CHAPTER 11

Understanding Network Security Zones

Security is an important consideration with any network. Some areas of a network are more vulnerable to attacks than other areas. This increased risk requires increased security. Different areas of a network are categorized in zones with varying levels of security required in different zones.

The Internet is the riskiest zone. Internal networks, or intranets, are the safest. Between these two, you can create perimeter networks as a buffer zone. One of the primary methods of separating the zones is with firewalls. This chapter covers these different zones and provides some information on firewalls in general and Microsoft firewalls in particular.

- **Understanding risks on the Internet**
- **Exploring an intranet**
- **Understanding firewalls**
- **Identifying a perimeter network**
- **Understanding extranets**

Understanding Risks on the Internet

I'm betting you've used the Internet once or twice, but it's still worth mentioning here. It's the largest network in the world and continues to grow by leaps and bounds with no end in sight.

Several things have been mentioned about the Internet throughout this book, and it's worth consolidating them here in the context of network security zones:

The Internet Is the Riskiest Security Zone Attackers from anywhere in the world can attack computers on the Internet, and they do. In 2009 and 2010,

malware authors created 20 million new strains of malicious software (an average of 63,000 a day). Infected systems join massive botnets and participate in attacks on other computers.

All Internet Addresses Are Public Internet Protocol (IP) addresses used on the Internet are public IP addresses. In other words, they are accessible from any other computer with access to the Internet. In comparison, IP addresses on internal networks are private.

The Internet Is TCP/IP Based The TCP/IP protocol suite is the standard used on the Internet. Most internal networks use the same TCP/IP protocol suite for easy interaction on the Internet.

The World Wide Web (WWW) Travels Over the Internet The primary protocol used to transfer web pages is the Hypertext Transfer Protocol (HTTP). Note that the WWW isn't the Internet. Rather, you can think of the WWW like a semitruck delivering goods and the Internet as the highway that the truck travels on. Other protocols traveling over the Internet include the File Transfer Protocol (FTP) and Simple Mail Transfer Protocol (SMTP).

Botnets and Malware

Malicious software (malware) includes viruses, worms, Trojan horses, and other software designed with malicious intent. In the early days of computers, malware would often cause harm to a user's computer such as destroying data or destroying a user's hard drive. Some were relatively benign and simply popped up a message like "Legalize Marijuana" on a certain day.

However, malware has changed. Today, the primary purpose of most malware is to have a computer join a *botnet*.

Botnet is short for *robot network*, implying an automated network. Infected computers become a member of a botnet as a clone or zombie. The terms *clone* and *zombie* are interchangeable. Botnets are networks of these clones or zombies that can be secretly controlled at will by the attackers. Attackers manage computers on the Internet with command and control software that can issue orders to them. These zombies check in periodically and do the bidding of the attacker. It's not unusual for the attackers to have almost as much control of the user's computer as the user does.

Zombies may send spam on behalf of the attackers, steal identities, or steal financial data. Zombies also participate in massive *distributed denial of*

(Continues)

BOTNETS AND MALWARE *(Continued)*

service (DDoS) attacks on the Internet. A DDoS is a simultaneous attack on a single system or server by multiple attackers.

Any computer with access to the Internet (even computers within private networks) can become a zombie. Users are often unaware their computers are infected as zombies. Indeed, this is one of the strengths of botnets. They don't harm the user's computer but instead enlist it in their army. Today, it's not unusual for a botnet to have tens of thousands or even millions of zombies at their beck and call.

The best defense is antivirus software that is always on and regularly updated.

Exploring an Intranet

An *intranet* is nothing more than a LAN by a different name. A stricter definition is that an intranet is a private network that uses TCP/IP protocols to share resources within the network.

You can also think of an intranet as an internal network that uses the same protocols found on the Internet.

From a network security perspective, the intranet is the safest network security zone. It includes clients on the internal network and has substantially fewer risks than computers placed directly on the Internet. Administrators control these computers and can implement many layers of security on them.

However, don't think that computers within an intranet are risk free. They aren't. The only way to keep a computer free of risks is to leave it powered off. Of course, it isn't very useful without power.

Intranets have private IP addresses. Chapter 5 listed these usable private IP address ranges, but as a reminder, here they are:

10.0.0.1 through 10.255.255.254

172.16.0.1 through 172.31.255.254

192.168.1.1 through 192.168.255.254

You may remember that private IP addresses can only be used on internal networks, and they are never used on the Internet. However, most users within intranets need to access the Internet. Since private IP addresses are used in intranets and public IP addresses are used on the Internet, networks need some method of connecting the two. Enter NAT.

Understanding Network Address Translation

Network Address Translation (NAT) is a service that translates private IP addresses to public IP addresses and translates public back to private.

Consider Figure 11.1. It shows a private intranet with connectivity to the Internet via a router that is running NAT. The router does basic routing, and the NAT service translates the private and public IP addresses.

The router in the figure also has a firewall. Firewalls are explained in more depth later in this chapter, but it's common to implement firewalls between the Internet and intranet.

