14 *Setting the program control to Medium causes ZoneAlarm to prompt you before permitting programs to access the Internet and respond to external requests. When a program asks for permission, think carefully before permitting it. The biggest problem with ZoneAlarm is that many users click Yes without thinking. If you open an e-mail message, for example, and ZoneAlarm asks whether you want to permit a new program to have Internet access, it could be malicious code coming in via the e-mail.* ■

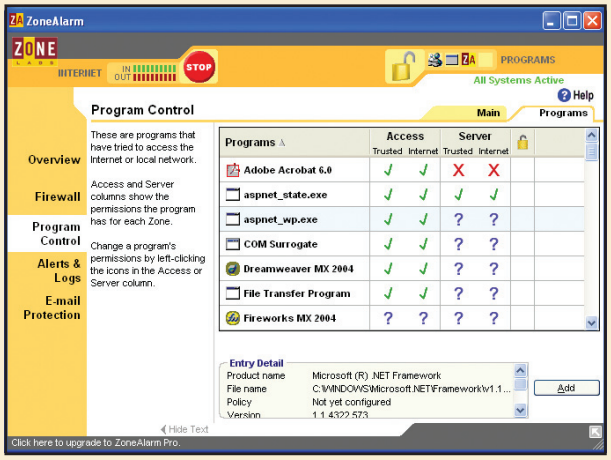

FIGURE 13-15 *The Programs tab of ZoneAlarm's Program Control category reveals the programs to which you have allowed or denied access. The question mark means that the program needs to ask for permission. A green check mark means the program does not need to ask, and a red X means the program is denied. To change a setting, click it to open a popout menu from which you can choose a green check mark, a red X, or a question mark.* ■

Transmitting Network Data Securely

Whenever you transfer data that needs to remain confidential, you should encrypt the data, especially if you are transmitting it over the public Internet. To **encrypt** means to encode the data stream by manipulating the symbols with a set of rules called an algorithm that makes the message appear scrambled and unintelligible. To decipher the data, the person who receives the message must have the **encryption key**, the secret algorithm comprising the rules used to encode the message.

Consider a simple example of an encryption key "123" that shifts each successive character in the message by 1, 2, or 3 characters in the alphabet. So encrypted, the plaintext message *Hello world* appears in cyphertext as *Igomq zptoe*. To decrypt the cyphertext, the person receiving the message needs to have the encryption key. Because "123" is a very simple key, however, a cracker could easily figure it out. In practice, encryption keys are much longer, and the encoding process is so complex that you would need a supercomputer to crack the encryption key.

Symmetric Cryptography and Secret-Key Encryption

Encryption systems are either symmetric or asymmetric. **Symmetric cryptology**, also called secret-key cryptology, uses the same secret key for both encryption and decryption. The most popular symmetric encryption standard is the Data Encryption Standard (DES). Figure 13-16 illustrates the DES algorithm. Even with a supercomputer and a staff of talented cryptologists, cracking such encryption can take many days or weeks. Due to advances in supercomputer technology that crackers may use to discover a DES key, the U.S. government uses triple DES (3DES) encryption, in which a message undergoes the DES process three times. Other symmetric encryption algorithms include RC2, RC4, RC5, and RC6, which Ronald Rivest developed for RSA Security. RC stands for Ron's Code. To learn more about the RC algorithms, go to www.rsasecurity.com/rsalabs/faq. Another popular secret-key algorithm is the International Data Encryption Algorithm (IDEA), a 64-bit iterative block cipher with a 128-bit key. More information about IDEA is at en.wikipedia.org/wiki/International_Data_Encryption_Algorithm.

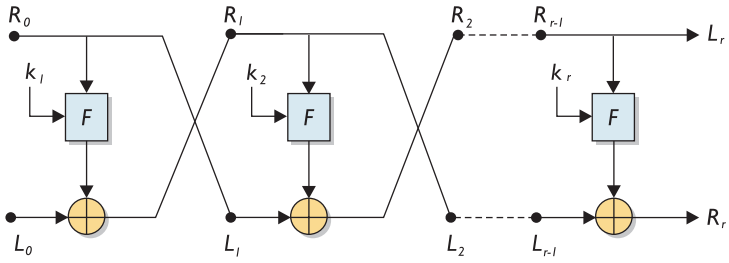

FIGURE 13-16 *DES uses a 16-round Feistel cipher based on a 56-bit encryption key. The block of plaintext being encrypted goes through an iterative process that calculates the ciphertext by repeated application of the Feistel transformation.* ■

The obvious weakness of secret-key cryptology is that the privacy of the information is only as good as the secrecy of the key. If a cracker learns the key by sniffing it in an Internet transmission, for example, or by fooling an insider into divulging it, the mathematical elegance of the encryption algorithm offers little protection.

Inside Info

Advanced Encryption Standard

By the time you read this, the U.S. government will be transitioning to a newer encryption scheme called the Advanced Encryption Standard (AES), which supports key sizes of 128 bits, 192 bits, and 256 bits.

Asymmetric Cryptography and Public Key Infrastructure (PKI)

To avoid the security risk posed by having users share secret keys, you can use a **public key infrastructure (PKI)**, a certificate authority system that assigns to each user a digital certificate containing a key pair consisting of a public key and a private key. The person sending a message uses the public key to encrypt the message, and the person receiving the message uses the private key to decrypt it. Because the key that encrypts the message is different from the key that decrypts it, this process is called **asymmetric cryptography**.

In a PKI, many users within an organization can share the public key. The security comes from the private key, which the user never transmits or shares. Thus, the PKI eliminates the weakness of symmetric cryptography's reliance on a shared secret, because the sender and receiver never transmit a secret key that crackers can sniff. Most PKI implementations use X.509 certificates, following standards issued by the Internet Engineering Task Force's Public Key Infrastructure X.509 (PKIX) working group, whose charter is at www.ietf.org/html.charters/pkix-charter.html.

