Chapter 18
Hardening Networks: Firewalls

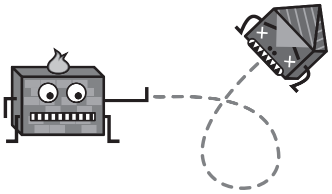

Firewalls are devices that can restrict information traveling in and out of a network.

Technology Overview

Don't you hate it when you are at a party, meeting someone, getting to know them better, and all they want to do is talk about firewalls? How their network gets scanned by hackers 200 times a day, how they've survived worms and viruses, and how good they are at security? Aren't they special?

Thankfully, party talk is not always that bad, but the term *firewall* has nonetheless broken the geek-speak barrier and obtained social buzzword status. Firewalls have also entered the world of the business mainstream. Businesses and home computer users have become comfortable with the concept of a firewall and its perceived role in the network. Perhaps too comfortable.

What people think: Firewalls solve all of our security problems. We are ahead of the game (sophisticated) because we have one or more firewalls.

What we think: Firewalls have become a defense that most perpetrators can circumvent with great ease.

One technology that has greatly contributed to common knowledge of firewalls is broadband (high-speed) Internet access. Broadband is now available to people in their homes throughout many major cities around the globe. Before broadband, it could be argued that home computers were rarely on the Internet because modem connections were typically short. Home modem users didn't feel like they were exposed to hackers (although they actually *were* primary targets). However, a broadband connection is always on, and therefore creates a large stationary target for hackers. Home users now see themselves as potential targets (which they still are). As a result, firewalls are now marketed to both home and corporate consumers as necessary security solutions.

What Is a Firewall?

Not all firewalls are created equal. Software firewalls and low-end "firewall appliances" (such as those built into cable/*Digital Subscriber Line* (DSL) modems) are radically different from the type of high-end firewalls sold by major commercial vendors. The difference is so significant that they should actually be called by different names. Nevertheless, since the public is comfortable with the "f" word, it has been applied to nearly every type of security networking product.

Most firewalls have one thing in common besides their name. They all perform *packet filtering*. A packet filter is simply specialized software that filters the information traveling between an outside network (such as the Internet) and a private internal network or computer system. To oversimplify the process, if incoming data does not meet certain criteria then it is not permitted to pass. More specifically, packet filters have the capability to accept or reject individual *Internet Protocol* (IP) packets based on information contained in the header of the packet. If that statement is confusing, skim over the few paragraphs on IP in Chapter 17, "Network Lingo."

Early packet filters treated each packet as a separate entity. They had no ability to deal with groups of related packets. For example, when a Web browser makes a request, a Web server sends many packets in response. Packet filters saw no relationship between the response and the request. This was a problem, because administrators wanted the ability to say, "If the outgoing request passes our filter, then all the related responses should be automatically allowed through the filter."

Stateful inspection is a concept that was added to packet filters to solve this problem. It compares certain parts of the packet to a database of trusted information. Information traveling from inside the firewall to the outside is monitored for specific defining characteristics, and then incoming information is compared to these characteristics. If the comparison yields a reasonable match, the information is allowed through. Otherwise, the packet is discarded. This procedure allows for a dynamic, on the fly, approach to firewalling.

Packet filtering software usually runs on routers, because firewalls often sit between two or more networks. Another related technology often integrated with routers and packet filters is *network address translation* (NAT), which is described in the following section of this chapter. In fact, most packet filters are capable of performing routing and NAT. There are some exceptions, though, such as software-based packet filters and *bridging firewalls*. Zone Alarm and Black Ice Defender are two examples of PC-based packet filters.

For many years, packet filtering, routing and network address translation were the only things that firewalls could do. Stateful Inspection provided a breath of fresh air, but it still left much to be desired. One glaring issue was that a packet filter has a blind eye toward the content of a packet. A malicious web packet looks the same as a harmless one. This is because firewalls were not supposed to look at the actual application data within each packet. That's a job for a different type of solution, called a proxy server, described in the sidebar.