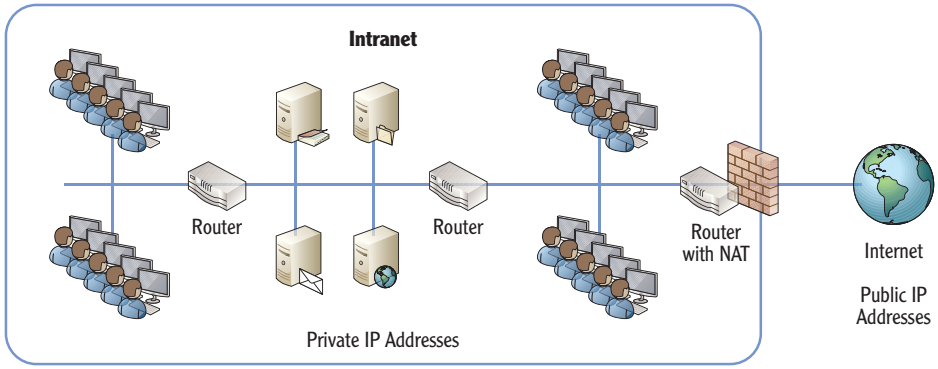

FIGURE 11.1 An intranet connected to the Internet

All the computers on the intranet have private IP addresses, and of course, the Internet has public IP addresses. The router with NAT has a private IP address assigned to the interface connected to the intranet and a public IP address assigned to the interface connected to the Internet.

Port Address Translation (PAT) is a popular way that NAT is implemented. PAT is sometimes called Network Address Port Translation, but more often than not, it's simply called NAT. The following explanation shows how the PAT version of NAT works.

Imagine that a user named Dawn on the intranet is trying to access Bing.com via the router. NAT will take the following actions:

1. It receives the request and logs the source IP address and port (Dawn's computer) and logs the destination IP address and port (Bing.com) in an internal table.
2. NAT then creates a new packet to forward the request to Bing.com. It keeps the destination IP and port but changes the source IP address to its own public IP address. It also changes the source port to an unused

port. At this point, the NAT table, with only one entry, looks something like this:

Source IP	Source port	Destination IP	Destination port	NAT source port
192.168.1.5	49155	Bing.com	80	49212

3. NAT sends the request to Bing.com. Bing.com returns the web page to the NAT server with the NAT source port (49212) included.
4. NAT looks at the source port and compares it to its internal NAT table. It sees that it's mapped to Dawn's computer with an IP address of 192.168.1.5 and then sends the page back to her computer.

The NAT table is stored in the system's memory. If it's a router running NAT, it's stored in the router's memory. If it's a proxy server, it's stored in the server memory.

You may be wondering why NAT created its own source port. That's a great point. It needs a way to identify the original requestor, and it does so with different source ports. Suppose that Jack was accessing Bing.com searching about feng shui at the same time Dawn was accessing Bing.com searching about firewalls. The NAT server would receive two answers from Bing.com. Without changing the source port for each request, there wouldn't be any way for NAT to determine who should receive which response from Bing.com.

The following table shows the NAT table with two entries. In this example, Dawn has an IP address of 192.168.1.5, and Jack's computer has an IP address of 192.168.1.22. NAT creates different source ports for each request in the internal NAT table. When Bing.com returns the data on firewalls requested by Dawn, it includes the source port created by NAT. NAT then uses this information to ensure that the request is forwarded back to Dawn's computer.

Source IP	Source port	Destination IP	Destination port	NAT source port
192.168.1.5	49155	Bing.com	80	49212
192.168.1.22	49158	Bing.com	80	49213

Chapter 4 covered the different ports. It included the well-known ports between 0 to 1023, the registered ports from 1024 to 49151, and the dynamic ports to 65,535.

You may be wondering how the source ports are generated. Most systems generate source ports from the dynamic port range of 49,152 to 65,535. Only ports that aren't currently being used are selected.

NAT provides several benefits:

Hides Internal Computers Since the computers don't have public IP addresses, they can't be directly accessed by Internet sources.

Reduces Costs If NAT wasn't used, you'd have to purchase public IP addresses for all internal computers. This is simply an unnecessary cost since it's so easy to install NAT.

Extended the Lifetime of IPv4 Since companies can use a single public IP address for hundreds or thousands of internal computers, the public IPv4 address range wasn't depleted earlier.

Although NAT can use a single public IP address, it's also possible to use multiple public IP addresses. Consider a large network with thousands of users. A single connection to the Internet may not be enough to adequately serve all of these clients. Instead, additional connections can be added with additional public IP addresses.

Static NAT uses a single public IP address, and all connections are mapped to this single IP address. Dynamic NAT uses a two or more public IP addresses. Any user's request from a private IP address can be dynamically mapped to any one of the public IP addresses. One benefit of dynamic NAT is that it is able to balance the load among the different public IP addresses.

Microsoft sells a proxy server product called Microsoft Forefront Threat Management Gateway (Forefront TMG). It was previously called Internet Security and Acceleration Server (ISA).