The most popular form of public-key cryptography is the RSA public-key cryptosystem. RSA stands for Rivest, Shamir, and Adleman, the cryptographers who invented the system. The RSA algorithm works as follows:

> Take two large primes, p and q, and compute their product $n = pq$; n is called the modulus. Choose a number, e, less than n and relatively prime to $(p\text{-}1)(q\text{-}1)$, which means e and $(p\text{-}1)(q\text{-}1)$ have no common factors except 1. Find another number d such that $(ed - 1)$ is divisible by $(p\text{-}1)(q\text{-}1)$. The values e and d are called the public and private exponents, respectively. The public key is the pair (n, e); the private key is (n, d). The factors p and q may be destroyed or kept with the private key.
>
> *Source:* www.rsasecurity.com/rsalabs/node.asp?id=2214

Using mathematical methods known today, obtaining the private key d from the public key (n, e) is extremely difficult. A cracker who could factor n into p and q, however, could obtain the private key d. To prove how hard this is to do, an ongoing contest called the RSA Factoring Challenge offers prize money to anyone who can factor the kinds of large numbers that RSA uses to design its secure cryptosystems. When this book went to press, the prize money ranged from $10,000 for the 576-bit challenge to $200,000 for factoring a number consisting of 2048 bits. For more information, go to www.rsasecurity.com/rsalabs, click Challenges, and follow the link to the new RSA factoring challenge.

Digital Signatures and Hash Encryption

Besides keeping confidential data private, you also need a way to tell whether the message is authentic. In other words, did the message truly come from the person who appears to have sent it, and has the message been altered along the way? Digital signatures provide a way for you to know.

A **digital signature** is an identification method that binds a document to the possessor of a particular key by creating a message digest and encrypt-

ing the digest with the sender's key. When you receive a message that contains a digital signature, Figure 13-17 shows how you can inspect the signature to see where the message came from, who owns the key that signed the message, what certificate authority issued the key, and what algorithms created the signature and the message digest.

A one-way encryption method called **hash encryption** creates the message digest. Hash encryption performs mathematical calculations on the message as a whole and computes a fixed-length number called a **message digest**, which is the message's digital fingerprint. When a digitally signed message arrives at its destination, the receiver's computer employs the same hash algorithm to see if the resulting calculation matches the message digest. The formula is $md = H(m)$, where md is the message digest, H is the hash function, and m is the message. The message can be of any length, but the message digest has a fixed length. It is computationally infeasible to reconstruct the message from its digest. Thus, hash values can verify the in-