Firewalls

Illustration by SageSecure

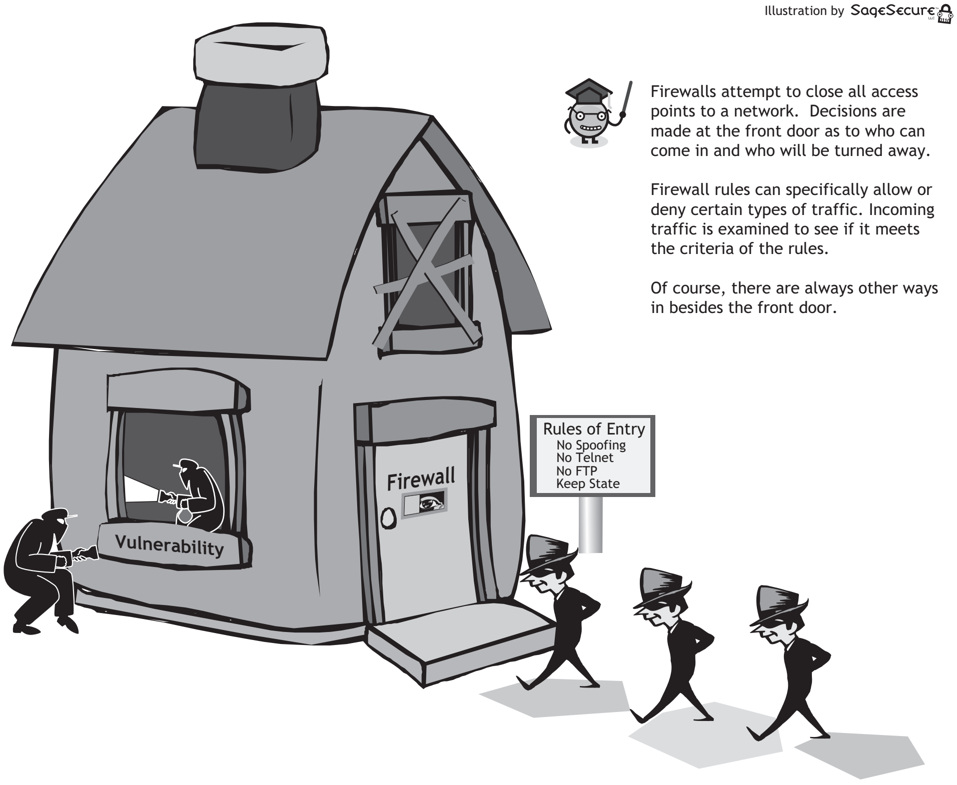

Firewalls attempt to close all access points to a network. Decisions are made at the front door as to who can come in and who will be turned away.

Firewall rules can specifically allow or deny certain types of traffic. Incoming traffic is examined to see if it meets the criteria of the rules.

Of course, there are always other ways in besides the front door.

Figure 18-1

Proxy Servers

A function that is often combined with a firewall is a proxy server. The proxy server acts as a middleman between a client and a server. It understands application data and can therefore make intelligent decisions about information.

When a computer protected by a proxy server makes a request to a remote server, the proxy server intercepts the request and inspects or manipulates the data. It then forwards the request on to the remote server (assuming it doesn't deny the request). When the remote server sends a response, the proxy server intercepts the information. Once again it has a chance to inspect, manipulate and potentially deny the data. The net effect is that the remote computer never comes into direct contact with anything on the protected network other than the proxy server.

Proxy servers can be used to control outgoing access as well as incoming access. For example, a web proxy server can require a password to access a general Internet web site. It can also prevent access to web sites that contain objectionable keywords, such as "sex," "x-rated," or "Rush Limbaugh."

Proxy servers can also make Internet access more efficient. Frequently requested data can be stored on the proxy server for rapid retrieval. Web proxies store frequently accessed web pages and images, speeding up retrieval time for popular web sites and reducing the amount of outgoing traffic. This means that the next time someone goes to that page, it doesn't have to load again from the web site. Instead it loads instantaneously from the proxy server.

