Chapter 1: Defending against Common Attacks with Basic Security Tools

In This Chapter

- **Reviewing internal network attacks**
- **Identifying external attack methods**
- **Implementing firewall rules to protect your network**
- **Working with other protection tools**

When dealing with security on your Cisco network, you can focus on a number of areas. In this chapter, you look at high-level devices in the form of firewall devices and proxies. In many cases, your firewall device may have proxy components built into it, or the proxy components will operate on a separate device or host. In my discussion on firewalls, I focus on the features found in the Cisco Adaptive Security Appliance (ASA) because they are fairly common when compared with the other devices on the market. (You find out how to manage and configure the ASA in Chapter 2 of this minibook.) You also look at the types of attacks to which you may be susceptible on your network. (For the specific functions of Network Address Translation [NAT] and Access Control Lists [ACLs], see Chapter 3 in this minibook.)

So sit back and review the types of attacks that the bad guys will likely attempt on your network. I follow that up with steps and tools you can use to prevent these attacks or to remediate the damage caused by such attacks. Remember, although these tools and attacks may be carried out by others with the intent to steal information or harm your network, you may also want to use the same tools and attacks as part of a network audit.

COPYRIGHTED MATERIAL

Knowing Your Enemy

In *The Art of War,* Sun Tzu said that "if you know your enemies and know yourself, you will not be imperiled in a hundred battles." By this, he means that the more you know about your enemies and how they will operate, the better you can avoid being drawn into unwanted battles. Knowing what your enemy is going to try provides you the opportunity to protect yourself. So in these sections, I show you some of the things that crafty (or sometimes not-so-crafty) enemies may try.

Two main types of attacks take place on a network: those that are run from inside the network and those that try to make their way in from the network's perimeter. I start by going through the most common internal attacks, and then I move on to describe common external attacks.

Handling attacks from within

Although everyone wants to trust the people that they work with, a large number of attacks occur from within your network. These attacks may be from employees or from non-employees who are in your building, and on your network. Although much of the focus on security deals with the perimeter of your network and the access points, you must not forget about the inside of your network and what you can do to defend yourself after the attacker is inside. The most common types of internal attacks are packet sniffing, man in the middle, cached credentials, masquerade, and network scanning. The following sections look at each of these attacks and what you can do to defend yourself.

Packet sniffing

Packet sniffing captures network traffic at the Ethernet frame level. After capture, this data can be analyzed and sensitive information can be retrieved. Such an attack starts with a tool such as Wireshark (described in Book I, Chapter 4). Wireshark allows you to capture and examine data that is flowing across your network. Any data that is not encrypted is readable, and unfortunately, many types of traffic on your network are passed as unencrypted data — even passwords and other sensitive data. Obviously, this situation represents a danger to your corporate data. Many applications that house corporate data (even those with slick Windows-based GUIs) still use Telnet as the data transfer mechanism. *Telnet* is a clear text, unencrypted data transfer mechanism. A person with a packet sniffer can view this data as it crosses your network. Figure 1-1 shows FTP logon data captured behind the FTP window, showing the user's password. Having your FTP password known allows the attacker to have your level of access to your FTP site, and any secret data that may be there; on top of that, many users who use the same password for all systems on the network. Now the attacker may have access to several of your corporate systems.

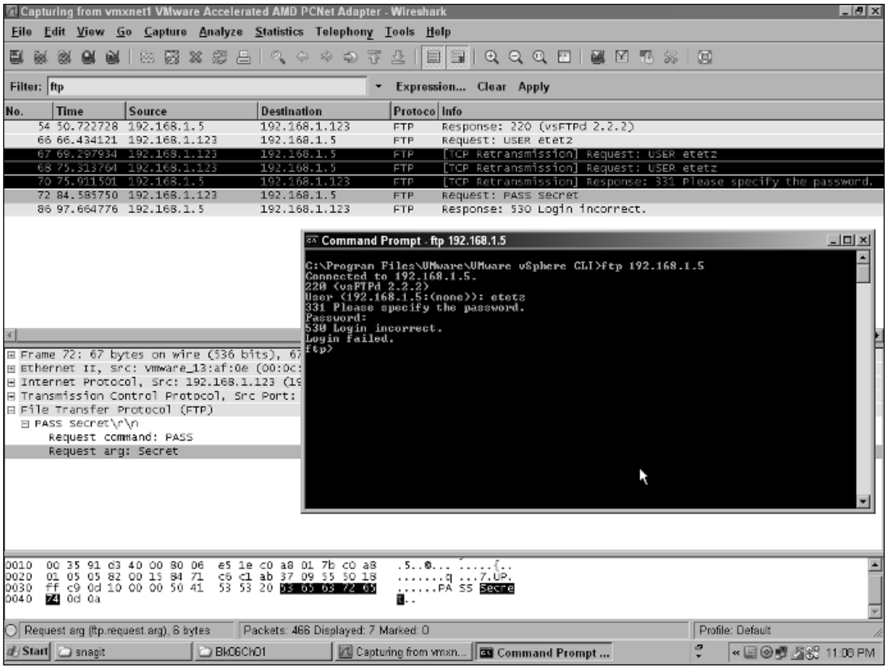
Capturing from vmxnet1 VMware Accelerated AMD PCNet Adapter - Wireshark
File Edit View Go Capture Analyze Statistics Telephony Tools Help
Expression... Clear Apply
Command Prompt - ftp 192.168.1.5
Start
snagit
BkO6ChO1
Capturing from vmxn...
Command Prompt ...
Profile: Default
11:08 PM

Filter: ftp

```
No.  Time         Source          Destination     Protoco Info
  54 50.722728    192.168.1.5     192.168.1.123   FTP     Response: 220 (vsFTPd 2.2.2)
  66 66.434121    192.168.1.123   192.168.1.5     FTP     Request: USER etetz
  67 69.297934    192.168.1.123   192.168.1.5     FTP     [TCP Retransmission] Request: USER etetz
  68 75.313764    192.168.1.123   192.168.1.5     FTP     [TCP Retransmission] Request: USER etetz
  70 75.911501    192.168.1.5     192.168.1.123   FTP     [TCP Retransmission] Response: 331 Please specify the password.
  72 84.585750    192.168.1.123   192.168.1.5     FTP     Request: PASS secret
  86 97.664776    192.168.1.5     192.168.1.123   FTP     Response: 530 Login incorrect.
```

```
C:\Progran Files\UMware\UMware vSphere CLI>ftp 192.168.1.5
Connected to 192.168.1.5.
220 (vsFTPd 2.2.2)
User (192.168.1.5:(none)): etetz
331 Please specify the password.
Password:
530 Login incorrect.
Login failed.
ftp>
```

```
⊞ Frame 72: 67 bytes on wire (536 bits), 67
⊞ Ethernet II, Src: Vmware_13:af:0e (00:0c:
⊞ Internet Protocol, Src: 192.168.1.123 (19
⊞ Transmission Control Protocol, Src Port:
⊟ File Transfer Protocol (FTP)
  ⊟ PASS secret\r\n
      Request command: PASS
      Request arg: Secret
```

```
0010  00 35 91 d3 40 00 80 06  e5 1e c0 a8 01 7b c0 a8   .5..@... .....{..
0020  01 05 05 82 00 15 84 71  c6 c1 ab 37 09 55 50 18   .......q ...7.UP.
0030  ff c9 0d 10 00 00 50 41  53 53 20 53 65 63 72 65   ......PA SS Secre
0040  74 0d 0a                                           t..
```

Request arg (ftp.request.arg), 6 bytes | Packets: 466 Displayed: 7 Marked: 0

Figure 1-1: FTP logon data in a packet capture.

