Firewall Design

As mentioned at the beginning of the chapter, a firewall is a device or devices that control traffic between different areas of your network. In a more robust design you typically see two or three firewall devices, as well as many other security components to protect company resources. In a firewall design, I refer to the security solution as a firewall system, indicating that many devices are being used to protect your resources.

As you will see in this section, you should follow some practical guidelines when developing a firewall system. These can include packet and application firewalls, application gateway and ATFs, host-based firewalls, and, more than likely, hybrid firewalls, as well as many other security devices, such as VPN concentrators, IDS devices, authentication security servers, and many other components.

After briefly covering the components in a security solution, I discuss some various designs that commonly are used to protect resources. Then I discuss their advantages and disadvantages and cover management issues.

Design Guidelines

You should follow five basic guidelines when designing a firewall system:

- Develop a security policy.
- Create a simple design solution.
- Use devices as they were intended.
- Implement a layered defense to provide extra protection.
- Consider solutions to internal threats that should be included in your design.

The following subsections cover these five key design points.

Developing a Security Policy

One of the first things you do when designing a firewall system is to create a security policy. The policy should define acceptable and unacceptable behavior, should state restrictions to resources, and should adhere to the company's business plan and policies. Without a security policy, it is practically impossible to develop a security solution that will meet your company's needs.

The key to a good design is basing it on a security policy. Basically, a policy defines who is allowed to access resources, what they are allowed to do with resources, how resources should be protected (in general terms), and what actions are taken when a security issue occurs. Without a security policy, it is impossible to design a firewall system that will protect your assets. In other words, if you don't have a security policy, what should you protect? How much should you protect resources? Who is allowed to access resources? If your policy does not define these items, it is hard to design and implement a solution based on hunches. Actually, without a security policy, the firewall system that you put in place might be creating a security risk: It might not be providing adequate protection to your company's resources.

Designing a security policy is beyond the scope of this book. However, it minimally should address the following items:

- The resources that require access from internal and external users
- The vulnerabilities associated with these resources
- The methods and solutions that can be used to protect these resources
- A cost-benefit analysis that compares the different methods and solutions

Designing Simple Solutions

A firewall system design should be kept simple and should follow your security policy. The simpler the design is, the easier it will be to implement it, maintain it, test and troubleshoot it, and adapt it to new changes. Many people like to call this the KISS principle: Keep it simple, stupid. The last kind of problem you want to deal with is a design or configuration error that leaves your network open to all different kinds of attacks.

CAUTION

Complex solutions are prone to design and configuration errors, and are difficult to test and troubleshoot. The simpler you can make the design, the easier it will be to manage it.

Using Devices Correctly

Network devices have functional purposes; they were built with a specific purpose in mind. For example, a Layer 2 switch is used to break up a collision or bandwidth domain, and it also uses VLANs to break up broadcast domains: It is typically not a good device to use to filter traffic because the filtering is done by creating filtering rules based on MAC addresses. The problem with this approach is that MAC addresses tend to change quite a bit: NICs fail, PCs and servers are upgraded, devices are moved to different locations in the network, and so on. Filtering is done best when logical addressing is deployed.

Using the wrong product to solve a security problem can open you to all kinds of security threats. For example, assume that you want to use an IDS to detect different kinds of network threats. You notice that your Cisco router has the capability in the Cisco IOS Firewall feature set, and you decide to enable it, feeling secure that your Cisco router will generate an alarm when an attack occurs. If you had taken time to read the security material related to Cisco routers, you would have realized that Cisco routers can detect only a few dozen different kinds of networking attacks (typically, the most common ones). Therefore, for all the other hundreds of kinds of attacks, your Cisco router will not be capable of detecting them, leaving you exposed. In this example, a better solution would have been to purchase an IDS solution that can detect hundreds of different kinds of attacks.

Using Old Junk to Solve New Problems

As a consultant, I repeatedly hear from customers that they want an inexpensive (that is, "cheap") solution that reuses a lot of the networking gear that they already have in their current design. Unfortunately, with a network that has equipment that is more than 5 years old, this presents a problem because that technology is outdated and probably useless.

When designing a network, especially as it relates to security, you first must come up with a simplified design that describes the basic components that are required. Then you must determine which actual products you will use for these components, based on the features and functions you need. Never try to approach this backward by trying to fit a product to your design: You create problems by doing this instead of solving them. Also remember that there is no such thing as a networking product that will do everything: Do not force yourself to cut corners by trying to make one device do everything.

Creating a Layered Defense

A security design typically uses a layered defense approach. In other words, you usually do not want one layer of defense to protect network. If this one layer is compromised, your entire network will be exposed.

The French and a One-Layer Defense

The French learned the one-layer defense limitation the hard way during World War II with the Maginot line, a physical barrier between Germany and France. The Maginot line extended only the distance of the border between France and Germany. The Germans were smart: They went north and west and came down through the Netherlands, bypassing this fortification and allowing them to easily conquer France.

Instead, you should use a multilayer defense in your firewall system design. With multiple layers, if one layer is compromised, you still have other layers behind it protecting you.

One of the examples that I like to use to describe a firewall system is the fortification systems that kings used to protected their castles in medieval times. Figure 2-22 shows an example of this. In this figure, the first line of defense is the moat surrounding the castle. For the second line of defense, spearmen are behind the moat, preventing anyone from trying to swim across it. Behind the spearmen is the third line of defense: the castle wall, which can be as little as 3 meters high but typically was much higher than this. On the wall are swordsmen, providing the fourth layer of defense. Inside the castle wall are the castle grounds and the castle itself. The castle is built with very high stone walls and turrets, providing the fifth layer of defense. And in the windows of the wall and on top of the turrets are the archers, providing the last layer of defense. As you can see from this system, an attacker must go through many layers of defense to capture or kill the king.