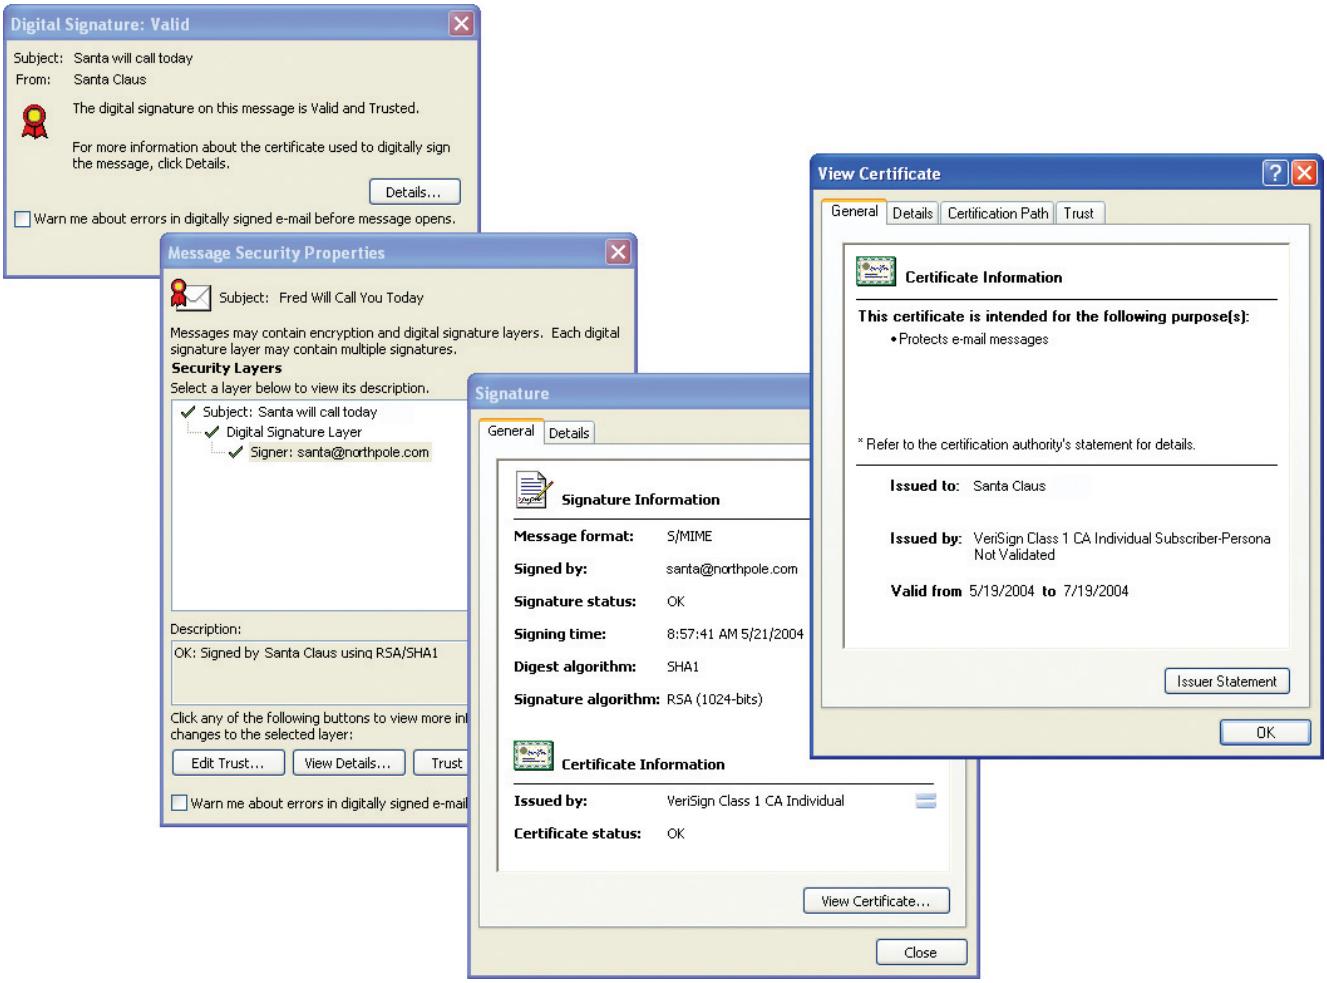

FIGURE 13-17 *When you receive a digitally signed e-mail message and click the digital signature to inspect it, you can see who sent the message and click the Details buttons to reveal the certificate authority that issued the Digital ID, the algorithms used in the signing, and the signing time.* ■

tegrity of a message without transmitting the message itself. The two most commonly used hash encryption algorithms include:

- **SHA-1** Invented by the National Institute of Standards and Technology (NIST), SHA-1 is a corrected version of the Secure Hash Algorithm (SHA), which takes a message up to 2^{64} bits in length and produces a 160-bit message digest.
- **MD5** Rivest, who is the R in RSA, invented the MD2, MD4, and MD5 message digest (MD) algorithms. The latest and greatest version is MD5, which creates a 128-bit message digest using an improved version of the MD4 algorithm. Details and source code for MD2, MD4, and MD5 are in RFC 1319, 1320, and 1321, respectively.

Using a Digital ID with Microsoft Outlook

Microsoft follows the X.509 digital certificate standard throughout its client and server products. If you have Microsoft Outlook, for example, you can get an X.509 certificate consisting of a key pair that enables you to digitally sign your mail and/or send mail encrypted. The certificate authority (CA) recommended by Microsoft is VeriSign, which offers a 60-day free trial, after which the cost of keeping the certificate is $14.95 per year. To get a VeriSign certificate for use with Microsoft Outlook, follow these steps:

1. Pull down the Microsoft Outlook Tools menu and choose Options to bring up the Options dialog.
2. When the Options dialog appears, click the Security tab. Near the bottom of the Security settings, click the button titled Get a Digital ID to bring up a Web page listing the digital ID services recommended by Microsoft. Figure 13-18 shows that these services include VeriSign, which provides digital IDs for secure e-mail.
3. Follow the link to VeriSign, one of the Internet's primary certificate authorities. Figure 13-19 shows that VeriSign offers a 60-day free trial. Click the link to the free trial, and follow the onscreen instructions to get a digital ID. These instructions take you through a four-step process that includes (1) filling out the enrollment form, (2) checking your e-mail for instructions, (3) picking up your digital ID, and (4) installing your digital ID.
4. After you fill out the enrollment form and follow the onscreen instructions to submit it, VeriSign will send you an e-mail message within the hour. Figure 13-20 shows that this message has a Continue button. Click the Continue button to proceed.
5. The final step of the installation process appears as illustrated in Figure 13-21. Click the Install button. A dialog comes onscreen warning you that installing certificates from untrusted Web sites is a security risk. When asked whether you really want to install the certificate, click OK.

note *The term digital ID refers to an X.509 certificate containing a key pair that consists of a public key and a private key.*

note *The purpose of this e-mail verification step is to ensure that you are the person who applied for this digital ID.*

6. To activate your Digital ID for use in sending e-mail, pull down the Outlook Tools menu, choose Options, click the Security tab, check the option to *Add digital signature to outgoing message*, and click the Settings button. When the Change Security Settings dialog appears, use the Choose buttons to select the certificate you want used for signing and encrypting messages. If you have no other certificates, the only choice will be the one you received

Click here to find out about obtaining a digital ID from the VeriSign certificate authority.

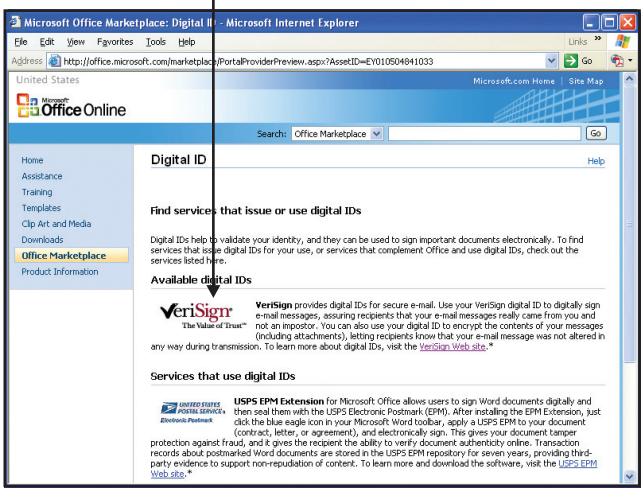

FIGURE 13-18 *Microsoft recommends VeriSign as the certificate authority from which you can obtain a Digital ID consisting of an X.509 certificate comprising a key pair that contains a public key and a private key. You can use this key pair to send encrypted e-mail. You can also sign and seal MS Word documents to prove you are the author and verify that the contents have not been tampered with.* ■

FIGURE 13-19 *The certificate authority VeriSign has a 60-day free trial whereby you can get a Digital ID that you can use to send encrypted e-mail messages.* ■