Understanding Proxy Servers

Instead of just using NAT, many organizations use *proxy servers*. A proxy server acts on behalf of the client computers in the internal network to retrieve web content from the Internet. A proxy server often includes NAT, but it does more.

A proxy server isn't a replacement for antivirus software within a company. However, it is useful as part of a defense-in-depth security strategy.

Proxy servers provide three important benefits:

Caching If one user requests a page from a site, the proxy server will retrieve the page and return it to the user. It also keeps a copy of the page in its local memory, or cache. If another user then requests the same page, the proxy server retrieves the page from memory and serves it to the second user. This saves Internet bandwidth since the same content doesn't have to be retrieved repeatedly.

Filtering The proxy server can use filtering lists to restrict access to certain websites. For example, if an organization wants to ensure that employees don't access gambling sites, a filter list can list these sites, and the proxy server will then block all access to these sites.

Content Checking Some proxy servers can verify that the content is valid. For example, the proxy server can check web pages for malicious content, such as

embedded malware or malicious scripts. If the web server includes a certificate for secure HTTPS pages, the proxy server can check the certificate's validity.

Consider Figure 11.2. Notice in the figure that a proxy server is between the Internet and the other computers in the intranet. This is a common configuration for many midsize and large organizations.

The proxy server will retrieve any requests that are allowed. It will block requests for pages identified in its block list.

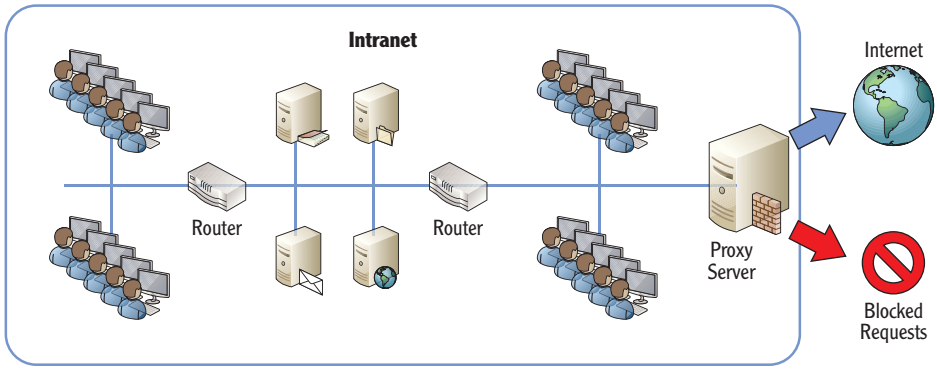

FIGURE 11.2 An intranet connected to the Internet via a proxy server

Proxy Server Filters

Some companies sell subscriptions to filter lists. These companies have web bots that constantly crawl the Web to identify content. The content is categorized, and the web pages are then added to specific lists. For example, one list might be for gambling and include all known gambling sites. Another list might be for pornography and include all known pornography sites.

Organizations can then subscribe to the different lists. These lists are added to the proxy server, and any requests to access a site on a list are blocked.

Some organizations are more proactive and create lists of only acceptable websites. If a user tries to access any website that isn't on this list, access is blocked.

For example, if a client wanted to access a web page on the Internet, the internal process would take the following steps:

1. The client computer forwards the request to the proxy server.

2. The proxy server checks the internal filter.
    a. If the page is on a block list, the request is not filled. Instead, the user will usually see a web page indicating that accessing this page is against the company policy.
    b. If the page is allowed, the web server will attempt to retrieve it from the Internet. It often uses the same NAT process shown previously in this chapter.
3. When the web page is received, the proxy server checks the content to ensure it's valid. Suspect content can be blocked with a warning to the user that the page is suspect.
4. The proxy server places valid web pages in cache. Pages in cache are served to other users from cache without retrieving them from the Internet again.
5. The web page is sent to the client that originally requested it.

Client computers need to be configured to use the proxy server. For example, most Windows computers use Internet Explorer. Figure 11.3 shows the proxy server settings on Internet Explorer. In this example, the IP address of the proxy server is 192.168.1.251, and it is listening on port 8080.

Administrators can set these settings manually or automate the settings.