Figure 2-22. A Medieval Firewall System

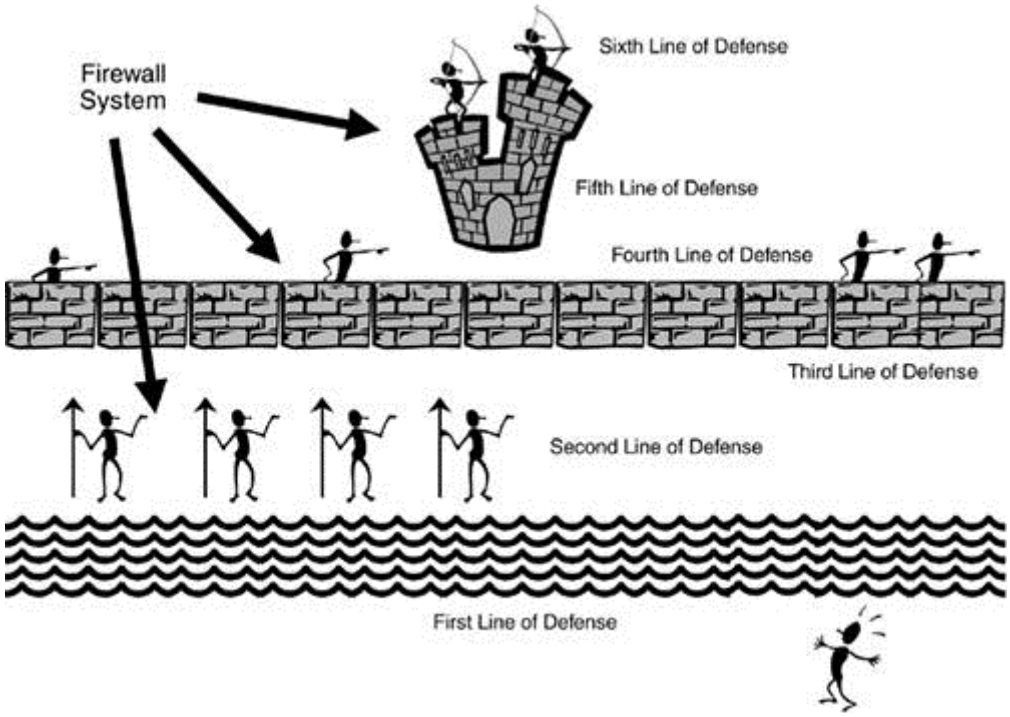

Dealing with Internal Threats

Too often, security personnel are concerned about protecting a company's resources and assets from outside threats. Remember that it is much easier to attack your assets from within; plus, most threats and attacks (60 to 70 percent) are internal attacks. Therefore, a good firewall system not only protects you from external threats, but also allows you to minimize internal threats.

DMZ

Most firewall systems use a demilitarized zone (DMZ) to protect resources and assets. A DMZ is a segment or segments that have a higher security level than that of external segments, but a lower security level than that of internal segments. DMZs are used to grant external users access to public and e-commerce resources such as web, DNS, and e-mail servers without exposing your internal network. A firewall is used to provide the security-level segmentation among the external, DMZ, and internal resources. Basically, the DMZ acts as a buffer between different areas in a network.

DMZ Rules and Traffic Flow

To help enforce security more easily, each area in the firewall system is assigned a security level. This could be something as simple as low, medium, and high, or something more sophisticated, such as a number between 1 and 100, where 1 is the lowest security level and 100 is the highest. Typically, traffic from a more secure (higher) layer is permitted to a lower layer, but not vice versa.

For traffic to go from a lower layer to a higher layer, it must be permitted explicitly: In other words, you must set up a filtering rule that allows this traffic to go from a lower level to a higher level. If two areas have the same security level, such as medium, the traffic between the two areas is either permitted or denied, based on the process that the product uses.

CAUTION

When allowing traffic to go from a lower security level to a higher one, you should be very specific about what traffic is allowed. For example, if you want Internet users to access a web server, specify both the destination IP address of the web server, such as 200.1.1.2, the protocol (TCP), and the destination port number (80), in the filtering rule. By being very specific, you are opening a hole in the firewall system for only traffic that is necessary; all other traffic (including other types of traffic to the web server) is blocked by the configuration of your security levels.

It is important to point out that these rules sometimes are built into the firewall or must be configured manually on the firewall. In either situation, this gives you a lot of flexibility in assigning a level of threat to particular areas in your network and configuring your firewall system appropriately.

NOTE

The Cisco PIX firewall uses a numbering scheme approach to security levels, providing a very flexible system for separating areas by their level of threat. Higher-number areas freely can send traffic to lower-number areas, but not vice versa, and areas with the same level cannot send traffic to each other. On Cisco IOS routers, you must configure filtering rules manually to define levels of security.

The network shown in Figure 2-23 illustrates how security levels work.