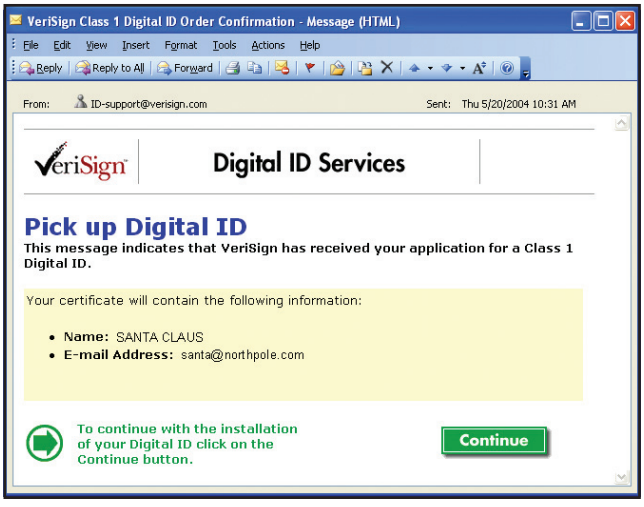

FIGURE 13-20 *Within an hour of submitting the Digital ID application form, VeriSign sends you an e-mail message. When you receive this message, click the Continue button, which brings up the installation screen illustrated in Figure 13-21.* ■

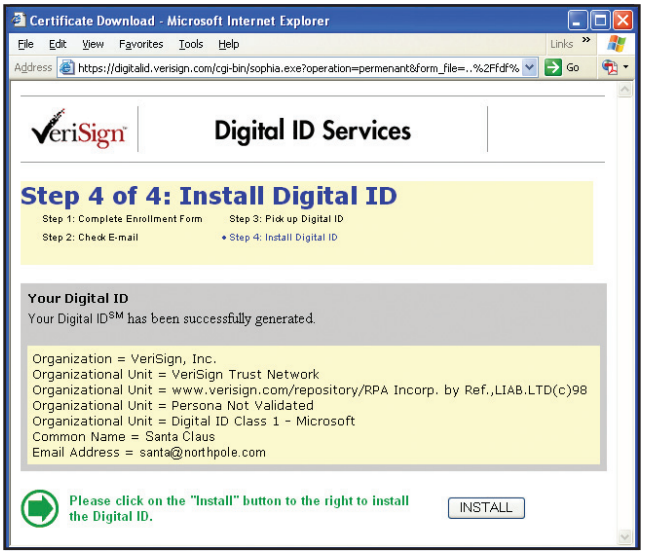

FIGURE 13-21 *The final step in the VeriSign Digital ID procurement process is to click the Install button and follow the onscreen instructions.* ■

when you worked through this exercise. After you finish reviewing the settings, click OK to close the dialogs.

7. Use Outlook to send yourself an e-mail message. When you click the Send button, Outlook informs you that the message is being signed. When you receive the message, you should see a new header named *Signed by* telling you who signed the message.

8. Alongside that header, note the Digital Signature button on the right edge of the message. Click that button to see validation information regarding the signature.

9. When someone who has a key sends you a digitally signed message and you want to add their key to your contact list, right-click their name in the From box and choose Add to Contacts. If the person already has a contact entry, select *Update new information from this contact to the existing one*.

10. To send mail that is encrypted, pull down the Outlook Tools menu, choose Options, click the Security tab, and check the option to Encrypt contents and attachments for outgoing messages. If you then send a message to someone who does not have a key pair, a dialog will warn you that the person does not have a key pair and ask if you want to send the message unencrypted.

11. In the future, if you ever need to change your Digital ID settings or learn more about the things you can do with your Digital ID, go to www.verisign.com/client/guide.

Secure Sockets Layer (SSL) Handshake Protocol

So far, this chapter has presented encryption methods that operate on blocks of information in e-mail messages, for example, or in word-processed documents. Algorithms designed to encrypt blocks of text are called **block ciphers**. There is another important class of algorithms called **stream ciphers**, which operate at the byte (*i.e.*, character) level to encrypt real-time communications. When you use a browser to conduct financial transactions at an online store, for example, the e-commerce site uses a stream cipher to encrypt the session. The most popular stream cipher is RC4, which Ronald Rivest designed for RSA Security.

One of the protocols that use RC4 is the **secure sockets layer (SSL)**, a handshake protocol that defines how a server establishes a secure session in response to an end-user's request to transact. During this first part of the handshake, the server sends its certificate and cipher preferences, which the client uses to create a master key. After encrypting the master key with the server's public key, the client sends the encrypted master key to the server, which authenticates itself to the client via the master key. For the remainder of the session, the client and the server encrypt subsequent communications with keys derived from the master key. Periodically, the server may send a challenge to the client, which authenticates by sending its digital signature and public key certificate in response to the challenge.

SSL supports many encryption algorithms that may be used after the handshake, including RC2, RC4, IDEA, DES, and triple-DES.

Transport Layer Security (TLS) Handshake Protocol

The IETF is working on a successor to SSL called **transport layer security (TLS)**. SSL and TLS have a lot in common. Both run on port 443, the Internet port reserved for secure HTTP sessions. Like SSL, TLS supports RC2, RC4, IDEA, DES, and triple DES. In addition, the IETF is building into TLS support for AES, which will eventually replace DES.

When you visit a Web site running secure over SSL or TLS, the URL in the Web address field begins with *https* instead of http. The letter *s* in https stands for secure. During the secure session, the browser displays the Security icon in the browser's status bar. This icon signals to the user, for example, that confidential forms data and credit card information are being encrypted instead of passing over the Internet in clear text.

The transport layer security working group of the IETF is responsible for creating the https protocols. For the latest on SSL and TLS, visit the IETF working group at www.ietf.org/html.charters/tls-charter.html.