Feature-Hungry Firewalls

Commercial firewalls are specialized computers running specialized operating systems with nice management interfaces. These devices are cleverly marketed and sold as complete one-stop solutions. In the process, the lines between firewalls, proxy servers, and other related security technologies have been blurred. In their race for one-upmanship, vendors advertise their firewalls as having capabilities such as virus scanning, web filtering, network performance enhancement, virtual private networking, and even intrusion detection. Different vendors provide different assortments of extra features. This makes it difficult to say exactly what a firewall does, since no two firewalls do exactly the same thing. The only sure thing is that a packet filtering router (usually with stateful inspection built-in) is at the core of every solution. Many of the other "features" are provided by integrated proxy servers and other software bundled with the firewall.

Nothing is particularly magical about an expensive commercial firewall. With a bit of knowledge, a very effective firewall can be built from scratch using old computer equipment.[1] All of the features of the most expensive firewalls would be available to this low-cost solution. However, commercial firewalls offer guarantees of reliability and service. A major vendor can make verifiable statements about the reliability of their products because millions of them are in use every day. It's not as easy to make the same statements about a homegrown solution, even if it is technically and functionally identical (or even superior).

Reliability is a significant concern when implementing a firewall solution. A firewall, whether hardware or software based, commercial or homemade, will ultimately fall into your gateway chain. This means that many, if not all, of the computers inside your network will be passing packets through the firewall. If the firewall should fail to function in any way the whole network can lose external connectivity. Things can get even more painful when firewalls used to protect internal network resources fail.

Some commercial firewalls are designed to perform automatic fail over when they are purchased in pairs. This is a feature that is more difficult to implement with homemade firewalls. This offers a big advantage to corporations who are willing to shell out big bucks for peace of mind.

Firewalls in Practice

The most direct use of a firewall is to protect a network from outside intruders. A typical company might have hundreds of computers that are networked together. Such a company will often have one or more connections to the Internet through some type of broadband connection. Without a firewall in place, all those computers are directly accessible to anyone on the Internet. Data packets can leave the inside to get to the outside and packets can come in freely. A malicious individual can probe those computers just as easily as those computers can access the Internet.

A firewall creates a doorway to the network, that can be used to prevent unwanted guests from entering the house. Firewalls are often placed between a broadband connection and an internal network. The firewall can then be used to enforce aspects of corporate security policies. For example, one of the security rules inside the company might be:

Out of the 20 servers and 450 workstations used at this company, only one of the servers can receive requests for web pages. This server will be permitted to act as a web server and people from the outside world can access web pages hosted on the server.

The firewall would enable this by blocking outside traffic to every machine, with the exception of the one web server. This firewall would only allow web requests to reach this machine, and no other types of traffic. Furthermore, only web responses would be sent back to the remote computer. In the event that server was compromised

[1] For most networks, firewalls do not require a large amount of computing power. A T1 Internet connection can be easily firewalled with just a Pentium II processor.

(through a web-based attack), it would not be possible to run other non-web services. All unauthorized communication would have to take place over the web channel.

Firewalls are also used to enforce network usage policies. They can control employee access to network resources and they can log activities in order to identify resource misuse. Companies commonly use firewalls to restrict access to the Web and other Internet resources. A firewall gives a company tremendous control over how people use the network.

What Firewalls Can and Can't Do

The key to understanding firewalls is knowing what they can actually accomplish. Thanks to confusing marketing techniques, many people think that a firewall is a complete security solution. Even the most advanced firewalls are only partial security solutions. To set things straight, we're going to look at some common security problems and the role that firewalls can (or can't) play in providing a solution.

Network vulnerabilities: Packet filters can help protect against attacks that exploit vulnerabilities in fundamental network protocols. For example, the path of a packet traveling through the Internet or any other network is usually chosen by the routers it meets along the way. However, the IP header of the packet can also specify the route that the packet should travel. This is called *source routing*. Hackers sometimes take advantage of source routing to make information travel through compromised systems. Normal applications almost never use this feature of the IP. Consequently, packet filters can prevent hackers from exploiting source routing by blocking packets that contain these spurious instructions.

Application vulnerabilities: Some programs have special features that allow for remote access. Others contain bugs that provide a backdoor or hidden access that provides some level of control of the program. Closely related to these are the next type of vulnerability.