Figure 2-23. Security Level Example

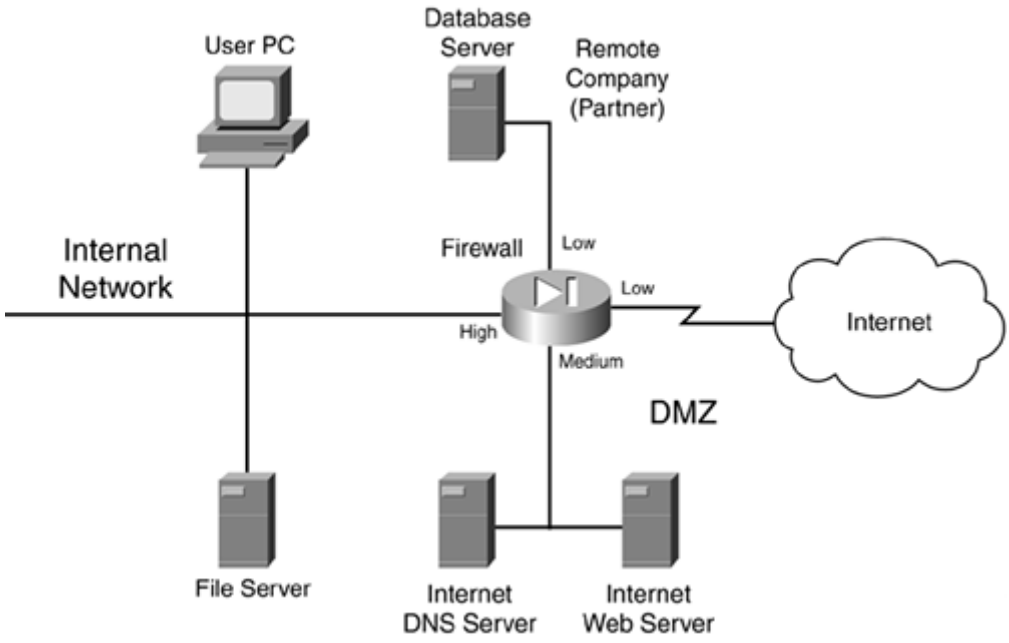

In this example, a firewall is used to separate different areas of a network. The firewall has the following four interfaces:

- A connection to the Internet, assigned a low security level
- A connection to the DMZ, where public servers are located, assigned a medium security level
- A connection to a remote company that is working on a project for them, assigned a low security level
- A connection to the internal network, assigned a high security level

This company has assigned the following rules:

- High- to low-level access: permit
- Low- to high-level access: deny
- Same-level access: deny

Given these rules, the following traffic is allowed automatically to travel through the firewall:

- Internal devices to the DMZ, the remote company, and the Internet
- DMZ devices to the remote company and the Internet

Any other type of traffic flow is restricted. One advantage of this design is that, because the remote company and the Internet are assigned the same level, traffic from the Internet cannot reach the remote company (which provides protection), and the remote company cannot use your Internet access for free (which saves you money).

Another advantage of security levels is that they create a layered approach to security. For example, for either hackers on the Internet or the remote company to access internal resources in

Figure 2-23, they probably first must hack into your DMZ and then use this as a stepping stone to hack into your internal network. Using this layered approach, you make the hacker's job much more difficult and your network much more secure.

DMZ Types

DMZs come in many types of designs. You can have a single DMZ, multiple DMZs, DMZs that separate the public network from your internal network, and DMZs that separate traffic between internal networks. The following sections show some of these implementations.

Single DMZ

Single DMZs come in two types:

- Single segment
- Service-leg segment

Figure 2-24 shows an example of a single DMZ with a single segment. In this example, there are two firewalls: a perimeter firewall and a main firewall, with the DMZ segment between the two. One disadvantage of this design is that two firewalls are needed: one to protect the DMZ from the Internet and one to protect the internal network from the DMZ and the Internet.

Figure 2-24. Single DMZ with a Single Segment

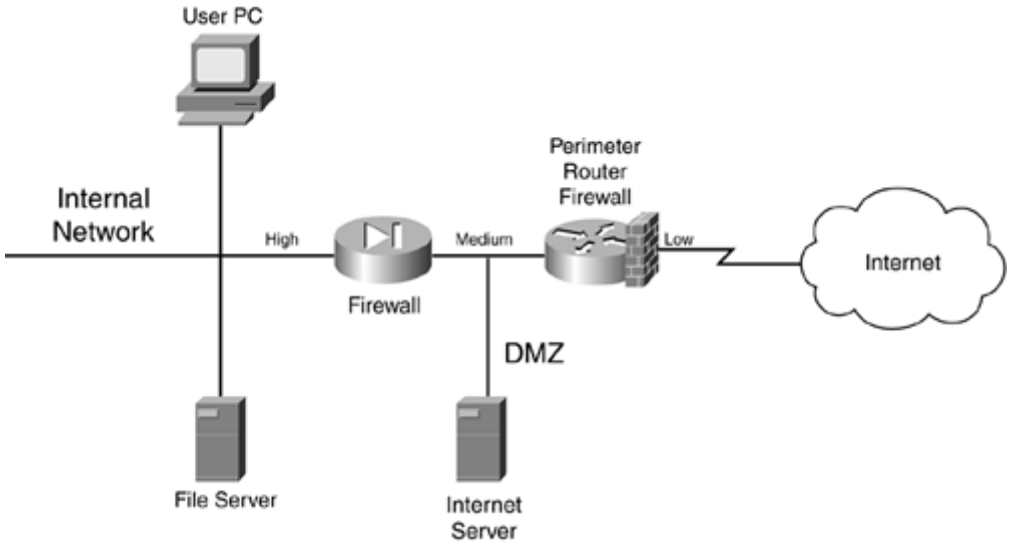

Most firewall designs use a service-leg DMZ, which is shown in Figure 2-25. In this example, a router is used to connect to the Internet. The design in Figure 2-25 has two advantages over the single-segment DMZ shown in Figure 2-24:

- The firewall sometimes can be connected directly to the Internet, removing the extra cost of the perimeter router.
- All security-level polices can be defined on one device (in a single-segment DMZ, you must define your policies on two devices).