IPSec and Virtual Private Networks (VPN)

The IETF's **Internet Protocol Security (IPSec)** working group is defining a framework of open standards that use cryptography services to ensure private, secure communications over IP networks. Intended to be the future standard for secure Internet communications, IPSec is already the de facto standard. Microsoft, for example, has integrated IPSec into the Active Directory service. Windows 2000, Windows 2003, and Windows XP support the use of IPSec for network-level peer authentication, data origin authentication, data integrity, encryption, and replay protection.

An important application of IPSec is the creation of a **virtual private network (VPN)**, a private data network that makes use of the public Internet's telecommunication infrastructure, maintaining privacy through the use of session keys and an HTTP tunneling protocol over which encrypted data passes. There are two basic types of VPNs. First, you can use a VPN to connect two private networks across the Internet. Second, you can use a VPN to create an extranet in which authorized users who have the necessary keys can access a private network from a remote Internet location that is outside the private net. Figure 13-22 illustrates how the IPSec driver operates at the datagram layer of the Internet by translating IP packets from the Transport Layer into secure IP packets for transmission over the Internet Layer. IPSec's Internet Key Exchange (IKE) framework governs the establishment and management of session keys. For more on IPSec and the RFCs it comprises, go to the Virtual Private Network Consortium (VPNC) at www.vpnc.org.

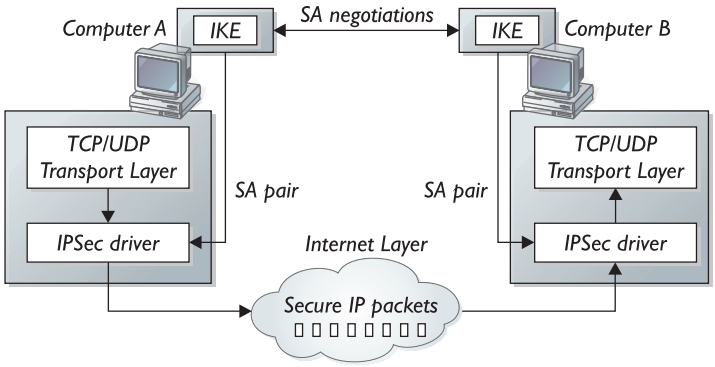

FIGURE 13-22 *The IPSec process consists of an Internet key exchange (IKE) through which two computers negotiate the keys that will encrypt the session, and an IPSec driver that translates IP packets from the Transport Layer into secure IP packets for transmission over the Internet Layer.* ■

Pretty Good Privacy (PGP)

Invented by Phil Zimmerman in 1991, **Pretty Good Privacy (PGP)** is a data integrity system that uses encryption, data compression, and digital signatures to provide for the secure transmission of e-mail messages and other kinds of store-and-forward file systems. PGP can work with many different algorithms, including Elgamal or RSA for key transport, IDEA or CAST5 for block data encryption, DSA or RSA for signing, and SHA-1 or MD5 for computing message digests. The shareware program ZIP compresses the messages to conserve bandwidth in data transmission and to save space in file storage. The PGP corporate Web site at www.pgp.com sells commercial versions of PGP in corporate, workgroup, and personal desktop versions. The IETF has an OpenPGP working group that is creating an open specification for Pretty Good Privacy at http://www.ietf.org/html.charters/openpgp-charter.html.

Pretty Good Privacy is actually very good. Perhaps the best testimonial is Microsoft's use of PGP signatures to verify the authenticity of messages that come from the Microsoft Security Notification Service. You can download Microsoft's public PGP key from www.microsoft.com/technet/security/bulletin/notify.mspx.

Publishing a Web Securely

As you learned in Chapter 1, FTP is the protocol for transferring files over the Internet. One of the most common uses of FTP is to publish files to the Web. You learned how to do this in the Web publishing tutorial at the end of Chapter 7. Many people do not realize, however, that when you log on to a Web site by using an ordinary FTP client, your user name and password traverse the Internet in plain text. This is not very secure, because a cracker can use a packet sniffer to see your password in clear text. To prevent this from happening, you should have a secure FTP (SFTP) program that uses the secure shell (SSH) protocol to prevent your password and data from passing over the Internet in plain text.

Using a Secure Shell (SSH) Protocol

The **secure shell (SSH)** protocol enables two computers to negotiate and establish a secure connection that uses encryption to thwart crackers who may try to sniff passwords and data that would otherwise traverse the Internet in clear text. Through a process called **tunneling**, other kinds of TCP/IP connections can funnel through the SSH connection, which provides a secure communication channel for doing mail, accessing the Web, logging on securely to remote sites, and publishing files via SFTP.

The IETF's secure shell working group is in charge of standardizing the SSH protocols. You can read the charter at www.ietf.org/html.charters/secsh-charter.html.

Making a Secure FTP (SFTP) File Transfer

Searching Google or Yahoo for the keywords *sftp client* brings up many SFTP clients you can use to publish a Web securely. You will be happy to discover that the Core FTP client you learned in Chapter 7 supports SSH. If the Web server to which you transfer your files also supports SSH, you can use Core FTP to make an SFTP transfer. To make Core FTP use SSH, follow these steps:

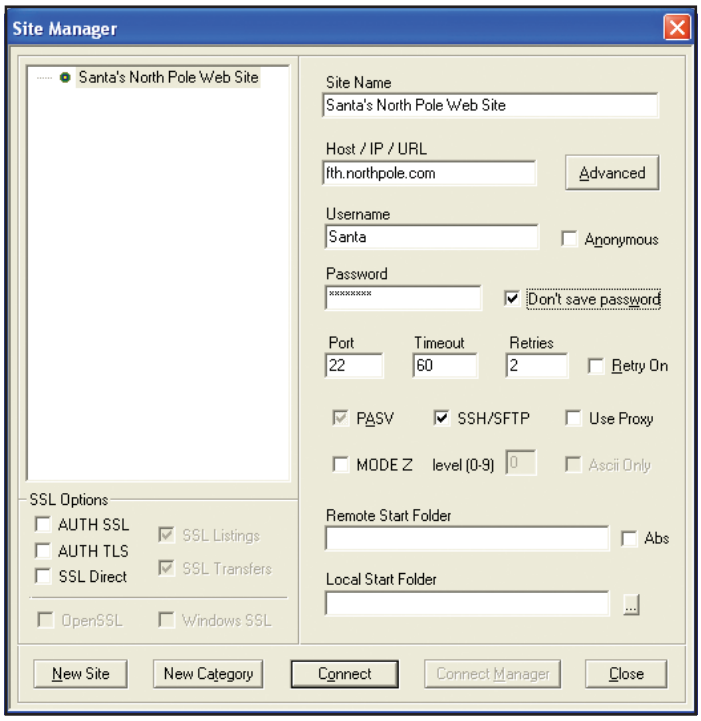

FIGURE 13-23 *To make the Core FTP program use the SSH protocol, check the box titled SSH/SFTP. When you check this box, the port automatically changes from 21 to 22, which is the SFTP port number.* ■

1. Click the Windows Start button and choose Programs | Core FTP | Core FTP. In the Site Manager window, select the site to which you want to connect.
2. Near the middle of the Site Manager window, check the box titled SSH/SFTP. Figure 13-23 shows that when you do this, the port number will change from 21 to 22, which is the SFTP port number.
3. Click the Connect button to connect to the FTP site. If the site supports SFTP, you will get a message asking if you trust the host and want to add its key to your SSH cache. Click Yes to store the key and continue. By storing this key you will get a warning if the remote key ever changes, which may clue you to a cracker trying a **man in the middle (MITM) attack** in which a cracker tries to hijack your session.
4. From now on, your Core FTP sessions at this site will use the SFTP protocol.

Having learned how to make a secure FTP (SFTP) transfer, you may now conclude this book with a clear conscience.

Chapter 13 Review

Chapter Summary

After reading this chapter and completing the step-by-step tutorials and Try This! exercises, you should understand the following facts about Internet security:

Identifying Internet Security Issues

- Internet security risks fall into three general categories: (1) unauthorized access, (2) data manipulation, and (3) service interruption.
- About two-thirds of all security breaches come from inside an organization's network architecture. To help guard against inside attacks, you should locate equipment behind locked doors and limit access to authorized personnel.
- A social engineering attack exploits human weaknesses to gain access to the organization's network.
- Depending on the nature of the organization, some security threats will be more of an issue than others. Because it is impossible to prepare for all possible attacks, identifying assets helps prioritize the need for different kinds of defenses.
- Viruses top the list of attacker's tools. Viruses propagate by attaching themselves to software programs or other kinds of files. When a user executes or opens a file thus infected, the user's computer catches the virus.
- Another kind of malicious code is the worm, which can propagate across the Internet and infect other computers without attaching itself to other programs. Worms can wreak havoc by replicating themselves on an infected system to tie up all its resources and then following that system's network connections to attack other computers.

Applying Internet Security Safeguards

- For millions of end users and server administrators who have the Windows operating system, subscribing to one or more of the Microsoft Security Newsletters is an excellent way to keep up with Microsoft's security strategies and guidance.
- Microsoft runs a service called Windows Update that can automatically download the latest security patches to your computer. You can configure the service to install the patches right away or to notify you when the patches are ready to install.
- The most frequent kind of attack is Denial of Service (DoS), in which the attacker seeks to consume so much of a server's resources that the host cannot respond to legitimate requests.
- In a so-called brute force or front door attack, a cracker programs a computer to look up words in a dictionary and generate variants with which the computer keeps trying to log on until it discovers a password that gets in.
- Another security risk created by well-intentioned IT staff is the trapdoor attack, in which crackers find a way into your computer by running diagnostic tools that your IT staff left on the system after troubleshooting some kind of problem.
- Security auditing is a proactive process that considers the risks associated with security assets, predicts the methods crackers may use to exploit each risk, and takes protective steps to thwart them.
- A firewall is a combination of hardware, software, and security policies that block certain kinds of traffic from entering or leaving an entire network, subnet, or individual host computer. Depending on the configuration, the firewall has one of four topologies: (1) packet filtering firewall, (2) single-homed bastion host firewall, (3) dual-homed bastion host firewall, or (4) screened subnet firewall with DMZ.

Transmitting Network Data Securely

- To encrypt means to encode the data stream by manipulating the symbols with a set of rules called an algorithm that makes the message appear scrambled and unintelligible. To decipher the data, the person who receives the message must have the encryption key, the secret algorithm comprising the rules used to encode the message.

- Symmetric cryptology, also called secret-key cryptology, uses the same secret key for both encryption and decryption.
- To avoid the security risk posed by having users share secret keys, you can use a public key infrastructure (PKI), a certificate authority system that assigns to each user a digital certificate containing a key pair consisting of a public key and a private key. The person sending a message uses the public key to encrypt the message, and the person receiving the message uses the private key to decrypt it. Because the key that encrypts the message is different from the key that decrypts it, this process is called asymmetric cryptography.
- A digital signature is an identification method that binds a document to the possessor of a particular key by creating a message digest and encrypting the digest with the sender's key. A one-way encryption method called hash encryption creates the message digest.
- Microsoft follows the X.509 digital certificate standard throughout its client and server products. The certificate authority (CA) recommended by Microsoft is VeriSign, which refers to digital certificates via the term Digital ID.
- The Secure Sockets Layer (SSL) is a handshake protocol that defines how a server establishes a secure session in response to an end user's request to transact. SSL supports many encryption algorithms that may be used after the handshake, including RC2, RC4, IDEA, DES, and triple-DES. The IETF is working on a successor to SSL called transport layer security (TLS).
- The IETF's Internet Protocol Security (IPSec) working group is defining a framework of open standards that use cryptography services to ensure private, secure communications over IP networks. An important application of IPSec is the creation of a virtual private network (VPN), a private data network that makes use of the public Internet's telecommunication infrastructure, maintaining privacy through the use of session keys and an HTTP tunneling protocol over which encrypted data passes.
- Invented by Phil Zimmerman in 1991, Pretty Good Privacy (PGP) is a data integrity system that uses encryption, data compression, and digital signatures to provide for the secure transmission of e-mail messages and other kinds of store-and-forward file systems. Microsoft uses PGP signatures to verify the authenticity of messages that come from the Microsoft Security Notification Service.