In addition to capturing cleartext sessions, such as login traffic, an attacker can have an application that captures only specific data from a network, such as network authentication packets, which she then reviews to crack network passwords.

If you are using switch-based network, you make packet sniffing a little tougher. On a switch-based network, the sniffer will see only data going to and from the sniffer's own network device or broadcast traffic, unless the attacker uses a monitoring port on a switch. If you have not secured your switches and your switch configuration documentation with a strong password, you are leaving yourself open to a packet-sniffing attack.

A packet-sniffing attack on a switch-based network happens like this: The attacker connects to a switch and uses information from that switch to locate his own MAC address. The attacker locates his MAC address via `show address-database`, which lets him know what port the address is seen on. The attacker can follow the path until he finds the switch to which he is connected. From there, the attacker can enable a monitor port as the port to which he connected. Now he can see all the traffic on that switch and can start a packet capture of data.

REMEMBER

Switch security is the first line of your network security from internal hacking. Switch security is the path attackers must go through to get to the rest of your network. If you can keep attackers from connecting or restrict their ability to gain sensitive information, you beat them.

Man in the middle

Man in the middle is a type of masquerade attack. A man-in-the-middle attack works like this: If the attacker places herself between you and the server to which you are talking, the attacker can see all the data (encrypted or not) that you are sending to the server. This particular attack is very disconcerting because they can easily see the data that you expect is 100-percent secure, even your HTTPS dealings with your bank.

In this attack, the attacker takes over the role of a device between you and the system you are talking to. This device could be a router, where the attacker confuses the switch ARP table and has data destined for the router to be sent to her. Then she relays the data to the router. In this way, the attacker can still deal with the router and the server on the other side of the router, but the attacker sees all the traffic. This setup allows an attacker to capture passwords, even for secure sites, such as banking.

Tools that can conduct this type of attack are freely available. One such tool is Cain & Abel from `www.oxid.it`.

WARNING!

Man in the middle is one of the most insidious attacks, because you may not even know it is happening. For this reason, any unsecured network should be considered hostile or even broken.

Cracking cached credentials

Another source for attackers to gain access to system passwords and sensitive data is right on your workstations. Many users leave their computers unattended for periods of time. Even if the computer is powered off, an attacker can boot from a USB device and access your entire hard drive, or if left on and unlocked, he can do whatever he wants. One set of files attackers are often after hold the credential cache. That file contains all the passwords you have told Windows to save on the system.

Cain & Able has the unique ability to extract data from your system, such as passwords that are stored in a variety of password caches. Figure 1-2 shows a list of server passwords that have been cached on the local workstation. Attackers accomplish this type of data retrieval by reversing the hashing (or encrypting) process that hid the password in the first place.

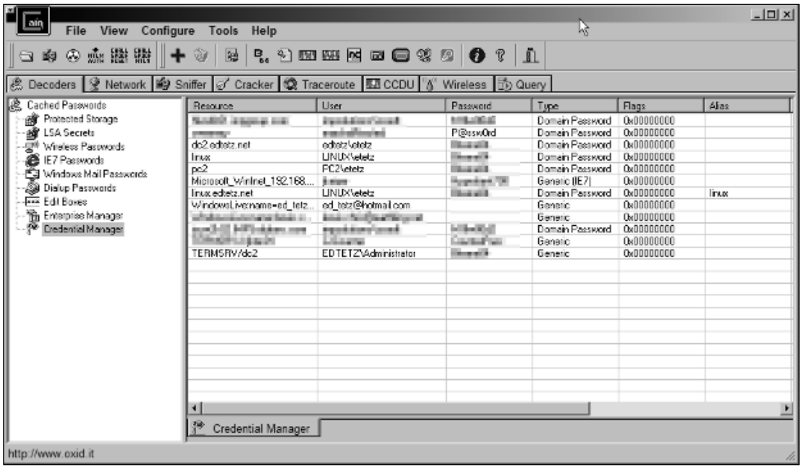

Resource	User	Password	Type	Flags	Alias
[redacted]	[redacted]	[redacted]	Domain Password	0x00000000	
[redacted]	[redacted]	P@ssw0rd	Domain Password	0x00000000	
dc2.edtetz.net	edtetz\etetz	[redacted]	Domain Password	0x00000000	
linux	LINUX\etetz	[redacted]	Domain Password	0x00000000	
pc2	PC2\etetz	[redacted]	Domain Password	0x00000000	
Microsoft_WinInet_192.168....	[redacted]	[redacted]	Generic (IE7)	0x00000000	
linux.edtetz.net	LINUX\etetz	[redacted]	Domain Password	0x00000000	linux
WindowsLive:name=ed_tetz...	ed_tetz@hotmail.com		Generic	0x00000000	
[redacted]	[redacted]		Generic	0x00000000	
[redacted]	[redacted]	[redacted]	Domain Password	0x00000000	
[redacted]	[redacted]	[redacted]	Generic	0x00000000	
TERMSRV/dc2	EDTETZ\Administrator	[redacted]	Generic	0x00000000	

Figure 1-2: A list of passwords retrieved from the local system by Cain.

Masquerade

A *masquerade attack* is a type spoofing attack where the attacker pretends to be someone or some device which he is not. (The man-in-the-middle attack can be considered a masquerade attack because the man in the middle pretends to be a router or some other type of middle device.) E-mail addresses, URLs, and network devices such as routers can all be spoofed. Masquerade attacks often succeed because people see what they expect to see.

One effective masquerade is to create a fake Windows server. Clients on that network automatically attempt to authenticate to this fake Windows server with their current logon credentials. A client does this authentication by accepting a random challenge word from the server, encrypting it using her password as the encryption key, and sending that newly encrypted string to the server. The attacker running the masquerade server knows what word was originally sent as the challenge, so he can compare that encrypted string with the string he gets from a series of password attempts. When the attacker finds a matching string, he knows the password.

Network scanning

Network scanning allows you to find out what systems are on your network, what services they may be offering, and sometimes a fair bit more than that, such as services with known vulnerabilities or systems that the IT staff thought were removed from the network years ago. One of the most common general purpose network scanners is Nmap, or network map, with its Windows-based Zenmap, available from `http://nmap.org`. From the attack perspective, this tool is part of most attacker's information-gathering arsenal. With a list of systems, operating systems, and running services, she can pick the weakest members of your network herd.