Operating system vulnerabilities: Like applications, some operating systems have backdoors. Others provide remote access with insufficient security controls or have bugs that an experienced hacker can exploit.

Firewalls can be configured to block all unrequested external attempts at connecting to internal systems and servers. That's easy for a packet filter to accomplish. However, a packet filter can't tell the difference between a legitimate web page and an exploit page designed to crash the web browser and install a trojan. If a vulnerable application running on an internal system makes a connection to an external server, it's possible for that server to exploit the application.

Viruses: Probably the most well known threat is a computer virus. A virus is a small program that can embed itself in other programs and copy itself to other computers. Viruses can spread quickly from one system to the next via floppies, CDs, and the Internet. Most viruses come as attachments in email or through a file downloaded off the Web. Some spread by exploiting application vulnerabilities. Those that spread using network applications are called worms. Some viruses simply display harmless messages; others can erase all of your data.

Firewalls can usually stop some worms from spreading by preventing them from entering or leaving your network. Some commercial firewalls come bundled with virus scanners. These use integrated web, file transfer, and email proxies to intercept data that might contain viruses. But in general, packet filters can't tell the difference between healthy data and a virus. The only way a standard firewall can protect against viruses is by blocking off access to the Internet in general. Even this does nothing against viruses brought in via CDs and floppies.

Trojans: Trojans are closely related to viruses. The difference is that they provide some degree of remote access. Trojans often instigate connections to remote machines, opening the doorway from the inside and letting in the intruder.

Trojans disguise their communications to look like harmless traffic to packet filters. If stateful inspection is being used, the trojan will effectively create an open command channel. Packet filters alone are relatively useless against trojans. High-end firewalls with virus scanners and other features may provide somewhat better protection.

Spam: Typically harmless but always annoying, spam is the electronic equivalent of junk mail. Spam can be dangerous though. Quite often, it contains links to web sites. Be careful of clicking on these because you may accidentally install a Trojan that provides a backdoor to your computer.

Firewalls can do little about spam, other than blocking out email entirely. A better solution is to use a spam filtering solution. This is either a proxy that can be integrated with a firewall or an add-on for mail server software.

Denial of Service (DoS): You have probably heard this phrase used in news reports relating to major ISPs, such as AOL, and major web sites. What happens is that a hacker creates a situation where the target machine can no longer process information. There are many ways of instigating a DoS. A common technique is to flood the target system with requests. The target system becomes so overwhelmed by hacker requests that it can't process normal traffic. The attack effectively takes a system offline. Sophisticated DoS attacks leverage flaws in software and protocols to dramatically increase the impact of the attack. These attacks can crash servers and possibly do permanent damage.

Years ago, DoS attacks came from no more than a handful of systems. Firewall rules could be adjusted to block traffic from these networks until the attack subsided. Attacks that are more modern use thousands of "zombie" systems all over the world. These are machines that have been previously compromised for the purpose of launching distributed attacks. This type of attack is nearly impossible to counter. There's no way to block all the zombies—they're too numerous and they constantly change. Furthermore, the zombie systems aren't the actual hackers. Zombies are ordinary users who don't realize that their computer is participating in a massive attack.

One way to combat DoS attacks effectively is to work with ISPs and backbone providers to limit traffic in DoS conditions. They can use traffic shaping and proxy techniques to prioritize normal traffic, minimizing the effective impact of the attack.

Firewalls are not useful for stopping DoS attacks, but they can be used to prevent your network from unwillingly participating in one. A good firewall administrator will know how to create rules that will prevent abusive traffic from originating from within your network.

E-mail bombs: An e-mail bomb is a DoS attack usually focused on a specific person. Someone sends the same email hundreds or thousands of times until the recipient's email system cannot accept any more messages. A similar technique would be to fill up a person's voice mail system or answering machine. This technique could also be broadened to overload an entire company's mail server.

If the attack is coming from a single IP address, a temporary firewall rule can block email from the address until the problem is resolved. If the attack is distributed across thousands of IP addresses, packet filtering is not going to solve the problem. A proxy is needed that can filter out email messages that match a certain pattern or that are overly repetitive. Some spam prevention products can be configured to provide this protection. If the firewall includes spam filtering, it may be able to combat this attack. Otherwise, additional systems will be necessary.