Publishing a Web Securely

- Many people do not realize that when you log on to a Web site by using an ordinary FTP client, your password traverses the Internet in plain text. To keep this from happening, you should have an FTP program that uses the secure shell (SSH) protocol to prevent your password and data from passing over the Internet in plaintext.
- The Core FTP client supports the SFTP protocol. To make Core FTP use the SFTP protocol, check the SSH/SFTP box in the Site Manager.

Key Terms

application level gateway *(13)*
asymmetric cryptography *(19)*
bastion host *(14)*
block cipher *(23)*
boot record virus *(5)*
brute force attack *(10)*
circuit level gateway *(12)*
demilitarized zone (DMZ) *(15)*
Denial of Service (DoS) *(10)*
digital signature *(19)*
dual-homed bastion *(14)*
encrypt *(18)*
encryption key *(18)*
file infector virus *(6)*
firewall *(12)*
front door attack *(10)*
hash encryption *(20)*
inside attack *(3)*
IP address spoofing *(4)*
Internet Protocol Security (IPSec) *(24)*
macro virus *(6)*
man in the middle (MITM) attack *(26)*
masquerading *(4)*
message digest *(20)*
packet filter *(13)*
Pretty Good Privacy (PGP) *(25)*
public key infrastructure (PKI) *(19)*
replay attack *(4)*
screened-subnet firewall *(15)*
secure shell (SSH) *(25)*

secure sockets layer (SSL) *(23)*
security auditing *(11)*
single-homed bastion *(14)*
social engineering attack *(4)*
stateful inspection *(13)*
stream cipher *(23)*
symmetric cryptology *(18)*
transport layer security (TLS) *(24)*
Trojan horse *(6)*
tunneling *(25)*
virtual private network (VPN) *(24)*
Windows Update Service *(9)*
worm *(6)*

Key Terms Quiz

1. _____ is a masquerade that happens when an intruder uses the IP address of a trusted system to gain access rights granted to that system.
2. A(n) _____ is malicious code that can propagate across the Internet and infect other computers without attaching itself to other programs.
3. In a brute force or _____, a cracker programs a computer to look up words in a dictionary and generate variants with which the computer keeps trying to log on until it discovers a password that gets in.
4. _____ is a proactive process that considers the risks associated with security assets, predicts the methods crackers may use to exploit each risk, and takes protective steps to thwart them.
5. The simplest kind of firewall uses a(n) _____, which monitors the headers of all incoming or outgoing packets and can block transmissions based on source or destination ports or IP addresses.
6. A(n) _____ is a computer that sits on the perimeter of a local network and serves as an application level gateway between the external network and the internal client workstations.
7. The _____ topology establishes a demilitarized zone (DMZ) by placing packet filtering routers on both the Internet side and the private network side of the application gateway.
8. _____, also called secret-key cryptology, uses the same secret key for both encryption and decryption.
9. To avoid the security risk posed by having users share secret keys, you can use a(n) _____, a certificate authority system that assigns to each user a digital certificate containing a key pair consisting of a public key and a private key.
10. The _____ is a protocol that enables two computers to negotiate and establish a secure connection that uses encryption to thwart crackers who may try to sniff usernames, passwords, and data that would otherwise traverse the Internet in clear text. Through a process called tunneling, other kinds of TCP/IP connections can funnel through this connection, which provides a secure communication channel for doing mail, accessing the Web, logging on securely to telnet sites, and publishing files via FTP.

Multiple-Choice Quiz

1. When a cracker uses a packet sniffer to record a logon sequence to play back at a later time in an attempt to gain access, what kind of an attack is it?
   a. Brute force
   b. Inside
   c. Masquerade
   d. Replay
2. What kind of virus runs code from the startup device when you turn the power on?
   a. Boot record
   b. File infector
   c. Macro
   d. Trojan horse
3. Which safeguard can download and install security patches automatically?
   a. Microsoft Security Notification Service
   b. Passport
   c. Process Explorer
   d. Windows Update Service

4. Which firewall strategy can detect a hijack by keeping track of when a port opens, what session is using it, and how long it stays open?
   a. Application level gateway
   b. Circuit level gateway
   c. Network address translation
   d. Stateful inspection
5. Which firewall topology uses two NICs on which IP forwarding is disabled, thereby creating a complete physical break between the internal and external networks?
   a. Dual-homed bastion
   b. Packet filtering
   c. Screened subnet
   d. Single-homed bastion
6. Which one of the following applications is popular as a personal firewall product?
   a. IPSec
   b. Process Explorer
   c. SSH secure shell
   d. ZoneAlarm
7. Which of the following is a hash algorithm used for one-way encryption?
   a. AES
   b. DES
   c. MD5
   d. RC4
8. What protocol does X.509 define?
   a. Digital certificates
   b. SSL
   c. SFTP
   d. SSH
9. What kind of network can you create by using the IPSec protocol to create an Extranet in which authorized users who have the necessary keys can access a private network from a remote Internet location that is outside the private net?
   a. IKE
   b. SSH
   c. STP
   d. VPN
10. Which of the following products is a data integrity system that uses encryption, data compression, and digital signatures to provide for the secure transmission of e-mail messages and other kinds of store-and-forward file systems?
    a. DES
    b. PKI
    c. PGP
    d. SSL