As an internal auditing tool, I regularly use Zenmap to verify available IP addresses on a network. By providing Zenmap a network ID and few seconds, it can provide you with a list of used IP addresses, matching MAC addresses, DNS names for those systems, open ports on those systems, and even the OS type for the hosts that it has found. The following code is an example of the type of information you can see from a Zenmap or an Nmap scan of a system. It discovered the following:

- This is an Ubuntu Linux computer.
- This machine shares files out to Windows-based computers.
- This machine hosts a website.
- This machine is running VMware Server.
- This host supports SSH and VNC as remote access methods.
- This host is running a mail server and an FTP server.

```
Starting Nmap 5.21 ( http://nmap.org ) at 2011-04-15 02:01 Atlantic Daylight Time
NSE: Loaded 36 scripts for scanning.
Initiating ARP Ping Scan at 02:01
Scanning 192.168.1.5 [1 port]
Completed ARP Ping Scan at 02:01, 0.30s elapsed (1 total hosts)
Initiating Parallel DNS resolution of 1 host. at 02:01
Completed Parallel DNS resolution of 1 host. at 02:01, 0.00s elapsed
Initiating SYN Stealth Scan at 02:01
Scanning 192.168.1.5 [1000 ports]
Discovered open port 445/tcp on 192.168.1.5
Discovered open port 111/tcp on 192.168.1.5
Discovered open port 5900/tcp on 192.168.1.5
Discovered open port 53/tcp on 192.168.1.5
Discovered open port 21/tcp on 192.168.1.5
Discovered open port 80/tcp on 192.168.1.5
Discovered open port 22/tcp on 192.168.1.5
Discovered open port 25/tcp on 192.168.1.5
Discovered open port 443/tcp on 192.168.1.5
Discovered open port 139/tcp on 192.168.1.5
Discovered open port 8222/tcp on 192.168.1.5
Discovered open port 902/tcp on 192.168.1.5
Discovered open port 8009/tcp on 192.168.1.5
Discovered open port 8333/tcp on 192.168.1.5
Discovered open port 1984/tcp on 192.168.1.5
Discovered open port 2049/tcp on 192.168.1.5
Completed SYN Stealth Scan at 02:01, 1.53s elapsed (1000 total ports)
Initiating Service scan at 02:01
Scanning 16 services on 192.168.1.5
Completed Service scan at 02:03, 116.14s elapsed (16 services on 1 host)
Initiating RPCGrind Scan against 192.168.1.5 at 02:03
Completed RPCGrind Scan against 192.168.1.5 at 02:03, 0.03s elapsed (2 ports)
Initiating OS detection (try #1) against 192.168.1.5
NSE: Script scanning 192.168.1.5.
```

```
NSE: Starting runlevel 1 (of 1) scan.
Initiating NSE at 02:03
Completed NSE at 02:03, 25.06s elapsed
NSE: Script Scanning completed.
Nmap scan report for 192.168.1.5
Host is up (0.0014s latency).
Not shown: 984 closed ports
PORT     STATE SERVICE         VERSION
21/tcp   open  ftp             vsftpd 2.2.2
22/tcp   open  ssh             OpenSSH 5.3p1 Debian 3ubuntu4 (protocol 2.0)
| ssh-hostkey: 1024 5b:6d:35:57:65:42:7f:8a:73:7e:00:e3:89:f9:15:bf (DSA)
|_2048 4d:6e:be:c4:3b:0c:55:f5:46:dd:b8:05:05:1c:94:ea (RSA)
25/tcp   open  smtp            Exim smtpd 4.71
| smtp-commands: EHLO linux Hello isc-l0065.local [192.168.1.137], SIZE 52428800,
    PIPELINING, HELP
|_HELP Commands supported: AUTH HELO EHLO MAIL RCPT DATA NOOP QUIT RSET HELP
53/tcp   open  tcpwrapped
80/tcp   open  http            Apache httpd 2.2.14 ((Ubuntu))
|_html-title: Ed's Web Page Test Zone
111/tcp  open  rpcbind         2 (rpc #100000)
| rpcinfo:
| 100000  2        111/udp  rpcbind
| 100003  2,3,4   2049/udp  nfs
| 100005  1,2,3  43439/udp  mountd
| 100021  1,3,4  52866/udp  nlockmgr
| 100024  1      57570/udp  status
| 100000  2        111/tcp  rpcbind
| 100003  2,3,4   2049/tcp  nfs
| 100024  1      35177/tcp  status
| 100005  1,2,3  41859/tcp  mountd
|_100021  1,3,4  41980/tcp  nlockmgr
139/tcp  open  netbios-ssn     Samba smbd 3.X (workgroup: NET)
443/tcp  open  ssl/http        Apache httpd 2.2.14 ((Ubuntu))
|_html-title: Ed's Web Page Test Zone
445/tcp  open  netbios-ssn     Samba smbd 3.X (workgroup: NET)
902/tcp  open  ssl/vmware-auth VMware Authentication Daemon 1.10 (Uses VNC, SOAP)
1984/tcp open  bigbrother?
2049/tcp open  nfs             2-4 (rpc #100003)
5900/tcp open  vnc             VNC (protocol 3.7)
8009/tcp open  ajp13           Apache Jserv (Protocol v1.3)
8222/tcp open  http            VMware Server 2 http config
|_html-title: VMware Server 2
8333/tcp open  ssl/http        VMware Server 2 http config
|_html-title: VMware Server 2
MAC Address: 00:22:15:BA:93:1C (Asustek Computer)
Device type: general purpose
Running: Linux 2.6.X
OS details: Linux 2.6.19 - 2.6.31
Uptime guess: 11.438 days (since Sun Apr 03 15:32:20 2011)
Network Distance: 1 hop
TCP Sequence Prediction: Difficulty=203 (Good luck!)
IP ID Sequence Generation: All zeros
Service Info: Host: linux; OSs: Unix, Linux

Host script results:
| nbstat:
|   NetBIOS name: LINUX, NetBIOS user: <unknown>, NetBIOS MAC: <unknown>
|   Names
|     LINUX<00>            Flags: <unique><active>
|     LINUX<03>            Flags: <unique><active>
|     LINUX<20>            Flags: <unique><active>
|     \x01\x02__MSBROWSE__\x02<01>  Flags: <group><active>
|     EDTETZ.NET<1d>       Flags: <unique><active>
|     EDTETZ.NET<1e>       Flags: <group><active>
|_    EDTETZ.NET<00>       Flags: <group><active>
|_smbv2-enabled: Server doesn't support SMBv2 protocol
| smb-os-discovery:
|   OS: Unix (Samba 3.4.7)
|   Name: Unknown\Unknown
|_  System time: 2011-04-15 01:59:48 UTC-3

HOP RTT     ADDRESS
1   1.41 ms 192.168.1.5

Read data files from: C:\Program Files\Nmap
OS and Service detection performed. Please report any incorrect results at
    http://nmap.org/submit/ .
Nmap done: 1 IP address (1 host up) scanned in 147.66 seconds
           Raw packets sent: 1021 (45.684KB) | Rcvd: 1016 (41.416KB)
```

What does this information allow an attacker to do? Well, it gives an attacker a fairly complete list of services that are offered by this network device, and if he wants to find a way onto a network, he can examine this list of services offered for a service that is known to be weak and use that as a method or path to gain access to the system. For example, if an attacker has found a Windows computer telling him that TCP port 3389 is available, he can run Remote Desktop Connection (`mstsc.exe`) to connect to that computer and try a number of common passwords for the Administrator account, or he can run some tools or exploit some known weaknesses in the Windows OS.