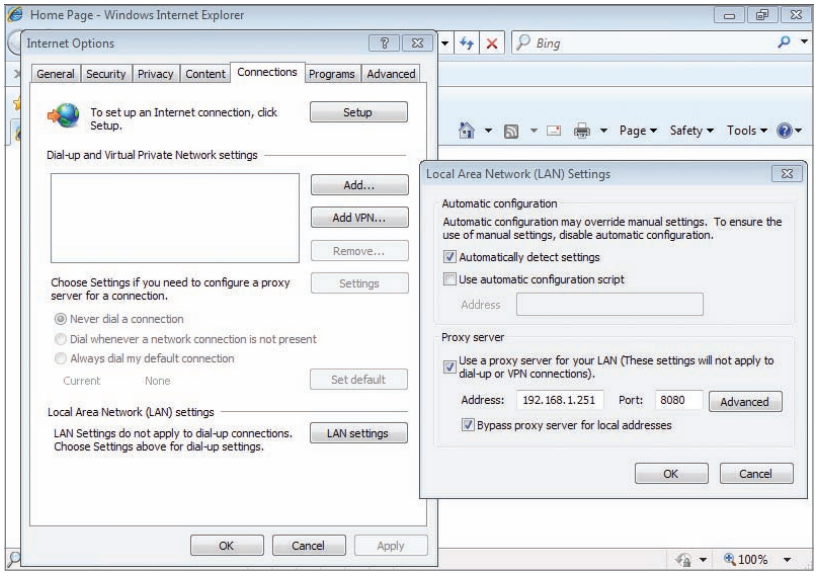

FIGURE 11.3 Configuring proxy server settings in Internet Explorer

An additional setting shown in the figure is Bypass Proxy Server For Local Addresses. This ensures that requests to web servers on the internal network (the intranet) don't have to go through the proxy server.

You can access the settings in Figure 11.3 by following these steps:

1. Launch Internet Explorer.
2. Select Tools and then Internet Options.
3. Click the Connections tab.
4. Click the LAN Settings button.

Understanding Firewalls

Chapter 2 introduced firewalls. As a reminder, a *firewall* provides protection to both networks and individual systems by controlling the traffic that can flow in or out. A host-based firewall controls the traffic for an individual host or computer. A network-based firewall controls the traffic for a network.

Microsoft's Forefront Threat Management Gateway (Forefront TMG) is a network firewall. It's an additional server product you can purchase and install on a server. Forefront TMG was previously known as Internet Security and Acceleration (ISA) Server.

Forefront TMG has multiple security purposes. As mentioned earlier, it can be used as a proxy server.

Firewalls have been widely improved over the years. The most basic firewall is simply a router with rules that define what traffic is allowed and what traffic is blocked. This is also known as a packet-filtering firewall.

Packet-Filtering Firewall A packet-filtering firewall filters packets based on IP addresses, ports, and some protocols. For example, if you want to allow only HTTP traffic (which uses port 80), you can create a rule to allow incoming traffic on port 80. If you only wanted to allow traffic through a firewall from specific computers, you could create rules based on their IP addresses.

Chapter 4 covered ports, including many of the commonly used well-known ports.

Stateful Filtering Traffic is filtered based on the state of the network connections. In other words, the firewall is able to examine packets in different conversations and make decisions based on connection states. Both TCP and UDP traffic is analyzed. If traffic isn't part of a known connection, it is blocked.

Content Filtering Some firewalls can block traffic based on the content. For example, malware is often delivered via spam embedded as a zip file and other types of attachments. Content filtering is often performed on email servers also in order to filter spam and its attachments.

Microsoft's Forefront TMG firewall performs packet filtering, stateful filtering, content filtering, and application layer filtering.

Application Layer Filtering Traffic is filtered based on an application or service. The firewall has a separate component for each application protocol (such as

HTTP or FTP) that it will filter. These firewall components examine the traffic using that protocol to allow and block certain types of traffic. For example, HTTP Get commands (which allow retrieval of documents or files) could be allowed, while Put commands (which would post documents or files) can be blocked. In practice, application layer filters are CPU intensive and used sparingly.

Most firewalls use an implicit deny policy. In other words, all traffic that has not been explicitly allowed is blocked. As an example, consider Figure 11.4. This shows a partial listing of programs and their Windows Firewall settings on a Windows Server 2008 system.

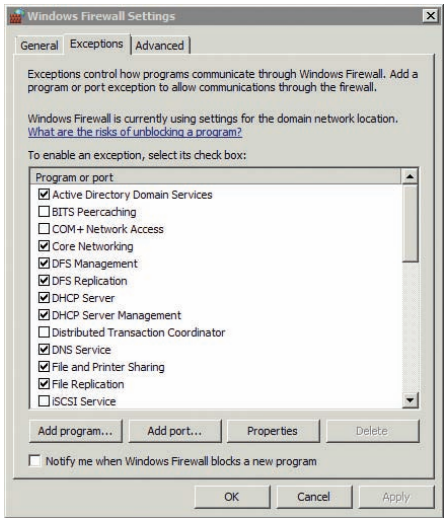

FIGURE 11.4 Allowing programs through the Windows Firewall

Each item that is checked is explicitly allowed. If an item is not selected, it is blocked.

You can access the screen shown in Figure 11.4 on a Windows Server 2008 system by following these steps:

1. Click Start ➢ Control Panel.
2. Click Security.
3. Select Allow A Program Through Windows Firewall.

Exploring the Windows Server 2008 Firewall

Today's Windows operating systems have the Windows Firewall built in as a host-based firewall. Following Microsoft's principle of secure by default, the Windows Firewall is enabled by default.

Early versions of the Windows Firewall allowed you only to create rules to control inbound packets. However, since Windows Vista and Windows Server 2008, you have been able to control both inbound and outbound traffic.