Social Engineering: A hacker can always use social engineering techniques to bypass a firewall. He or she might convince somebody to install some software on an internal PC. This would get a virus or trojan past any firewall.

Dedicated Hackers: A highly skilled, dedicated hacker can find ways into almost any network. The compromises that a network administrator makes on a daily basis result in many opportunities for hackers to gain access. Good security philosophies and policies are the only effective defense against smart, dedicated hackers.

How Packet Filters Work

Packet filters work on a simple principle: they inspect packets and either accept (pass) or deny (block) them. This can be approached in one of two ways:

- Accept everything except for specific things that should be blocked, or
- Block everything except for certain things that should be accepted.

These are called *default policies* because they describe what happens if the firewall doesn't have explicit instructions matching a given packet. When a bouncer at a club is making sure patrons aren't wearing jeans or sneakers, he's using the former policy. When he's checking for invitations, he's using the latter. In general, the latter of the two is much more secure and usually more consistent with security philosophies.

The list of things to accept or block is called a *rule set*. Each individual instruction is a rule. Packet filter rules can make decisions based on the various information present in the TCP/IP header. This includes: source and destination IP address/domain name, source and destination port, protocol (TCP, UDP, ICMP, and so on), and specific protocol options (IP options, TCP state, ICMP message type). Quite a bit can be accomplished using those few levers. Let's look at these filtering choices more closely:

Chapter 18 Hardening Networks: Firewalls

IP Addresses: Each machine on the Internet is assigned a unique address called an IP address. IP addresses are 32-bit numbers, normally expressed as four octets in a dotted decimal number. A typical IP address looks like this: 216.239.57.101. If a certain IP address outside the company is reading too many files from a server, the firewall can block all traffic to or from that IP address. The firewall can also block traffic coming from certain inside machines going outbound. Some companies actively block certain web sites (such as Hotmail, Yahoo mail, MSN mail, and so on) that they feel are inappropriate or harmful to the workplace environment. Others block everything and then allow access to a handful of pre-selected sites. The rest allow complete access to the Web.

Domain names: Most servers on the Internet have human-readable names known as domain names. These exist because it is hard to remember the string of numbers that make up an IP address, and IP addresses sometimes need to change. For example, it is easier for most of us to remember www.google.com than it is to remember 216.239.57.101. Furthermore, some of the larger domain names have large clusters of redundant servers, each with a different IP address. The domain name alternates between these IP addresses on a request-by-request basis to spread out the traffic load. When this happens, it's very difficult to block based on IP address alone. Most packet filters will allow filtering on the domain name, which gets around this problem. The only caveat is that looking up a domain name takes a little time and may noticeably slow down the firewall's performance.

Protocol: The protocol is the pre-defined way that someone who wants to use a service talks with that service. The "someone" could be a person, but more commonly is a computer program like a web browser. Protocols are often text, and simply describe how the client and server will have their conversation. *Hypertext Transfer Protocol* (HTTP) is the protocol of the World Wide Web. Firewalls can filter based on low-level protocols such as *Transmission Control Protocol* (TCP), *User Datagram Protocol* (UDP), and *Internet Control Message Protocol* (ICMP). Higher-level protocols can be filtered using ports.

Ports: Any server machine makes its services available to the Internet using numbered ports, one for each service that is available on the server. For example, if a server machine is running a web HTTP server and an FTP server, the web server would typically be available on port 80, and the FTP server would be available on port 21. A company might block port 21 access on all machines but one inside the company. No hard rule on ports exists though—a web server can run on port 21 and an FTP server can run on port 80. This usually doesn't happen because it's very confusing, but hackers might put their own server on port 80 in order to make a firewall think that it's a harmless web server.

Security Considerations

Firewalls are only as good their configuration, and configuring them can be quite difficult and time consuming. In fact, many firewalls in many companies are completely ineffective at protecting the networks on which they reside. Direct attacks against a network become possible when firewalls are poorly configured.