Figure 2-25. Single DMZ with a Service-Leg Segment

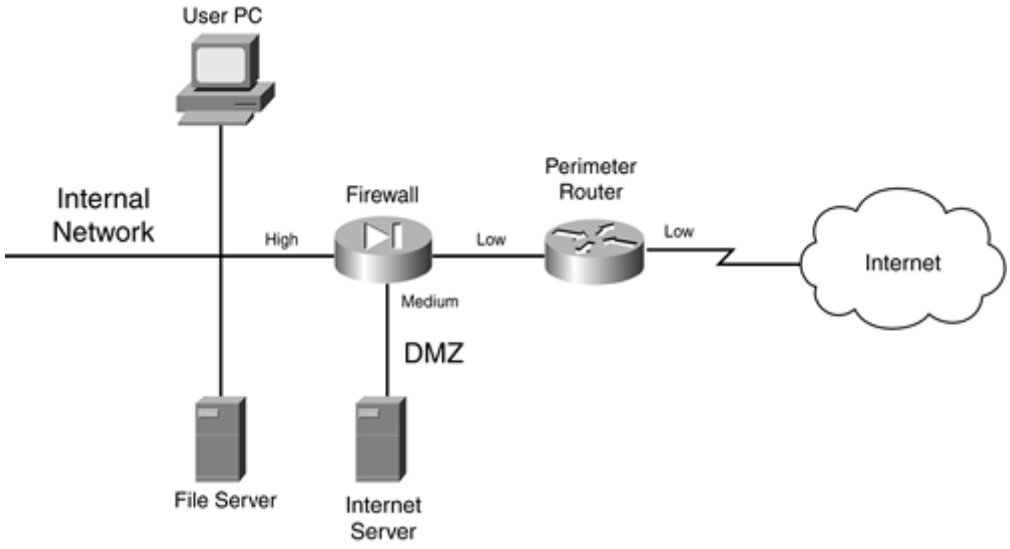

The main problem with this approach, however, is that, because a single firewall is handling all security policies, a successful DoS attack can degrade the firewall's performance. In the best case, only your throughput is affected; in the worst case, the firewall might crash. With a single-segment DMZ, because the policies are spread between the two firewalls, there is less likelihood of an overload occurring.

This is especially true for traffic between the DMZ and the internal network. If a hacker is attacking a web server on the DMZ, the perimeter firewall takes the brunt of this attack, which allows the internal firewall to handle DMZ/internal traffic without affect.

Multiple DMZs

A firewall system can be used to separate multiple areas of your network, including multiple DMZs. Figure 2-26 shows an example of a network with multiple DMZs. In this example, a firewall is used to break up a network into four areas: the internal network, a DMZ for Company A's server, a DMZ for Company B's server, and the Internet.

In this example, the internal network is assigned a high security level, both company servers are assigned a medium security level, and the Internet is assigned a low security level. Assume that high-to-low access is allowed by default but that same-to-same is denied. In this example, the internal network can access any resource, and each company server can access the Internet, but not the other company's server. You would need to set up security rules to allow Internet access into the two servers on the medium-security segments, as well as communication between the two servers, if this is desired.

Figure 2-26. Multiple DMZ Example

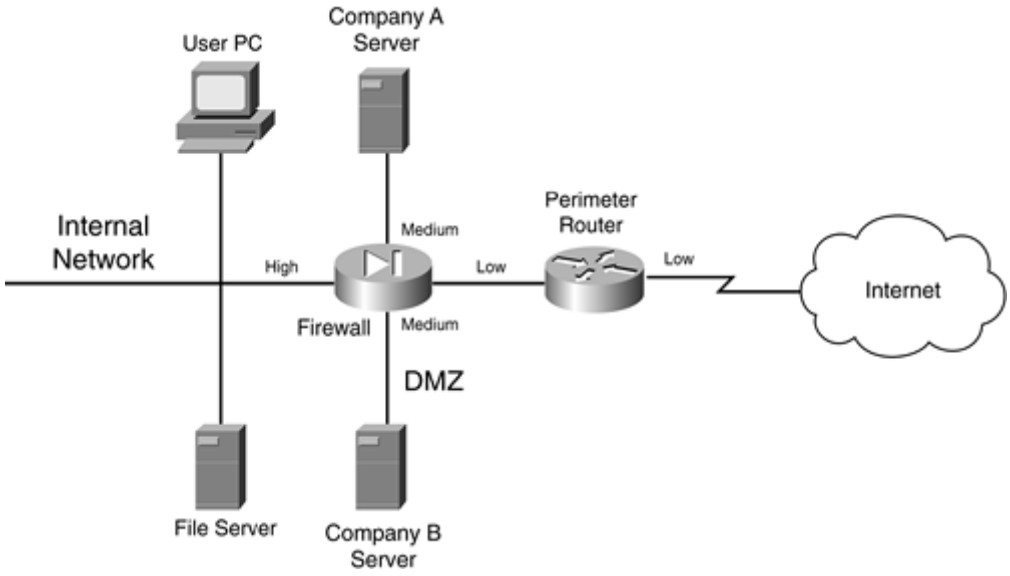

Actually, this type of design is very common in ISPs. Most ISPs offer web-hosting services and use this type of design to separate each company's server(s) from the others.

Internal DMZ

Another type of DMZ is an internal one. An internal DMZ enables you to provide separation between different parts of your internal network. Most people assume that a DMZ is used to separate your internal services from those that you offer to the public, such as a web or e-commerce server; however, they can be used effectively to protect resources in one part of your company from another.

Figure 2-27 illustrates the usefulness of an internal DMZ. In this example, a firewall is used to separate the internal network (in the right of the figure) from the engineering and accounting users.

In this example, both engineering and accounting are assigned a medium level of security. Assume that same security level to the same security level is denied by default; in other words, if two interfaces have the same security level, they cannot, by default, communicate with each other.

With this configuration, accounting and engineering are allowed to send traffic to the internal network but not themselves. Internal users cannot access either of these two groups because the internal users are on a lower security level interface.

To allow these last types of access, you would need to configure security rules on your firewall to allow same-to-same or low-to-medium levels of access.