Another feature of Windows Firewall in current Windows operating systems is the use of different rules based on where your computer is operating. For example, you could have a Windows 7 computer running in a home network, in a corporate domain network, or in a public wireless networks such as a coffee shop or airport. Each of these network locations has different levels of risk. Windows sometimes automatically detects this network location. Other times, you identify it when you first connect. Either way, Windows implements firewall rules to increase or decrease security based on the network location settings.

The Windows Firewall has been included since the release of Windows XP. It has been enabled by default since Windows XP Service Pack 2 (SP2).

Network Discovery in Windows

Windows systems allow computers to discover each other. When network discovery is enabled, your computer can discover other computers on the network, and other computers can discover your computer. When it's off, it prevents other computers from seeing your computer.

Network discovery doesn't prevent connections. For example, a computer with network discovery disabled on a public wireless network will still be able to access the Internet by going through a known wireless router. However, network discovery does enable specific firewall rules, which makes it more difficult for other computers to discover a Windows computer running in a public network.

The different network locations are as follows:

Public This is a public location such as in a coffee shop or airport. Users often connect via wireless connections, and other users are completely unknown. The other users could be friendly or malicious. Attackers can try to hack into systems in a public network to steal data. Since a public network is the riskiest network location, the Windows Firewall provides the highest level of protection and helps prevent computers from being discovered on the network. Network discovery is disabled.

Home This indicates a small, protected network where you know and trust other devices on your network. Network discovery is enabled. Users in a home network can join a homegroup, which is a special type of workgroup in newer Windows operating systems.

Work This is similar to the home network location. Network discovery is enabled, and computers can discover each other. Computers can be a member of a workgroup but not a homegroup.

Domain Computers that are joined to a domain are automatically configured for a domain network location. Administrators control these settings using domain tools.

In Windows Vista and Server 2008, home and work network locations are the same and expressed as home/work. In Windows 7 and Server 2008 R2, they are separated.

There are two basic graphical user interfaces (GUIs) you can use to manipulate the firewall in Windows Server 2008. The basic GUI is in the Control Panel, and the second tool is the Windows Firewall with Advanced Security GUI which is located in the Administrative Tools section.

Figure 11.5 shows the Control Panel view of the firewall in a Windows Server 2008 R2 system. Notice that the connection for the domain networks is Connected. This indicates the computer is joined to a domain and that the firewall is using the settings for a domain. You can also see that the firewall is On, and it's configured to block all incoming connections that haven't been explicitly allowed.

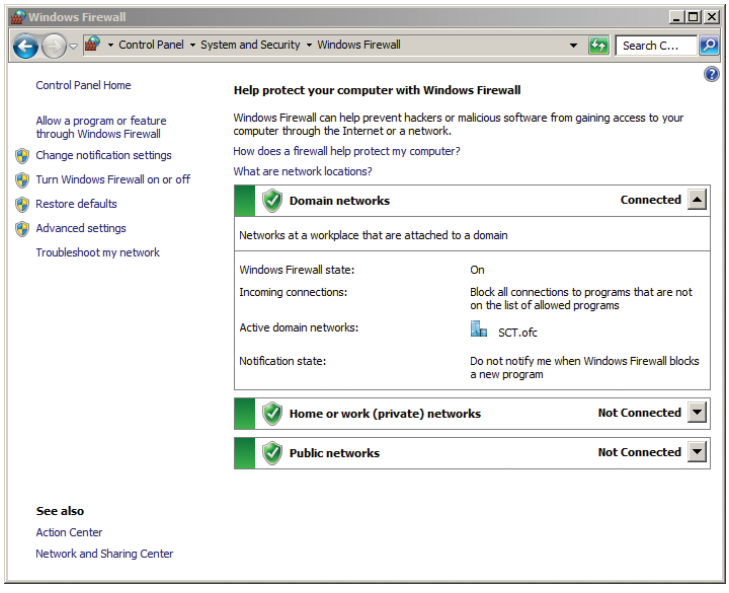

FIGURE 11.5 Basic Windows Firewall GUI in Windows Server 2008 R2

You can access the screen shown in Figure 11.5 by following these steps:

1. Click Start, and select Control Panel.
2. Click Check Firewall Status.

Figure 11.6 shows the Windows Firewall with Advanced Security GUI in Windows Server 2008 R2 with the New Inbound Rule Wizard started. The inbound rules are selected, and all the rules with a green circle are enabled to allow the traffic. The ones that are grayed out are not enabled. Notice in the left pane that there are also outbound rules and connection security rules that you can manipulate.

Connection Security Rules use Internet Protocol Security (IPSec). IPSec can encrypt data traveling on the wire.