Essay Quiz

1. Explain how a network administrator could have used a packet filtering firewall to stop Slammer, which attacked through the Microsoft SQL resolution service that runs on port 1434.
2. Describe the architectural difference between single-homed and dual-homed bastion firewall topologies. In particular, describe the nature of the second layer of defense you can achieve by using two NICs instead of one inside the bastion host.
3. Explain why you must warn users about the danger of following links to non-trusted sites. In particular, explain the Web page technology through which a cracker can run malicious code on the user's computer.
4. How does a public key infrastructure (PKI) avoid the security risk posed by having users share secret keys? If the key is public, what prevents a cracker from using the public key to decrypt the message?
5. What does it mean when a Web address begins with https instead of http? Tell what the *s* in https stands for, and explain the handshaking that goes on behind the scenes.

Lab Projects

• Lab Project 13-1: Choosing a Virus Scanner

Schools and companies can lose a lot of time and money when viruses strike. It is critically important for both the servers and the client computers in your workplace to be protected from viruses. Imagine that you work for a school or company that has recently undergone a bad virus attack. Your employer wants to prevent such an attack from happening again. Your employer has asked you to recommend the brand of virus scanner that should be installed on all the machines at your workplace. You have also been asked to look into protecting the computers your fellow employees have at home, to minimize the risk that employees might inadvertently transmit to the workplace a virus from their home computer. In adopting a virus scanner for use in your school or company, consider these issues:

- Dangerous viruses can spread quickly across the Internet. The virus scanner you recommend should have an update service that automatically updates the virus definitions when new viruses come on the Net.
- Home computers need to be protected as well as machines in the workplace. If school children are using an employee's computer at home, for example, viruses from school can be transmitted to the employee's home computer, from which the infection could propagate to the workplace.
- If there is a mix of Windows and Macintosh machines in the workplace and in coworker homes, you need to consider virus protection for both brands of operating systems. Also consider other operating systems that may be used on your workplace network.
- Viruses can be caught both coming and going. Consider whether the virus scanner you are considering can scan outgoing as well as incoming messages.
- Consider all the ways information can come and go, including E-mail, IM, FTP, and peer-to-peer file sharing. Check to see whether the virus scanner you are proposing scans all these ways of transmitting viruses.
- New ways of transmitting viruses may have been discovered or invented since this book went to press. Check the virus alert centers at www.sarc.com and www.mcafee.com to see if any new transmission modes have arisen.

Use a word processor to write up your virus scanner recommendation in the form of a brief essay. Report the brand names of the virus scanners you considered, identify the one you recommend for adoption in your workplace, and explain why you selected it instead of the others. If your instructor has asked you to hand in the report, make sure you put your name at the top; then save it on disk or follow the other instructions you may have been given for submitting this assignment.

• Lab Project 13-2: Determining Network Vulnerabilities

In a series of Microsoft white papers titled "Best Practices for Enterprise Security," Benson published a framework for determining network vulnerabilities. The framework consists of a series of questions organized according to the three categories of (1) physical security, (2) data security, and (3) network security. Imagine that your employer has asked you to use the Benson framework in determining the vulnerabilities of your school or workplace network. Use your word processor to write an essay in which you answer the questions in the Benson framework and make recommendations for shoring up vulnerabilities you uncover. The questions to answer are as follows:

Category I: Physical Security

1. Are there locks and entry procedures to gain access to servers?
2. Is there sufficient air conditioning and are air filters being cleaned out regularly? Are air conditioning ducts safeguarded against break-ins?
3. Are there uninterruptible power supplies and generators, and are they being checked through maintenance procedures?
4. Is there fire suppression and pumping equipment and proper maintenance procedures for the equipment?
5. Is there protection against hardware and software theft? Are software packages and licenses and backups kept in safes?
6. Are there procedures for storing data, backups, and licensed software off-site and onsite?

Category II: Data Security

7. What access controls, integrity controls, and backup procedures are in place to limit attacks?
8. Are there privacy policies and procedures to which users must comply?
9. What data access controls (authorization, authentication, and implementation) are there?
10. What user responsibilities exist for management of data and applications?
11. Have direct access storage device management techniques been defined? What is their impact on user file integrity?
12. Are there procedures for handling sensitive data?

Category III: Network Security

13. What kinds of access controls (Internet, wide area network connections, etc.) are in place?
14. Are there authentication procedures? What authentication protocols are used for local area networks, wide area networks, and dialup servers? Who has the responsibility for security administration?
15. What types of network media (e.g., cables, switches, and routers) are used? What type of security do they have?
16. Is security implemented on file and print servers?
17. Does your organization make use of encryption and cryptography for use over the Internet, Virtual Private Networks (VPNs), e-mail systems, and remote access?
18. Does the organization conform to networking standards?

If your instructor asked you to hand in your answers to these questions, make sure you put your name at the top of the essay; then copy it onto a disk or follow the other instructions you may have been given for submitting this assignment.

Note: *The full text of the Benson white paper that contains these questions is at www.microsoft.com/technet/security/bestprac/bpent/sec1/secstrat.mspx. More security resources are at www.microsoft.com/technet/security.*

Main text file is embedded
because of difficulty w/ keyword references.
PLS. always export the text after done with the chapter.
(File - Export Text ...)

PLEASE PLACE THIS NOTE AT END OF CHAPTER,
AFTER BUILDING AS A REMINDER
DURING CORRECTIONS