Figure 2-27. Internal DMZ Example

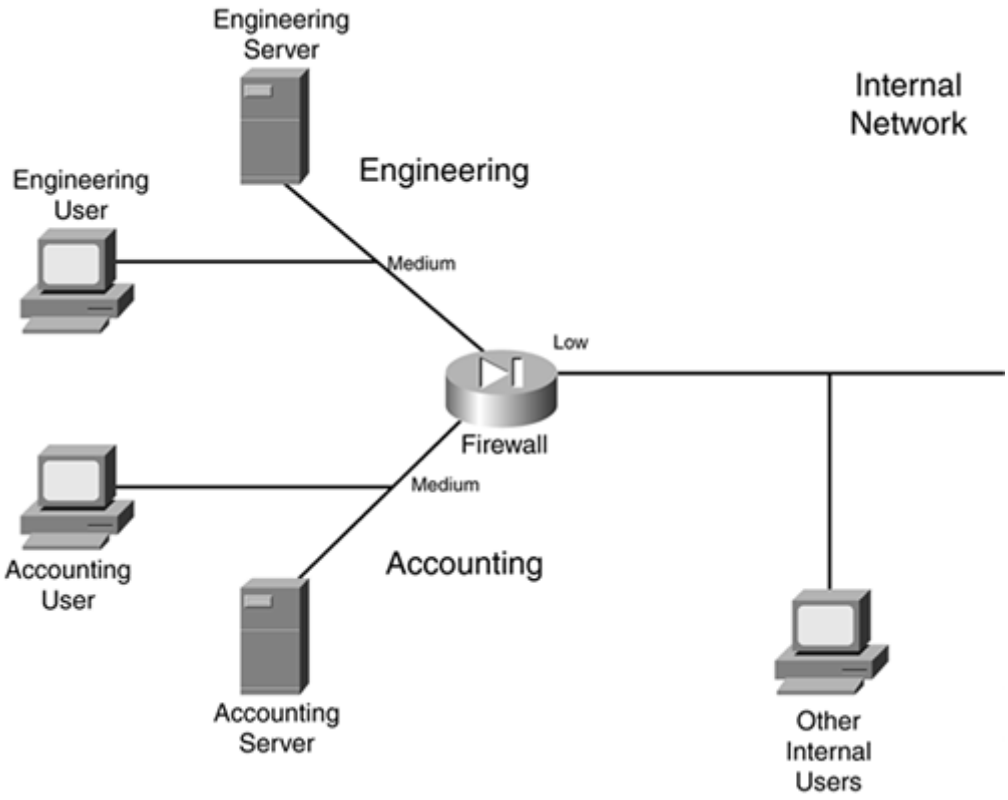

As you can see in this example, you easily can protect important resources in your network from other internal users. Internal DMZs enable you to accomplish the following:

- Control traffic between areas
- Localize security problems

I already have discussed the first bulleted point. For the second, imagine that a hacker somehow could compromise your perimeter defenses and access an internal resource such as a file server. By using an internal DMZ, you still are protecting important resources. In Figure 2-27, the accounting and engineering resources are protected from the hacker if an internal resource connected to the lower-level internal interface was compromised.

Components

Now that you have a basic understanding of firewall system design practices, this section takes a closer look at the components that make up a firewall system. A good firewall system typically contains the following components:

- Perimeter router
- Firewall
- VPN
- IDS

It is important to point out here that I used the word component, not device, to describe what is included in a firewall system. This is because many devices can support multiple components. For example, a Cisco IOS router can act as a perimeter router or a main firewall, can terminate VPN connections, and can perform IDS?all in one chassis. Remember my caution, though, if you are using a single layer for defense: Using multiple layers provides for a better design and protection. The following sections cover the different components used in a firewall system.

Perimeter Router Component

The main purpose of the perimeter router is to provide a connection to a public network, such as the Internet, or a different company. It is used to convert data-link layer media types from a LAN to either a WAN or MAN medium.

The functions of the perimeter router can include the following:

- Routing through static routes or a dynamic routing protocol
- Filtering through either packet filtering or stateful filtering
- Terminating VPN connections
- Providing address translation

TIP

Remember that the main function of the perimeter router is to access a different network: It is your first, but not your last, line of defense. Therefore, you do not want to overload it with a bunch of security functions that another component in the firewall system can handle better.

Firewall Component

The main purpose of the firewall component is to separate your network into different security levels and control traffic between these levels. Typically, you find a firewall component near the perimeter of the network, protecting you from external threats as well as providing controlled access to a public DMZ segment. However, you also might find firewall components inside your internal network, separating critical resources so that they are better protected.

The functions of the firewall can include the following:

- Stateful filtering
- User authentication of connection with CTPs
- Connection filtering with CGFs
- Address translation

VPN Component

The main purpose of the VPN component is to provide a protected connection between two devices, two networks, or a device and a network. This protection can include encryption, authentication, and packet-integrity checking, preventing eavesdrop attacks from prying eyes. VPNs are a cost-effective, remote-access solution because they enable you to use a public network, in a secure manner, to connect two private networks. This is cheaper than purchasing private WAN links to provide connectivity.

The functions of the VPN component can include the following:

- Protecting (encrypting) traffic between LAN sites and remote access users
- Assigning addressing information to remote access clients
- Using simple packet filters to restrict traffic flow

IDS Component

The main purpose of the IDS component is to detect, and possibly prevent, reconnaissance, DoS, and unauthorized access attacks. To understand the different kinds of network attacks that your company is facing, you need an intimate understanding of the different kinds of traffic flowing through your network, as well as the intentions of this traffic.

Most traffic entering or traversing your network has a valid purpose: to access web pages with HTTP, resolve names to addresses with DNS, send e-mail with SMTP, and so on. However, a small percentage of traffic has malicious intentions. In these cases, a hacker might be executing a reconnaissance attack to determine what kinds of resources are available in your network, and then might execute a DoS attack to affect their level of service or carry out an unauthorized attack to open a back door into your network. An IDS solution should be capable of detecting these kinds of threats.

IDS Overlooked

IDS components often are overlooked or are seen as unnecessary in a security solution if your network has a firewall component. I once performed work for a company that took this position on IDS solutions, and it was shaken when I set up a test IDS system that detected more than 1000 network threats directed against the network system over a single day. You will not know what kinds of attacks your network is facing unless you monitor it. Plus, new types of attacks are appearing at a steady rate. A good IDS solution can help detect and prevent attacks.

IDS components fall under one of three categories:

- Anomaly-based
- Signature-based
- Hybrid-based

Anomaly-based solutions capture traffic over a period of time and use this as a reference for what is valid. These systems then compare new traffic to what is considered to be "normal" and look for anomalies. One disadvantage of anomaly-based solutions is that they tend to generate a lot of false positives (they trigger alarms that are not really attacks). This is because traffic patterns change; if you do not stay up-to-date on a database of normal traffic flows, false positives are bound to happen.