Dealing with external attacks

Attacks that start from outside a network fall into a couple of categories: They tend to be either denial of services (DoS) or attempts to gain access and exploit a system. In many cases, these are both one and the same. When your devices are running correctly, they have space to log data and access attempts, and applications — especially the security applications — all have enough memory to operate. Many of the attacks in the DoS category flood the systems with so much data that these data logs overflow (so you cannot see what the attacker is attempting), and security applications or processes run out of memory and possibly shut down or malfunction. When your system has nonfunctioning security applications and a lack of logging, the attacker can take control of that system to further her access on your network. The following sections describe common external attacks you need to be aware of and prepared for.

SYN flooding

In Book II, Chapter 2, I describe what a three-way handshake looks like. What would happen if I did not complete the process? The situation would be like getting a phone call at home, but when you answer the call and say “Hello” (the second step in the process), the person on the other end does not respond or tells you to hang up. In this case, you stay on the phone for some period of time, during which you cannot take any more calls.

In the TCP world, your network devices are capable of handling a limited number of connections. It is a high number, but it is limited based on the device and its configuration. Figure 1-3 shows what happens in a *SYN flood,* an attack where the attacking device sends a series of SYN requests with the goal of overwhelming the system. What the attacking system does not do is respond to any of those returned SYN-ACK packets. Because you have a limited number of listening connections on your system, for a relatively short period of time, you cannot accept a new connection because all the lines are busy waiting for ACK packets from the person who opened all the connections.

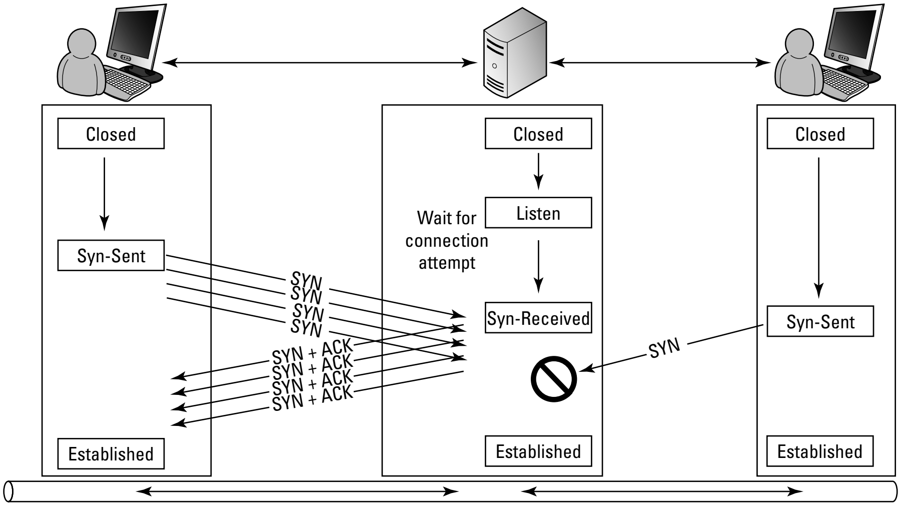

Figure 1-3: When all the Listening connections are consumed, legitimate users are prevented from connecting.

SYN flooding is a denial-of-service attack because legitimate users of the system cannot connect and do what they would typically be able to. This attack may interrupt services, or it may be an attempt to fill log files so that the actual attack does not leave any trances. After one of your systems has been targeted by a SYN flood, you may be able to connect to the flooded system and clear these half-opened connections rather than waiting for the system to time them out and clear them on its own schedule. Although SYN flooding is an old attack, it is still an effective attack on many systems.

Cisco devices allow you to do a few things to reduce the effectiveness of these attacks:

- Increasing the TCP backlog
- Reducing the SYN-RECEIVED timer
- Implementing a SYN cache
- Implementing SYN cookies

Smurf attacks

Smurf attacks are popular DoS attacks, likely named because of its use of a large number of small ICMP packets. The goal of this attack is to create a crushing amount of traffic. This attack came about as a function of ICMP and the network broadcast address. If an attacker has a large network segment that he is aware of, he can send a ping or an ICMP Echo Request to that broadcast address. Each host on that network should take that because the broadcast address was used, though the Echo Request is actually destined for itself. What it should then do is generate an Echo Reply back to the attacker's computer. This Echo Reply could cause a full Class B address block to generate up to 65,000 replies for one request packet.

This Echo Reply floods the attacker's computer with replies that could then cripple his computer with only a few Echo Requests being sent out. Now, what if the attacker modifies that Echo Request when it goes out, and instead of specifying his address as the source of the packet, he specifies another address? This trick allows an attacker to cause that crippling number of replies to be sent to an innocent third party, who would be the actual target of the Smurf attack.

Distributed denial of service (DDoS)

In a *distributed denial-of-service (DDoS)* attack, the attacker takes the same basic process as with any denial of service attack: He generates a sufficient amount of traffic to overwhelm the targeted device. The method may be a SYN flood attack or other attack, but what makes a DDoS attack unique is that the attack comes from more than just one device.

A virus may turn your computer into a robot waiting for commands to be run from a controlling server or device. This allows your computer to play a small role in a DDoS attack on some poor unlucky server or network. This way, if your defense to the DoS attack was to identify and block offenders, you can no longer defend your system because the attack comes from many places at the same time.

So if an attacker has an army of robot computers to generate SYN flood or Smurf attacks, she can wield a crippling level of power. (Insert maniacal laughter.)

Password attacks

Sometimes people attempt password attacks on a running system; but with passwords that lock accounts out after a few failed logon attempts, this attack is not very productive. More typically, password attacks capture RAW logon traffic from the network or break into a backup of a domain controller or workstation on the network. If an attacker reboots a workstation on a network from a CD or USB key, he can quickly grab a copy of the Windows SAM and security files from the Windows directory.

With these files in hand, the attacker can spend as much time as he wants trying to guess the passwords that are found in these files. In the case of the workstation, these security files give an attacker the local passwords on a computer, such as the local Administrator account; which he can then use to get a hold of network passwords which will provide more access to the network.