Misconfiguration isn't the only problem with firewalls. Changing business needs and changing technology infrastructures have minimized the value of firewalls in preventing certain types of attacks. Most of today's hackers attack services that few companies are willing to give up or block with a firewall. On average, these services center on web access and email access. The hackers have responded to the relentless installation of steel doors by going in through the window. web-based exploits, email viruses, trojans, and simple social engineering are all tricks used by modern hackers that can effectively bypass firewalls.

Commercial firewalls are highly marketed to consumers and businesses. They make you think that all you need to do is buy one, or maybe two, drop it on your network, and with a little help from their technician your network will be safe. This could not be further from the truth. Although commercial firewalls themselves will suffice as firewalls, they need to be configured based on a company's specific business and network needs. If you do not take a long hard look at all your network services, you can never configure a firewall properly.

To make an analogy, imagine if you wanted to pay a police officer to keep certain cars from entering a tunnel that you own. The cars you want to keep out are all makes and models that have a chassis of a certain weight or size. In order to know what kind of cars you want to prohibit entry, you need to know all about the possible cars that may come down the road. Well, many car manufacturers produce many different kinds of cars all over the world. Any car you do not explicitly tell the officer to flag down from the highway will be ignored and inevitably pass through the tunnel. Don't you think it would take a good deal of time to become knowledgeable about every car and its chassis specifications?

It takes even more knowledge for a network engineer or a security specialist to tighten down a firewall rule set. Even with the knowledge of every harmful packet of information, tuning a firewall still requires trial and error, experience, and a long-term strategy.

Even firewalls that appear to be configured properly can leave open doors that intelligent hackers can exploit. For example, stateful inspection is often used to allow outbound access while blocking all unrequested inbound traffic. The assumption is that all outbound requests are made by users, and therefore are safe. But a trojan server can establish a connection to a remote client by simply sending web packets out through the firewall. The server response will be allowed back in because the outbound connection looked just like a Web request. In reality, a hacker now can use this connection to directly control internal machines from outside the network, bypassing the firewall. A heavy steel door is useless if the door is held open for the intruder.

Making the Connection

Connecting Networks: Firewalls are often devices used to connect networks, or pieces of networks together. Core networking protocols are used by firewalls to perform basic networking services as well as advanced packet-filtering techniques.

Detecting Intrusions: Intrusions can occur at the network entrance point a firewall guards. Intrusion detection systems are often integrated into firewalls so blocking and watching can occur on the same physical device.

Best Practices

The level of security your philosophy and policies require will determine how many threats can be stopped by your firewall. The highest level of security would be to simply block all traffic. Obviously, that defeats the purpose of having a network connection. But a common practical rule of thumb is to block everything, and then select traffic to allow. This is a good rule for businesses that have an experienced network administrator who understands what the needs are and knows exactly what traffic to allow through. For most of us, it is probably better to work with the defaults provided by the firewall developer unless there is a specific reason to change it.

Transparent Bridging: Firewalls can be configured as bridging devices. This means that they can be transparently placed between two network devices. Traffic from one device is filtered and forwarded to the other device. If the traffic is allowed through, it happens as if the firewall isn't there. If it's not allowed through, the connection simply doesn't work.

The advantage to a bridging firewall is that it's invisible to devices on the network. There's no way to attack the firewall directly because it doesn't have an IP address. The only downside to transparent bridging firewalls is that they are more difficult to remotely manage.

Demilitarized Zone (DMZ): Sometimes you may want to have network segments that are deliberately insecure, or less secure than the private part of your network. These less secure segments are called demilitarized zones.

In the military, a DMZ is a buffer zone that sits in front of a well-defended, completely secure area. Although the firewall concept of a DMZ pulls its name from this example, it is not a very accurate analogy. A more accurate analogy would be a slightly less well known park in New York City called Gramercy. Gramercy Park is an anomaly that sits in the middle of a city full of anomalies.