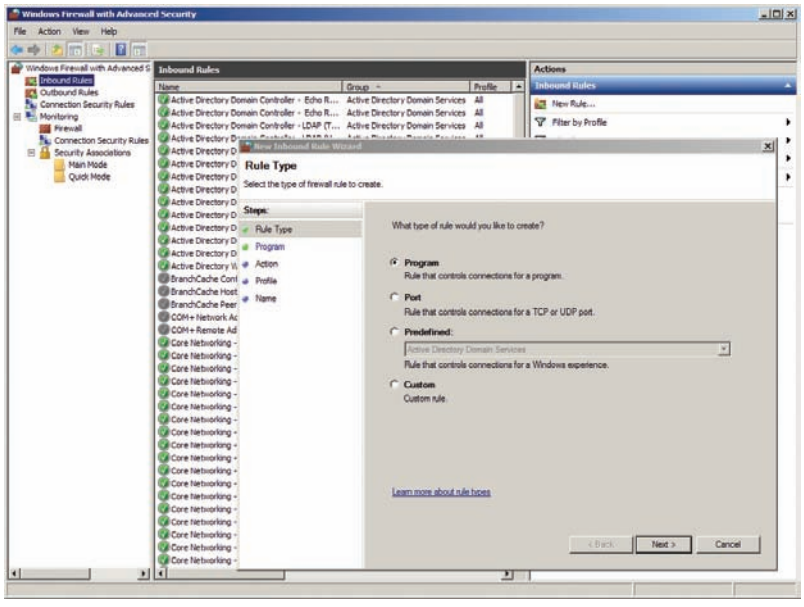

FIGURE 11.6 Windows Firewall with Advanced Security GUI in Windows Server 2008 R2

Although Windows Server firewalls include many built-in rules, you can also add your own rules. In the figure, the New Inbound Rule Wizard was started by clicking New Rule in the Actions pane (on the right).

You can access the screen shown in Figure 11.6 by following these steps:

1. Click Start ➢ Administrative Tools ➢ Windows Firewall With Advanced Security.
2. Select Inbound Rules.
3. Click New Rule in the Actions pane (on the right).

Identifying a Perimeter Network

Internet-facing servers are any servers accessible from the Internet. They include web servers, mail servers, FTP servers, and more.

A perimeter network is often called a *demilitarized zone* (DMZ) or a buffer zone. This is especially true when it's created with two firewalls.

A *perimeter network* is an area between the Internet and an intranet that hosts servers accessible from the Internet. It provides a layer of security protection for these Internet-facing servers and isolates these servers from the internal network.

Consider Figure 11.7. It shows a perimeter network hosting a web server and a mail server. Notice that the perimeter network is between two firewalls. This is a common configuration, but there are others.

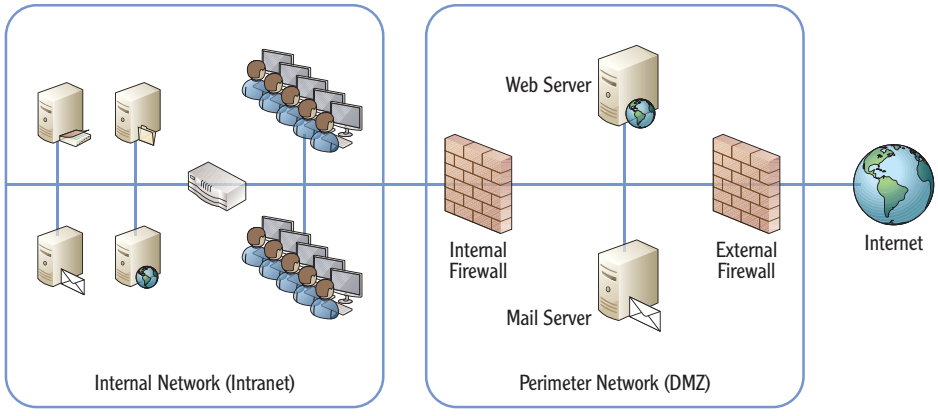

FIGURE 11.7 Using a perimeter network

An important point to realize about the perimeter network is that servers placed here are accessible from anywhere on the Internet to anyone who has access to the Internet. However, the perimeter network does provide protection.

As an example, consider the web server. A typical web server serves web pages using HTTP on port 80 and HTTPS on port 443. The external firewall will filter traffic to this web server and can block all traffic to this server that isn't using either port 80 or port 443. This can prevent many potential attacks from ever reaching the server.

From a risk perspective, the perimeter network is a little safer than the Internet. However, since servers in the perimeter network are still accessible from anywhere on the Internet, there is still a significant amount of risk, especially when compared with the intranet. Additionally, if a server in the perimeter network is compromised, the internal firewall will protect resources on the intranet.

You can also create a perimeter network with just a single firewall. Figure 11.8 shows an example of perimeter network created with just a single firewall. Notice

that the mail server and web server are still isolated from both the Internet and the intranet. The firewall controls what packets can reach the perimeter network and what data can reach the intranet.

Although this configuration is less expensive since only a single firewall is used, it's also much more complicated to configure. An administrator must configure rules to route traffic to specific NICs. Since these rules are more complex than the rules for two firewalls, there's a greater chance of error.

A significant benefit of a two-firewall perimeter network is that you can use two separate vendors. For example, one firewall can be Microsoft's Forefront TMG firewall, and another firewall can be from another vendor. Although vulnerabilities may occur in any system, it's unlikely that both firewalls will be vulnerable at the same time. Also, although an attacker may be an expert on either of the firewalls, it's less likely that an attacker will be an expert on both at the same time.