From the SAM file, an attacker can attempt to use two methods to crack these passwords:

- **Brute force attack:** With this attack or password-guessing technique, the cracking software goes through every password possibility from *a* through to *zzzzzzzzzz,* including all possible numbers or punctuation characters. This process of finding a password can take a very long time, longer for every extra character put into the password.
- **Dictionary attack:** This password-guessing technique can be done much faster, and it makes use of dictionary or word list files. These files are readily available on the Internet and include dictionaries such as the standard Oxford dictionary, every word found in the works of William Shakespeare, and even obscure or made-up language dictionaries such as Klingon. With these word lists in hand, the attacker can quickly compare these words to the hash values for the Windows passwords found in the SAM file. To speed things up even more, he can have his computer go through the dictionary files and create a password hash for every word, and then he just needs to compare the pregenerated hash values with those found in the SAM file. This process can give the attacker even more speed in finding these passwords. Despite warnings not to, many people still use standard dictionary words for their passwords.

REMEMBER

To protect your network from these possibilities, you need to provide some level of protection to the local OS installations, especially for workstations that are in the public or in areas of high public access. Workstations that must be in those areas should have physical security preventing them from being rebooted from custom media. Passwords used on these systems should be different from main domain client systems. Finally, use a strong password policy, which includes regular changes to the passwords and a requirement that the passwords should not be dictionary words.

Implementing Firewalls

A strong perimeter security helps to protect your network from external attacks. The main element on the perimeter security front is a firewall.

Types of firewalls

You can deploy several types of firewalls and other security options. The different types of firewalls include the following:

- **Packet filtering:** These firewalls use ACLs to inspect the data that they forward down to the IP layer. This inspection allows them to classify data based on the TCP or UDP ports, as well as the source and destination IP addresses. This filtering allows you to make forwarding decisions. Some organizations use packet filtering to allow only traffic that meets approved criteria to pass out of the firewall.
- **Stateful inspection:** Also known as Stateful Packet Inspection (SPI) firewalls, these firewalls not only allow packet filtering, but it pays attention to the flow of the packets. Rather than evaluating each packet as a separate entity, it looks at the flow of the traffic and identifies packets that are replies to others. SPI can evaluate packets that are suspicious and part of an attack profile.
- **Application layer firewall:** This firewall can be a specific firewall, but it tends to fall in the category of proxy and reverse proxy servers. In this case, there can be a deep packet inspection into the data to validate that it is not only allowed, but also not part of an attack on the systems that make up your network. These firewalls tend to be specific for the application layer protocol that they are protecting. Common choices here are HTTP, FTP, and SMTP.

Ingress and egress filtering

Most firewalls act as gatekeepers for networks or network segments and exist in a position where a router would exist. In fact, if the feature set has been enabled, your Cisco router can easily be called a firewall if it does any filtering of the traffic on your network. As a gatekeeper for your network, this device carefully filters out undesirable traffic that attempts to enter your network.

Although most people think of firewalls as protecting the network from incoming traffic, they can also prevent traffic from leaving your network. You can restrict your internal users from getting off of your network and going anywhere they would like. That is part of the egress filtering, which can be just as important as the ingress filtering.

I know some very paranoid people that use `deny` Access Control Lists (ACLs) as their basic network access rule on all firewalls in both directions, so all network traffic incoming or outgoing needs approval. I cover ACLs in Chapter 3 of this minibook. This method does take some commitment, but it ends up being very secure, if you manage to still keep it functional.

Defending data with the DMZ

A *demilitarized zone (DMZ)* is an area two opposing military forces have declared as a buffer zone between each other. Both sides agree that they will stay out of that area. For computer networks, the DMZ is an area where you have placed servers that the public at large — or at least people outside your network — need access to. These are placed outside your network and may have the ability to talk to your internal server. Although the servers are placed outside your network, they are not totally unprotected; they are still behind a firewall in a configuration similar to one of the options shown in Figure 1-4. The DMZ segment may be installed next to your current firewall or may be an actual zone between your network and the public network. Either is a valid DMZ option, each offering a benefit tradeoff between ease of configuration and security.

Finding holes in your firewall

You can run Nmap internally on your network to identify devices, but you can also run it from outside your firewall to identify holes in the firewall. Steve Gibson of Gibson Research Group (`www.grc.com`) offers a test for open ports from outside your network through his Shields Up! tool. The following figure shows what the results look like. This tool shows you what type information you are allowing to pass from your workstation to outside your firewall. Kind of spooky.

In the figure, the results show what I expected: No proof of a firewall existing on most ports, except for the ports on which I actually host services.

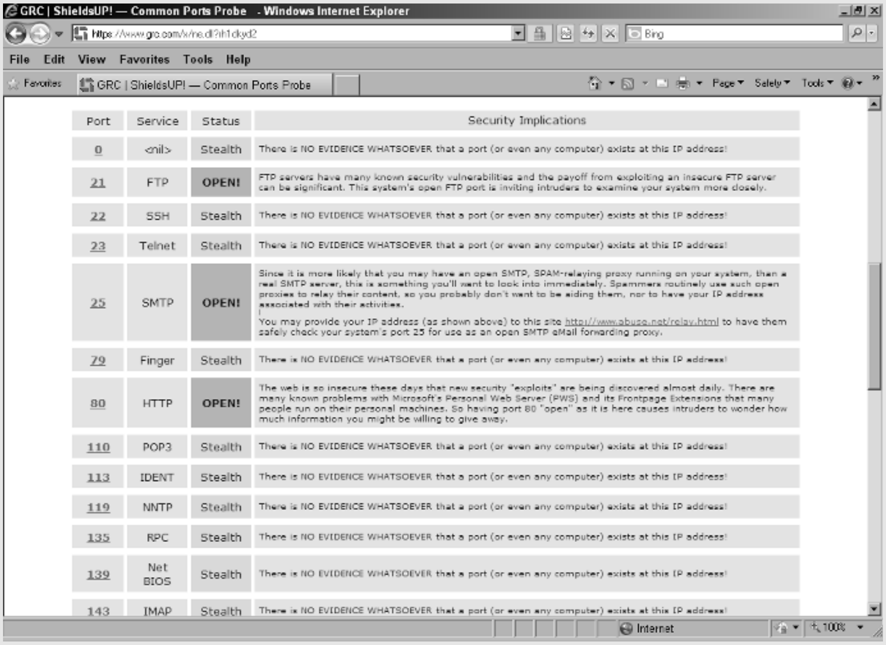
GRC | ShieldsUP! — Common Ports Probe - Windows Internet Explorer
https://www.grc.com/x/ne.dll?ih1ckyd2
Bing
File
Edit
View
Favorites
Tools
Help
Favorites
GRC | ShieldsUP! — Common Ports Probe
Page
Safety
Tools
Internet
100%