Signature-based solutions compare traffic to signatures to look for attacks. A signature is a static definition of things to examine in packet or packets; these can include header information as well as data. Signature-based solutions have a lot less false positive alarms. However, their main disadvantage is that they cannot detect new kinds of attacks unless you keep your signature database updated. This is where an anomaly-based solution shines: It can detect new kinds of attacks without a software upgrade.

In many of today's IDS solutions, a hybrid approach is used, with both signatures and anomaly detection used in tandem to provide the best possible intrusion detection.

IDS solutions come in two flavors:

- Network-based IDS
- Host-based IDS

A network-based IDS solution is a protocol analyzer on steroids: It plugs into your network at key points and monitors traffic. Network-based systems can be used to detect attacks against many different devices.

A good network-based IDS solution should have the capability, when an attack is detected, to access (log into) your firewall component and configure a temporary filter to block the malicious traffic. This is an excellent tool for shutting down the hacker's access even into public areas of your network.

A host-based IDS solution is IDS software running on a host, such as a PC or file server, that detects attacks only against that host. This can provide an additional measure of protection for critical servers that are not necessarily protected by a firewall or IDS component. One downside of host-based solutions is that they require extra processing power to examine packet information sent to the host.

CAUTION

IDS solutions are still in their infancy and should not be relied upon solely for security. For many IDS systems, a very wily hacker can slip by attacks without raising an alarm. Remember that a good firewall system has multiple components to it.

Your IDS component should play a key role in your firewall system. The functions of the IDS component can include these:

- Monitoring traffic for statistical purposes
- Examining traffic for network threats
- Reporting network threats and possibly taking action to prevent the threats

TIP

IDS systems can generate a lot of logging information. You should review and analyze this carefully daily. By doing this, you will gain an excellent understanding of the traffic patterns in your network, as well as what hackers are currently up to.

Component Placement

Now that you understand some of the components used in a firewall system, this section talks about where these components are placed in a network design. As you will see, you can design a network in many different ways, each with advantages and disadvantages.

NOTE

It is important to point out that there is no one correct way to design a secure network. Each network is unique, and its unique characteristics must be taken into account when designing a solution. However, the examples shown here are a good starting point for creating the appropriate design for your network.

Simple Firewall System Design

To help understand where components are placed in a firewall system, I use two examples. The first example, shown in Figure 2-28, is the simpler of the two designs.

Figure 2-28. Simple Firewall System Design

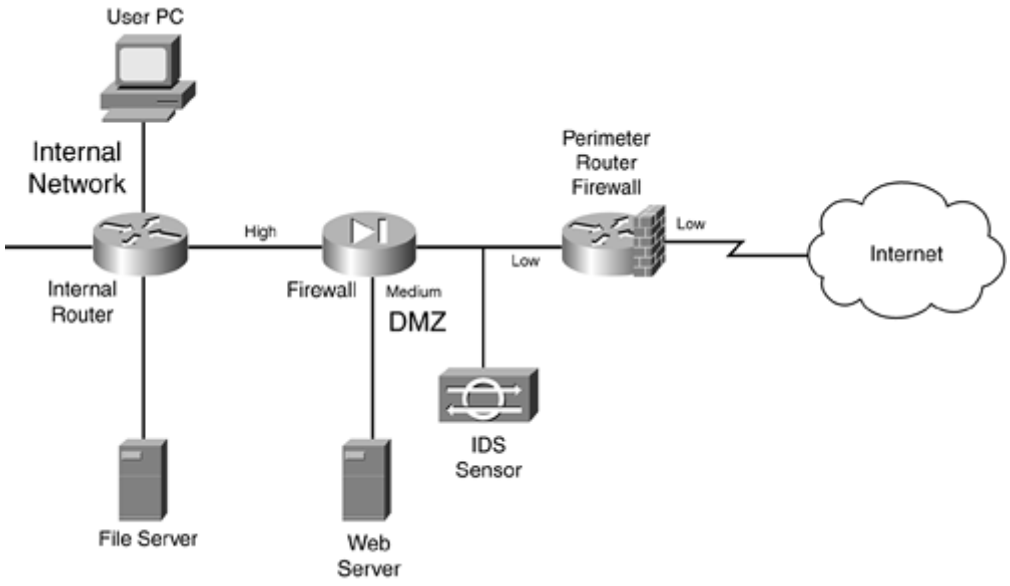

In this example, a perimeter router with basic packet filtering screens traffic as it enters the network. A standalone IDS device is used to detect attacks that the perimeter packet-filtering firewall did not filter.

The traffic then is processed by a stateful firewall. The stateful firewall has set up three security levels: low for the Internet side, medium for the DMZ, and high for the internal network. A security rule was added on the stateful firewall to allow traffic from the Internet to only the web server. All other traffic from a lower security level to a higher one is prohibited; however, higher-to-lower movement is permitted, allowing the web server administrator located on the internal network to log into the DMZ web server to update web pages.

An internal router in this design provides routing to internal segments. If you need to set up security levels and restrict access to areas of this network, you can use basic packet-filtering services on the router.

One of the advantages of this design is its simplicity: It has a minimum of three layers of defense: the packet-filtering firewall at the perimeter, the IDS, and the stateful firewall. Optionally, you can turn the internal router into a packet-filtering firewall.

This design has some disadvantages, however:

- Any attacks directed at the perimeter router/firewall are not seen by the IDS, which might be useful in determining who is trying to hack into the router and how they are trying to do it.
- No IDS exists on the inside the network, so you cannot easily determine whether internal attacks are occurring.
- The internal router might provide only simple packet filtering, which makes it difficult to implement security levels for internal users. However, this can be remedied by using a different type of firewall, such as a stateful firewall or an AGF.