Gramercy is a city-owned park that is semi-public and gated in as a result. To access it you need a key, but a copy of the Gramercy Park key is only given to residents who live in apartments surrounding the park! To the residents of Gramercy, the park is much like a DMZ on a network. It is an area that only residents have access to, but it is not as secure as their own private apartments. Although the residents prefer no

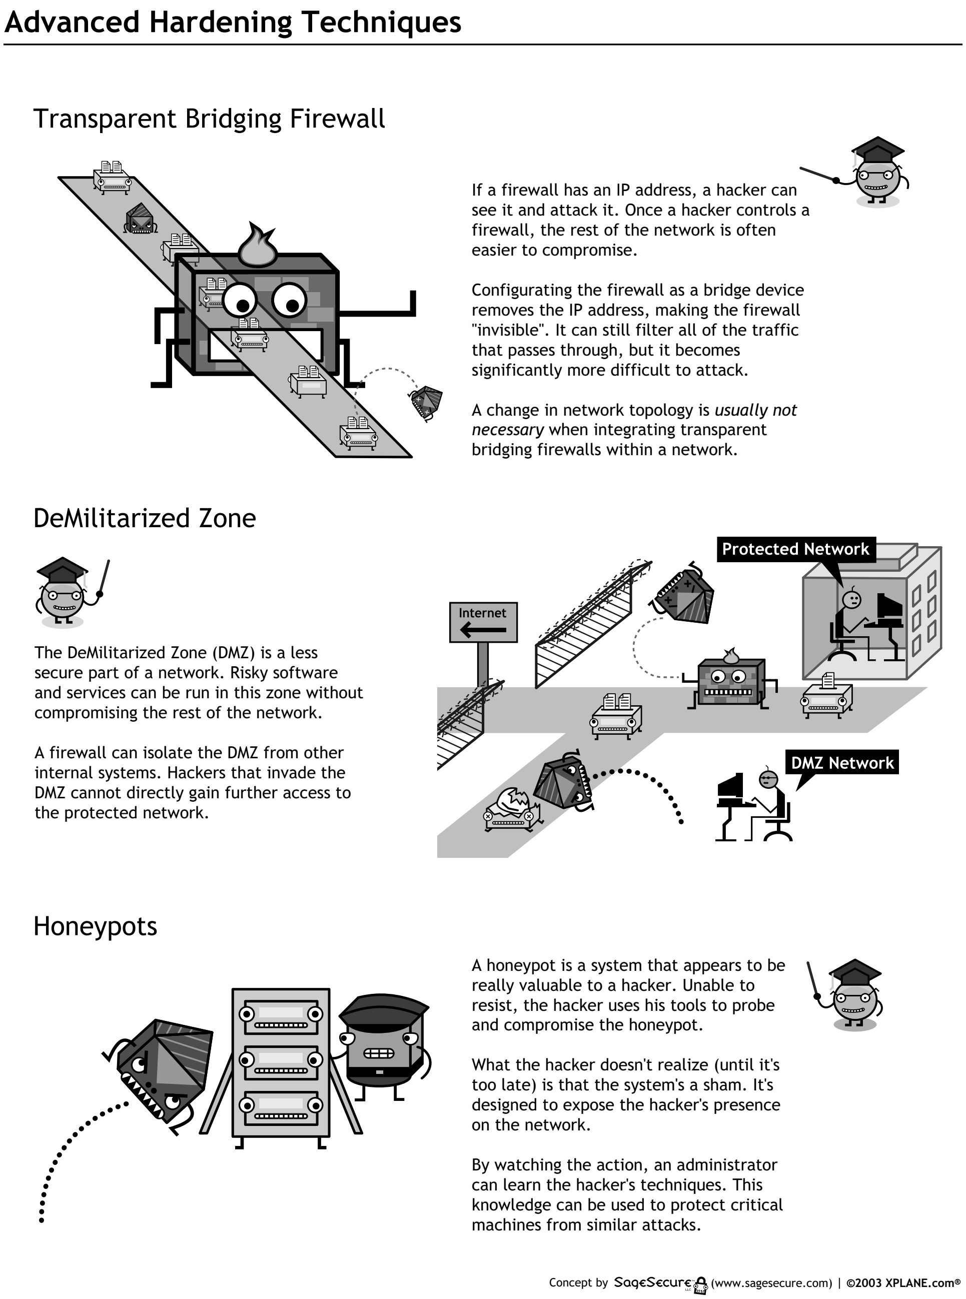

Advanced Hardening Techniques

Transparent Bridging Firewall

If a firewall has an IP address, a hacker can see it and attack it. Once a hacker controls a firewall, the rest of the network is often easier to compromise.