Port	Service	Status	Security Implications
0	<nil>	Stealth	There is NO EVIDENCE WHATSOEVER that a port (or even any computer) exists at this IP address!
21	FTP	OPEN!	FTP servers have many known security vulnerabilities and the payoff from exploiting an insecure FTP server can be significant. This system's open FTP port is inviting intruders to examine your system more closely.
22	SSH	Stealth	There is NO EVIDENCE WHATSOEVER that a port (or even any computer) exists at this IP address!
23	Telnet	Stealth	There is NO EVIDENCE WHATSOEVER that a port (or even any computer) exists at this IP address!
25	SMTP	OPEN!	Since it is more likely that you may have an open SMTP, SPAM-relaying proxy running on your system, than a real SMTP server, this is something you'll want to look into immediately. Spammers routinely use such open proxies to relay their content, so you probably don't want to be aiding them, nor to have your IP address associated with their activities. You may provide your IP address (as shown above) to this site http://www.abuse.net/relay.html to have them safely check your system's port 25 for use as an open SMTP eMail forwarding proxy.
79	Finger	Stealth	There is NO EVIDENCE WHATSOEVER that a port (or even any computer) exists at this IP address!
80	HTTP	OPEN!	The web is so insecure these days that new security "exploits" are being discovered almost daily. There are many known problems with Microsoft's Personal Web Server (PWS) and its Frontpage Extensions that many people run on their personal machines. So having port 80 "open" as it is here causes intruders to wonder how much information you might be willing to give away.
110	POP3	Stealth	There is NO EVIDENCE WHATSOEVER that a port (or even any computer) exists at this IP address!
113	IDENT	Stealth	There is NO EVIDENCE WHATSOEVER that a port (or even any computer) exists at this IP address!
119	NNTP	Stealth	There is NO EVIDENCE WHATSOEVER that a port (or even any computer) exists at this IP address!
135	RPC	Stealth	There is NO EVIDENCE WHATSOEVER that a port (or even any computer) exists at this IP address!
139	Net BIOS	Stealth	There is NO EVIDENCE WHATSOEVER that a port (or even any computer) exists at this IP address!
143	IMAP	Stealth	There is NO EVIDENCE WHATSOEVER that a port (or even any computer) exists at this IP address!

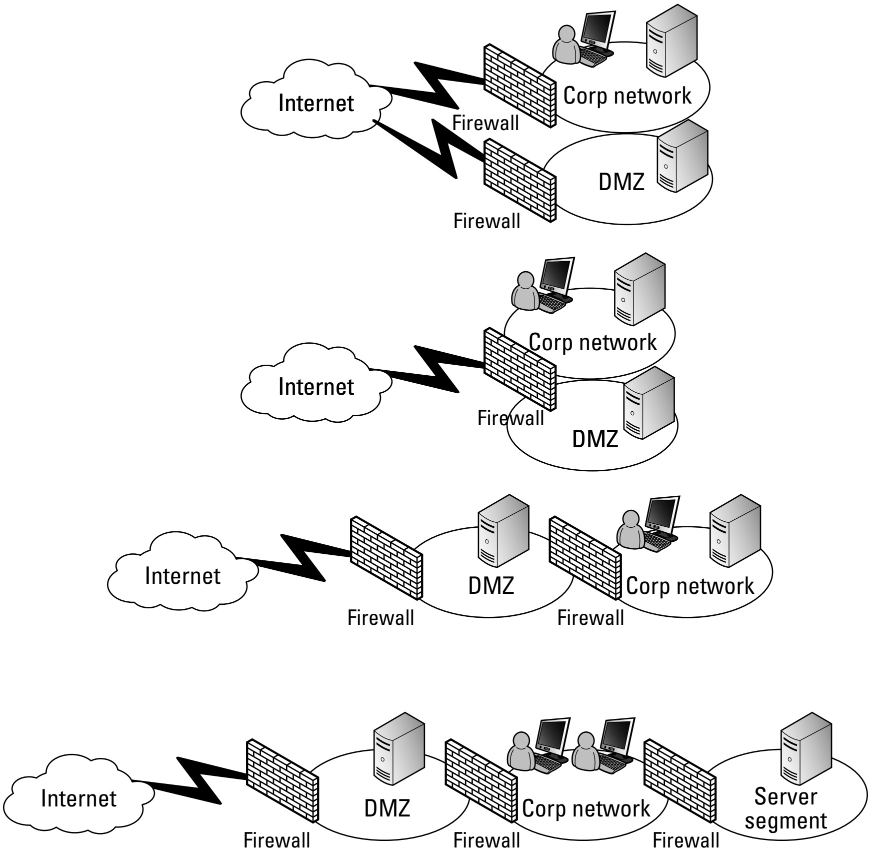

Figure 1-4: DMZ configuration options.

In the same manner as creating a DMZ, I have some clients who have isolated their servers on a separate network segment, with a firewall defending that segment. In this scenario, they have a DMZ protecting their public facing servers, their users behind the DMZ with protection from another firewall, and then their servers protected from their users behind yet another firewall. This scenario means that if an attack or a virus gains access to the user network, it will not get immediate access to the servers as well. Although this setup may seem a bit paranoid, it *is* very secure.

Defending Your Network against Attacks

As I mention earlier in this chapter, perimeter defense is one level of defending your network from attacks, and it works wonderfully to protect from external attacks. Perimeter defense is just part of the protection suite. I am a strong believer in defense in depth. *Defense in depth* means that several layers of security protect your network and its data, like Russian nesting dolls.

Perimeter defense

The job of perimeter defense basically falls to your firewall devices. Short of unplugging your network from the rest of the world, this perimeter defense seems to be a necessary evil. Now in some cases, people have gone with the unplugged solutions. (For example, the following scenario is not uncommon in the military: Several levels of networks operate, each with a specific protection level, the highest of which says you are not connected to any outside networks and no extra devices are allowed to be connected to this secured network.) This unplugged solution, however, does not work for most businesses.

So, as a necessary evil, you need to deal with a perimeter that you will attempt to defend. Limiting the number of connections from other networks, such as the Internet, to the network helps a lot because you have fewer connections for which you need to manage protection. Ideally, from the security perspective, you have only one connection to the Internet; whereas sometimes within a company, operation requirements mean that some groups within the organization require additional connections.

Active tools

In addition to managing your perimeter, you should implement intrusion detection systems (IDS) and intrusion prevention systems (IPS), which both offer a similar suite of options. In fact, you can think of IPS as an extension of IDS because an IPS system actively disconnects devices or connections that are deemed as being used for an intrusion.

IDS devices can be network-based devices, running as appliances or separate servers running software, which is performing the IDS role, but they can also be installed on client or network computers. The later is often referred to as host-based intrusion detection system (HIDS). These devices can reside inside your network, behind your firewall, detecting abnormalities there, and/or they can be placed outside your firewall. When they are outside your firewall, they are typically targeted for the same attacks that run against the firewall, thereby alerting you to attacks being run against your firewall.