Enhanced Firewall System Design

The second firewall system design is shown in Figure 2-29. As you can see, it has more components and rectifies some of the security deficiencies in the simple firewall system design. I examine the perimeter router component first. As in the last example, the perimeter router/packet-filtering firewall is performing basic filtering of traffic as it comes into the Internet. Nothing is different in this example except for what the bottom-right IDS device is doing: monitoring both the external Internet segment and the segment between the packet-filtering perimeter router and the stateful firewall. This allows the IDS to see what attacks are directed at the router, as well as what attacks are getting through the router.

Figure 2-29. Enhanced Firewall System Design

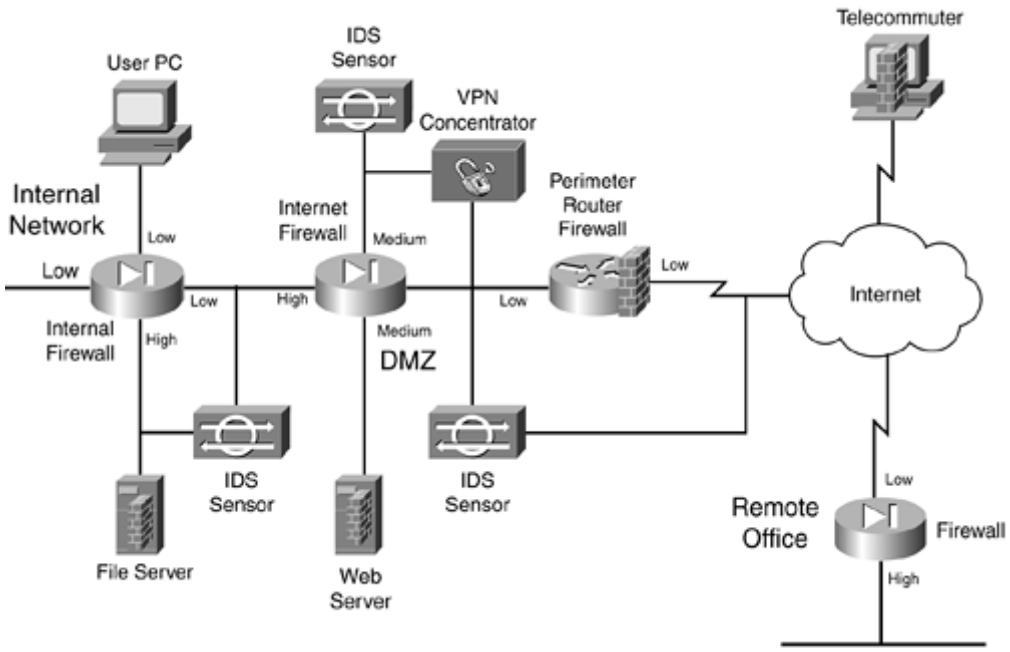

Next is the VPN concentrator. It is used to provide encrypted connections for the remote office connection (from the remote office firewall to the concentrator), as well as to terminate remote-access user connections from telecommuters and SOHO users. Notice that an IDS sensor behind the VPN concentrator is examining the unencrypted traffic. This is placed here just in case one of the remote access users or the remote office becomes compromised: The IDS can view the unencrypted traffic to detect network threats, which the IDS device connected between the perimeter router and the Internet firewall cannot because the traffic is encrypted at this point.

In addition, when the unencrypted traffic is sent to the Internet firewall from the external users, it is assigned a medium level of security, which means that it can be routed back to the Internet without any filtering and to the DMZ (assuming that same-to-same level of access is allowed). To access an internal resource, the internal firewall needs a security rule configured.

The Internet firewall provides a second layer of filtering after it passes the perimeter router/firewall. It handles traffic from the VPN concentrator, the Internet, and the DMZ. Notice that the server in the DMZ has a host-based firewall installed, adding protection to it.

The bottom-left IDS is monitoring the Internet Firewall-to-Internal Firewall segment as well as the internal high-security segment, detecting attacks that get through the respective firewalls. Plus, the host-based firewall software is installed on the internal file server.

Inside the network, an internal stateful firewall is used to provide security levels. In this example, the file-server connection has a high security level, and all of the other connections are set to low. This means that rules must be configured on the internal firewall to allow any type of traffic to reach the internal file server.

In all, this is a good security design. It uses IDS components at key places to monitor critical traffic, and it has a layered defense approach, with a packet-filtering firewall and two stateful firewalls. Plus, host-based firewall software is used on critical servers. A VPN also is used to protect traffic across the Internet.

This design also has some disadvantages, however:

- It costs a lot more than the simple design.
- The IDS systems are monitoring a lot of traffic and are generating a lot of logging information. If someone does not take the time to examine the IDS logs carefully, attacks could slip through the cracks.
- Much more configuration and management need to be done, to ensure that the correct security policies are implemented on the firewalls: the perimeter packet-filtering firewall, the two stateful firewalls, and the two host-based software firewalls. In this scenario, you definitely want to look at management software that enables you to define all your security policies from one desktop, and then have these polices converted into the configuration commands and downloaded and executed on your firewall devices.

TIP

As I learned a long time ago, there is no such thing as network nirvana. Every solution has advantages and disadvantages. The trick is to weigh these when factoring which solution meets your security policy requirements, provides the best cohesive yet manageable solution, and is the most cost-effective. Notice that I did not say that the solution has to be the most secure?only that it needs to meet your security policies.