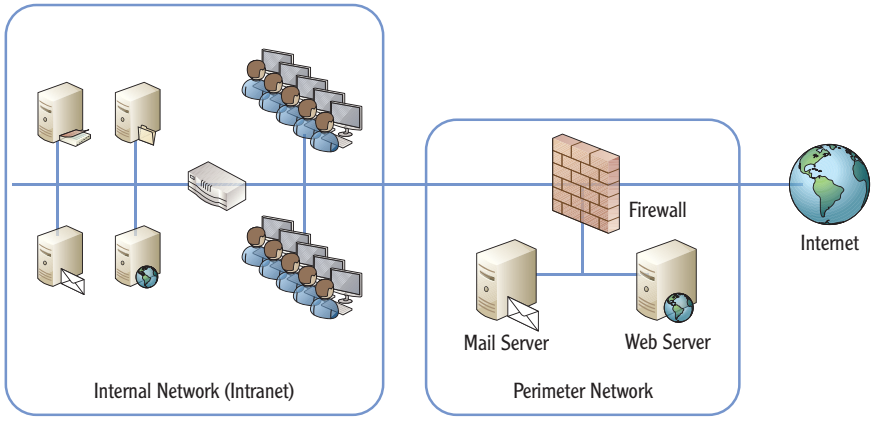

FIGURE 11.8 A single firewall perimeter network

Understanding a Reverse Proxy Server

Some organizations implement *reverse proxy servers* to increase security and performance of web servers. A reverse proxy server is an additional server in the perimeter network. It isolates these web servers from direct access on the Internet, providing a layer of protection from Internet attackers.

Consider Figure 11.9. This shows a reverse proxy server used with a web server. The reverse proxy server receives the requests from the clients and forwards them to the web server. The web server sends the web pages back to the proxy server, and the proxy server sends them to the clients.

Just as a regular proxy server can cache requests, the reverse proxy server can also cache requests. This reduces some of the load on the web server. However, since all the pages are still served over the Internet link, it doesn't reduce Internet usage.

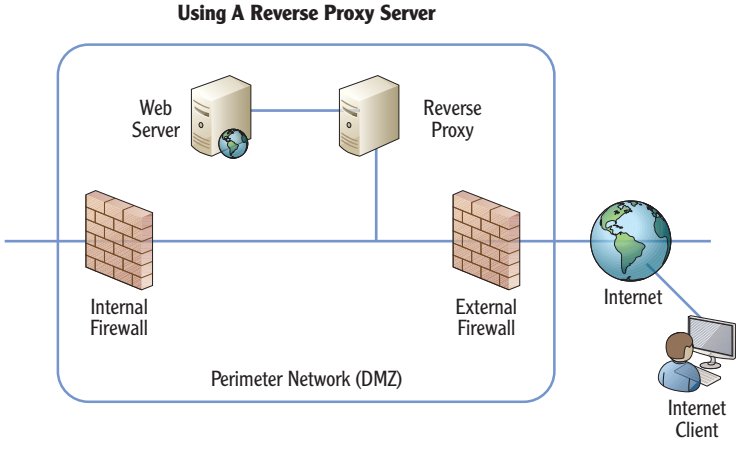

FIGURE 11.9 A single firewall perimeter network

Clients don't need to be configured to use a reverse proxy. Indeed, clients will rarely ever know a reverse proxy is in use. It is simply transparent to the end users.

Understanding Guest Networks

Guest networks are another type of perimeter network used by larger organizations. A guest network is an isolated portion of the internal network that can be used by guests or visitors.

Depending on how the guest network is configured, visitors may not need to provide any credentials to access the guest network. However, their access on the network is usually very limited. The primary access that is usually granted from a guest network is Internet access.

Guest networks are also becoming popular in home wireless networks. For example, Cisco's Valet Wireless Router allows you to create a separate password you can give to visitors without giving the primary password that is used for other connections. When the visitor leaves, you can change the visitor password or disable visitor access.

Figure 11.10 shows the Cisco Valet Wireless router's Guest Access Settings page. Notice you can also modify how many guests can connect.