Cisco offers several options for IDS and IPS systems and offers these as standalone systems or add-ons for your existing security products. The following are two such options:

- Cisco ASA Advanced Inspection and Prevention Security Services Module
- Cisco Catalyst 6500 Series Intrusion Detection System (IDSM-2) Module

IDS and IPS have several methods for working with detection. Similar to viruses on your network, intrusions and attacks have features that are recorded as a signature or behavior. So when the IPS system sees this type of data or behavior, the IPS system can swing into action. Suspicious behavior can also trigger these systems. This behavior can include a remote system attempting to ping every address on your subnet in sequential order, and other activity that is considered to be abnormal. When the IPS system sees this activity, the IPS can be configured to blacklist or block the source device, either indefinitely or for a period of time.

The other way these systems can identify suspicious traffic on your network is to have them run in a Learning mode for a period of time. Over the course of weeks, they can classify regular traffic patterns on your network and then limit traffic to those established patterns. If you introduce new software to your network, you may need to manually add appropriate rules or run a learning period and then put the system back into Prevention mode. This necessity is even true of the host-based systems because they update their rules from the management or policy server that is running on the network.

These systems help prevent the spread of *Day Zero attacks,* which are new viruses or network attacks that are different from all the previous network intrusions. Because these Day Zero attacks are new, you do not have a specific signature for the attack; but the attack still needs to perform the same suspicious behaviors, which can be detected and blocked. So even if I have never seen the attack profile in the past because it is new, I can still block it because it will do things that I have chosen not to allow.

Defense in depth

The best defense is a multilayered defense *(defense in depth).* If you review information about historic battles, in many cases, one side made a break through the front lines (the perimeter) and was able to move though a huge area because nothing was there to stop them. Other more defended areas had reserve troops behind the front lines that were able to quash the breakthrough.

So history shows time and time again that defense in depth is better than only perimeter defenses. The downside is that defense in depth tends to have higher costs. I had one client wonder whether he really needed to continue to pay to have his antivirus from one company scanning incoming e-mail on the outside of the firewall, then scan it again with a product on another the internal mail server, and then scan on the clients with a third product? Well right after pondering this, he was hit with a new Day Zero virus, which made it through the first two scans before being picked up at his client machines. The company is still running three layers of scanning.

So your cost for depth of defense may be a direct financial cost or may be in the time you spend to configure and manage it. However, because attacks can start from inside or outside your perimeter, you need to defend all areas of your network. And many of these attacks can start, willfully or accidentally, from inside your perimeter defenses.

Security Tools

You should use many different tools in your security arsenal. (My, more military metaphors!) Starting at your perimeter is your firewall. Other products you can consider in your defense toolkit are antivirus software, anti-malware software, spam filters, intrusion detection software, and vulnerability scanners. And, of course, you need to work with your network's users to make sure they understand and follow basic safe practices.

Personal firewalls

Just like your network firewalls, you can run internal firewalls on your network, even firewalls on your network devices. Microsoft has been nice enough to provide OS-based firewalls for all operating systems since Windows XP. So the cost of using this firewall is very reasonable, and it gives you a layer of protection right at each of your servers and client computers.

TIP

Although you get a free firewall with your operating system, some people choose to purchase a third-party firewall (such as Symantec Endpoint Protection) to give them enterprise network configuration and monitoring tools that will be used with the firewall. In many cases, the extra cost is worth the investment because you get easy-to-use management tools. The cost is outweighed by the cost of your time.

Antivirus software

Antivirus software keeps growing to defend against new threats. Rather than get into every type of anti software available, I focus on the three most common categories of threats that antivirus software protects against:

- **Viruses** are small, malicious programs that get installed without your knowledge, and they have specific goals, such as remote management of your computer or forwarding sensitive data from your device to another system over the network.

- **Worms** are small programs that tend to replicate over the network without much help from people to move them. Their goal is twofold: to spread and to accomplish whatever nasty business the worm writer intends.
- **Trojans** are small programs that need the computer operator's help to infect the computer or device. (The name is tied to the Trojan Horse of Greek legend, which would not have succeeded for the Greeks inside the horse, if not for people of Troy bringing the horse into the city.) Trojans are applications that are not what they claim to be. A common one is that pesky Windows Security Center Trojan that tells you that your computer is heavily infected, and you should install the software to fix it. Of course, installing the software infects your computer with the Trojan. Trojans also get packaged into other applications, so that when users install that application, they release the Trojans on their poor computers.

The first worm on the Internet was the Morris Worm back in 1988. It was claimed to have been a mistake and was supposed to help gauge the size of the Internet. It *infected* a system and started a task on that system to locate and infect other systems. Unfortunately, it would infect the same systems again and again, which eventually used up all CPU cycles on the infected servers, bringing them to a halt. The spread of this worm made use of known exploits on the servers and the fact that people tend to use common and weak passwords. These same principles are in use by virus writers today.

Because antivirus applications prevent the spread of viruses, many viruses these days make their first task disabling your antivirus software so they have free access to the system.

Anti-malware

Malware is software, hardware, or firmware that is malformed or malwritten such that it causes problems on computers or network devices. Malware are not viruses per se, but they are definitely undesirable and unwanted on any computing device. There is a bit of an overlap between anti-malware applications and antivirus applications, with antivirus applications typically ignoring code in the unwanted category and leaving that to the anti-malware applications.

Malware can slow a computer, give you unwanted pop-up ads, report Internet usage in the form of tracking cookies and other techniques, and do things that the world at large thinks is dodgy. Some applications make agreements to bundle malware with their products, and removing the malware could prevent the application from working. Some may say that is the price you pay for “free” software. (I wholeheartedly disagree!)

A variety of applications can help defend against malware and the sites that may try to push it. A few that I use regularly are Spyware Blaster (shown in Figure 1-5) and Spybot – Search & Destroy. I like these two tools because they both offer the ability to immunize your web browser against a huge number of infections. Spybot – Search & Destroy is also capable of scanning your system for files and registry entries that you want to have removed from your computer. I regularly install these tools on systems that I am called to clean or on systems that are going out to certain groups of users.

TECHNICAL STUFF

Spyware Blaster is free if you do not mind doing your updates manually. To update your system, click Update in the Updates window and then go to the Protection Status window and click the Enable All Protection link.

Malwarebytes, shown in Figure 1-6, is another tool that makes quick work of infections on your computer. The purchased version has auto updating and active protection against new infections Malwarebytes can run a scan of your entire computer and remove infections. I have found that multiple scans with all the tools, starting in Safe Mode booting and moving to Normal booting, typically gets rid of the most invasive programs. Do not forget your antivirus scans as well.