Configurating the firewall as a bridge device removes the IP address, making the firewall "invisible". It can still filter all of the traffic that passes through, but it becomes significantly more difficult to attack.

A change in network topology is *usually not necessary* when integrating transparent bridging firewalls within a network.

DeMilitarized Zone

The DeMilitarized Zone (DMZ) is a less secure part of a network. Risky software and services can be run in this zone without compromising the rest of the network.

A firewall can isolate the DMZ from other internal systems. Hackers that invade the DMZ cannot directly gain further access to the protected network.

Honeypots

A honeypot is a system that appears to be really valuable to a hacker. Unable to resist, the hacker uses his tools to probe and compromise the honeypot.

What the hacker doesn't realize (until it's too late) is that the system's a sham. It's designed to expose the hacker's presence on the network.

By watching the action, an administrator can learn the hacker's techniques. This knowledge can be used to protect critical machines from similar attacks.

Concept by SageSecure (www.sagesecure.com) | ©2003 XPLANE.com®

Figure 18-2

one without a key enter their park, if someone does it will not affect the security of anyone's private domicile.

Many reasons exist for having a DMZ on your network. Some companies use them as party zones. They place computers in the DMZ that they expect to be compromised. They use these computers to browse the Internet in a wildly carefree manner. This can be useful for testing, or research, or just plain fun. In other cases, DMZ's are used to provide services to the public Internet. These servers will be offered some protection in a DMZ, but not the same level of total protection afforded to the private network.

Some examples of services that a business might place in a DMZ are as follows:

- Web proxy
- News server
- FTP download and upload area
- Unrestricted web browsing clients
- Unrestricted peer-to-peer file-sharing clients

Setting up a DMZ will involve a slightly different process on every network. Sometimes a DMZ is created in the area between the Internet connection and the firewall. This can be achieved by simply placing a computer or server in this unprotected network pocket. In other cases, the DMZ will sit behind the firewall, but in a different subnet, with a less restricted set of firewall rules. Most of the software firewalls available will allow you to designate a directory on the gateway computer as a DMZ.

Honey Pots: A honey pot is a tasty looking trap, set for the unwary, starving hacker. The goal is to keep malicious folk away from the truly important stuff by hanging a big juicy steak in front of them. A honey pot machine is one that looks important, but is not used for anything. It might be running a fake database, or a CRM system populated with fictitious data. A honey pot has no legitimate users or traffic; any activity on the system is the result of an intruder. This can alert a network administrator to the presence of an intruder on the network. The authorities can be contacted while the intruder is busy gnawing on the tasty looking treats.

It is common for people to confuse honey pots and DMZs. A DMZ is not meant to attract hackers; it is an isolated network zone for services that need less security. Honey pots, on the other hand, are *explicitly* designed to attract hackers. A network administrator needs to have some type of plan for capturing and neutralizing any hackers drawn to the honey.

A study by Global Integrity published in September 2000 found in the following:

- Honey pots are an excellent method of detecting insider attacks.
- Honey pots sidetrack attackers' efforts, causing them to devote their attention to activities that can cause neither harm nor loss.

- Honey pots allow security administrators to study exactly what attackers are doing without exposing systems or networks to additional risk that results from compromised systems.
- A moderate proportion of attackers refrain from attacking systems within networks in which they know that measures have been taken to capture and monitor their actions.

Final Thoughts

We apologize for the extensive section on firewalls. Wes previously wrote a book on building firewalls. Jay tried to keep this section short and relevant, but Wes kept adding more and more stuff. After hours of arguing, Wes was sent to his room and Jay finished the chapter and submitted it to the publisher. But then Wes hacked into McGraw-Hill's network and kept adding stuff. See, Jay doesn't know I'm writing this, man will he be mad when he sees the printed book!