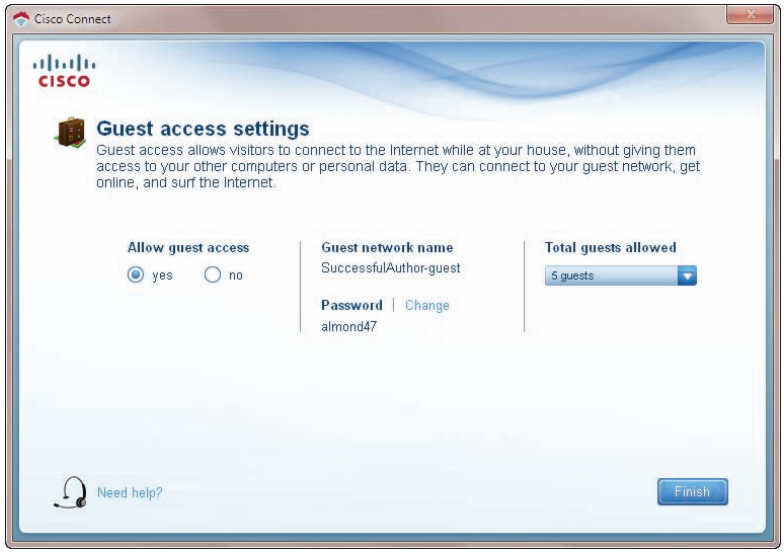

FIGURE 11.10 Guest network configuration

Understanding Extranets

An *extranet* is an area between the Internet and an intranet that hosts resources for trusted entities. These resources are available via the Internet. An extranet is often physically the same as a perimeter network. The difference is in the intent and the scope of access and resources that are made available. Specifically, an extranet is configured so that only trusted partners or customers have access to a company's resources in the extranet. These trusted partners typically need access to areas of a company's network such as private websites or databases that would not be accessible publically. This allows the company to extend access to their internal resources to trusted entities outside the intranet.

Figure 11.11 shows a drawing of an extranet. You may notice that this looks very similar to Figure 11.7. However, keep in mind that the difference between an extranet and a perimeter network is based on the intent. The perimeter network hosts servers that are accessible to any Internet clients from anywhere on the Internet. Extranets are available only to specific clients.

For example, a boating parts company sells and ships parts to boat builders. The parts company may want some customers to be able to access their accounts, check availability of parts, place orders, and track status. It can add a web server to an extranet and restrict access to specific customers.

Extranets are often created to share data between two companies that have business relationships or partnerships.

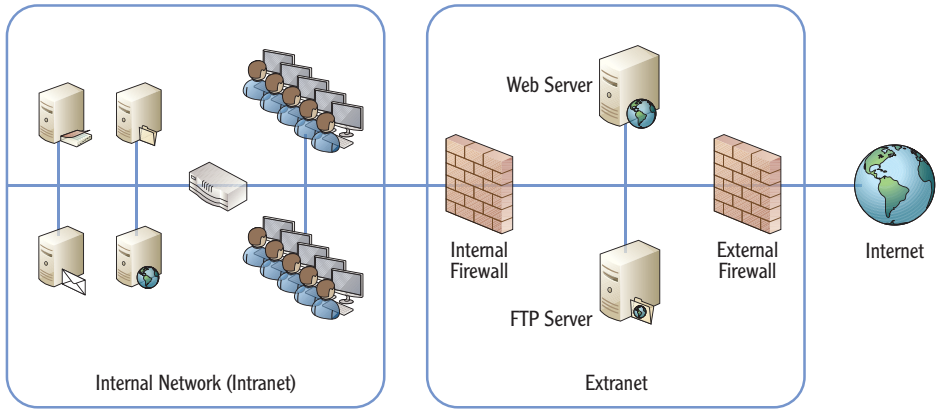

FIGURE 11.11 Using an extranet

By only allowing access to their web server via an extranet, they can control who is granted access to the extranet's website. This prevents unwanted users (such as competitors) from viewing information the company doesn't want to make public.

The Essentials and Beyond

This chapter covered basic security zones in networks. The Internet is the riskiest security zone. Any resources placed directly on the Internet are accessible from anywhere in the world and are subject to attack from anywhere in the world, as long as the attacker has access to the Internet. The intranet is an internal network and is considered the safest zone when compared to other zones. Firewalls typically separate the intranet from the Internet. Microsoft desktop and server operating systems include host-based firewalls built into the operating system. Additionally, Microsoft sells a network-based firewall server product called Forefront TMG. A perimeter network (also known as a DMZ) usually includes two network-based firewalls, and Internet-facing servers are placed between the two firewalls. The firewalls control traffic to and from resources in the perimeter network. Extranets are perimeter networks created to provide access to internal resources to specific trusted entities. Guest networks are perimeter networks created to provide temporary network access to visitors.

- Determine whether your computer is using a proxy server.
- Determine whether a software firewall is enabled on your computer.
- The network location determines what firewall rules are enabled on Microsoft operating systems. Determine what network location your computer is using.
- Draw a perimeter network, and draw an extranet. Describe the differences.

(Continues)

The Essentials and Beyond *(Continued)*

To compare your answers to the author's, please visit **www.sybex.com/go/networkingessentials**.

Review Questions

1. Which network security zone represents the highest risk?

   A. Internet **C.** Perimeter network
   B. Intranet **D.** Extranet

2. What service translates private IP addresses to public IP addresses and translates public IP address back to private?

3. An organization wants to restrict which web pages employees can access on the Internet using company computers. What should be implemented?

   A. NAT **C.** Proxy server
   B. Firewall **D.** Reverse proxy server

4. True or false. A DMZ provides a layer of security for Internet-facing servers.

5. How many firewalls are used to create a perimeter network? (Choose all that apply.)

   A. One **C.** Three
   B. Two **D.** Four

6. What allows computers to locate each other in a Microsoft network?

   A. Firewall **C.** Network discovery
   B. Public network location **D.** Proxy server

7. You want to provide access to some internal resources to a business partner via the Internet. No one else should have access. What should you create?