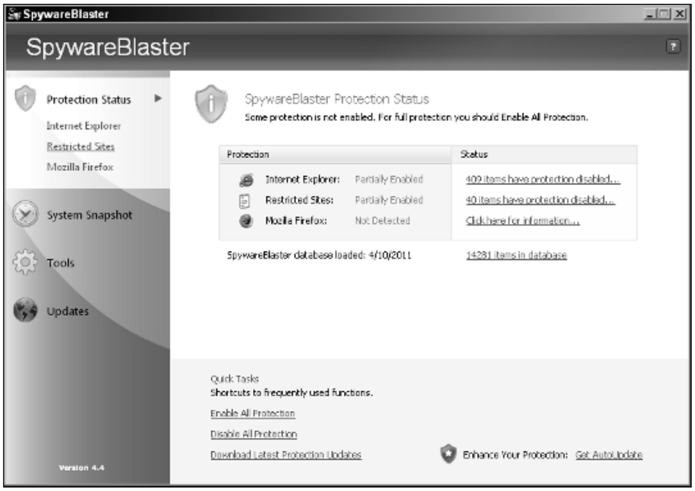

SpywareBlaster Protection Status

Some protection is not enabled. For full protection you should Enable All Protection.

Protection		Status
Internet Explorer:	Partially Enabled	409 items have protection disabled...
Restricted Sites:	Partially Enabled	40 items have protection disabled...
Mozilla Firefox:	Not Detected	Click here for information...

SpywareBlaster database loaded: 4/10/2011 14281 items in database

Quick Tasks

Shortcuts to frequently used functions.

Enable All Protection

Disable All Protection

Download Latest Protection Updates

Enhance Your Protection: Get AutoUpdate

Figure 1-5: Spyware Blaster has an easy configuration; just click Enable All Protection.

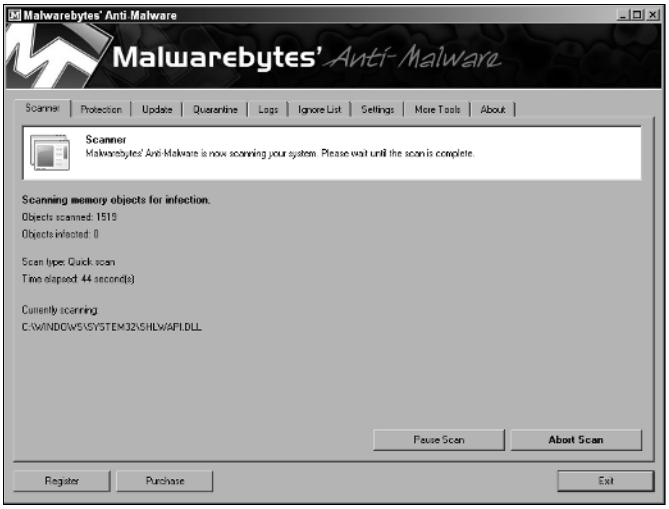

Scanner
Malwarebytes' Anti-Malware is now scanning your system. Please wait until the scan is complete.

Scanning memory objects for infection.
Objects scanned: 1519
Objects infected: 0

Scan type: Quick scan
Time elapsed: 44 second(s)

Currently scanning:
C:\WINDOWS\SYSTEM32\SHLWAPI.DLL

Figure 1-6: Malwarebytes runs full or quick scans of your computer for files or settings you may want to remove.

All these tools have been created with a lot of common sense, and it takes only a few mouse clicks to enable protection. Also with a few mouse clicks and several minutes of waiting, you can have the tools scan your system and then delete or fix any problems.

Spam filters

E-mail *spam* is all that gunk in your inbox that you did not ask for, which can include product promotions, viruses and exploits, and scams trying to get you to pay for a free trip to Italy. Spam e-mail is unsolicited, unlike e-mail from mailing lists that honor a request to unsubscribe.

E-mail spam (not the canned meat called SPAM) is a global problem. Most experts agree that well over 90 percent of all of the e-mail sent these days is spam. A recent Microsoft report claims that 97 percent of all e-mail is unwanted.

Without any spam filters, spammers would dump a tremendous amount of e-mail in users' mailboxes. The cost of spam is that to prevent it from filling users' inboxes, you need a frontend server to receive e-mail and filter out the spam. If you want people to be able to retrieve *false positive messages* (legitimate e-mail that was accidentally caught by the spam filter), you also need to store that spam for a period of time. So you have now paid for a server, hard drive capacity, power, cooling, and time — all to handle something that you do not want.

You can place spam filters on a filtering server outside the network or on a server inside the network. Personally, I like to keep this data outside my firewall. I can even have filtering on my local computer where I read my mail, but this is the last place where I typically want this job if I can avoid it.

Intrusion detection

Both intrusion detection systems (IDS) and intrusion prevention systems (IPS) have a place in my protection toolkit. These tools are like burglar alarms for a network or network devices. They can be installed outside the network, on the interior of the network, or even on specific hosts or computers on the network. This may well be a distant early warning (DEW) system for a network. When things start going down, these systems may prevent an attack from becoming a problem or at least let you know that it is happening before it has a chance to really take hold.

Vulnerability scanners

You can find several tools to scan your network for security holes. The premiere scanner is Nessus, which can perform a wide variety of scans and knows of many security holes in major products. It can generate detailed lists of items that you should check. Nessus is now distributed by `www.tenable.com`, and you can download and try it before purchasing it. Nessus used to be a free product but is now capable of doing approved audits for a variety of organizations, whose authorizations like to carry a licensing fee. In addition to Nessus, you can find a variety of other scanning tools listed at `http://sectools.org`.

User common sense

Making sure users follow basic, common-sense guidelines can really be one of the toughest security tools to properly implement. Even smart people can get duped into clicking that link, entering a credit card number, or believing the Nigerian royal family will give them money. (It makes me think of a line used in the BBC television series *Hustle:* “Everybody wants something for nothing, and we give them nothing for something.”)

Phishing for information

Remember: Legitimate companies *never* ask for personal data via unencrypted e-mail. If you receive such an e-mail, it is likely a *phishing scam* — a message from a scammer that looks like it comes from a legitimate company, such as a bank. These phishing e-mails, and the scam websites they link to, can look very convincing and may even include logos from the companies they purport to be. Do not click a link in an e-mail, or worse, enter any information on that linked website, because that is a favorite trick of the weasely people to get gullible folks to divulge personal information such as usernames, passwords, or credit card numbers.

If you need to update your personal information with a company, always be sure to go to the company's (secure) website by using your web browser.

Make sure users understand that if something does not seem right, or perhaps is too good to be true, very likely they should avoid it. Ask users to always be cautious and get the possible scammer's phone number and call him back, Google the name (or the Subject line of the e-mail), or just take a minute and think about it.

Do an Internet search for *Internet scams* to get a list of things you should avoid and then share the list with your network's users. If your search phrase includes a few details about a specific scam, you will likely find the one that you are concerned about. (I do have to hand it to some of these e-mails or websites trying to get your money. Some of them look pretty legit.)