Design Considerations

Here are some important points to remember when placing components in a firewall system:

- Use a packet-filtering firewall for the boundary router, to provide an extra level of protection. If you are really concerned about security, use a stateful firewall for this component.
- All servers that have publicly available resources should be placed in a DMZ. Servers that handle critical processes or financial transactions should have host-based firewall software installed on them. In addition, all unnecessary services on these servers should be disabled.
- For DMZ servers with sensitive information, consider using a multiple DMZ design. This is especially true if you have a web server and a database server that it interfaces with. Put the web server on a lower security level than the database server.
- If external users or remote sites that traverse a public network want to access internal resources, require a VPN. This ensures that any sensitive data is protected.
- VPN connections, as well as remote-access connections through private networks, should be terminated on their own DMZ on the Internet firewall. An IDS system should be used to examine this traffic after it is decrypted. This also allows the traffic to go right back out to the Internet, but because it is going through the firewall, you have more control over what is allowed.
- For critical internal resources, use an IDS component to monitor key segments to detect network threats. You can add extra security by segmenting your internal network into different security levels. This can compartmentalize your network and restrict access from the general population of users to areas that they have no business being in, such as accounting and R&D.
- For e-mail, you should have a public e-mail server in the DMZ that accepts all incoming and outgoing mail services. I highly recommend that you have antivirus, spam, and host-based firewall software running on this server. I also recommend that you use a limited form of a CGF that can examine mail content and make filtering decisions on it, to catch networking threats that the antivirus software cannot deal with. After e-mail is processed, it then can be forwarded to an internal server. I also recommend that all outgoing e-mail be forwarded through the DMZ e-mail server and have the same processes performed on it (remember that someone on the inside might try to use your e-mail server for malicious purposes).

I could list probably a dozen or so more items, but these are the more important ones. As you can see from this list, you have your work cut out for you.

Firewall Implementation

Now that you have chosen your components for your firewall system, you need to set them up and configure your security policies on them. Typically, you use either a command-line interface (CLI) or a graphical user interface (GUI) to perform the configuration.

Cisco products support both methods, but this book focuses on the CLI, which is the most common method used by Cisco administrators.

TIP

I have found that, at least with Cisco products, the Cisco development team for routers always implements features using the CLI first and then adds this function to GUI-based products, such as Security Device Manager (SDM) and CiscoWorks VMS. Therefore, I highly recommend that you become very comfortable with the Cisco CLI if you plan to keep your Cisco IOS routers up-to-date with the latest security technology.

Security Device Manager

Cisco has introduced a new web device manager called Security Device Manager (SDM) that provides an alternative to CLI configuration of a Cisco router. SDM is a web software component loaded into flash of a supported Cisco router that enables you to use a web browser to configure the router.

Not all routers support SDM, but most routers that Cisco currently sells today do, including the 831, 836, 837, 1701, 1710, 1711, 1712, 1721, 1751, 1760, 2610XM, 2611XM, 2620XM, 2650XM, 2651XM, 2691, 3620, 3640, 3661, 3662, 3725, 3745, 7204VXR, 7206VXR, and 7301. You also need Cisco IOS 12.2(11)T6 or later on your router.

If you buy one of the previously mentioned routers today, SDM automatically ships with it; however, you easily can download SDM from the Cisco web site and install it in flash on a supported router platform.

Figure 2-30 shows the main screen of SDM. With SDM, Cisco provides wizardlike functionality when configuring many Cisco IOS features, including LAN, WAN, firewall, and VPN features.

The use of SDM is beyond the scope of this book; however, if you are comfortable using the CLI of a router, you will find that using SDM makes configuring a Cisco router, including its security features, a very simple process.

More information about SDM can be found at http://www.cisco.com/en/US/products/sw/secursw/ps5318/index.html.

Figure 2-30. SDM's Main Window

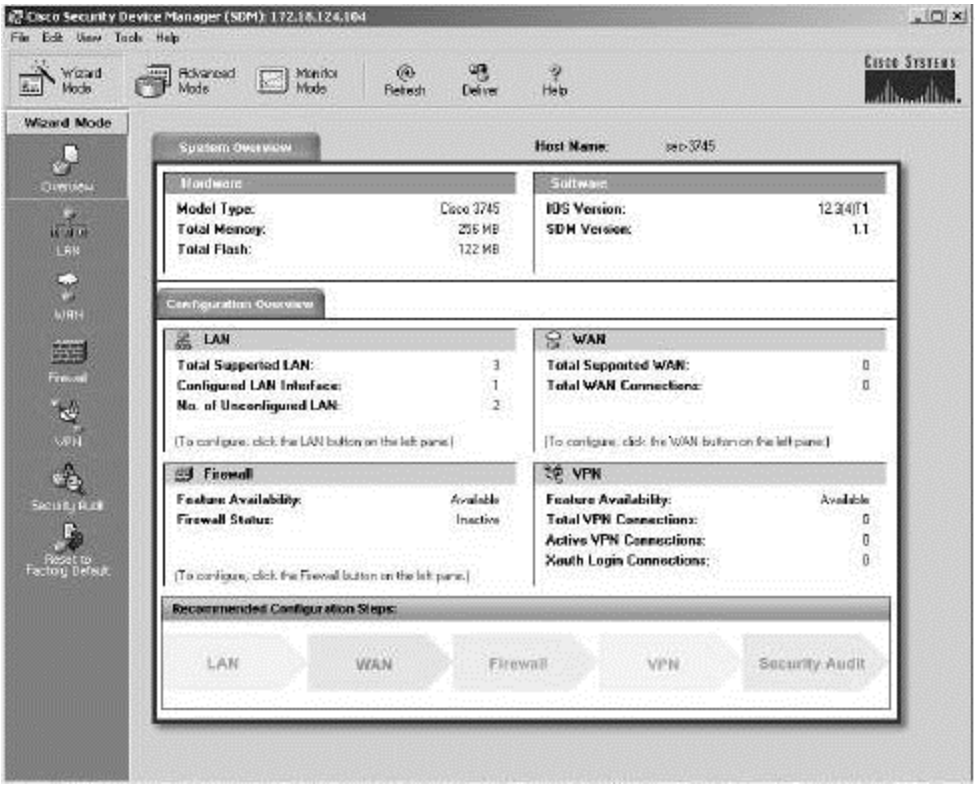

TIP

If you are not familiar with the CLI of Cisco routers, SDM is the right management tool for you. Using a GUI interface, SDM turns your input into actual router CLI commands. After you perform and apply a configuration in SDM on your router, you can use the CLI to view the actual commands that SDM implemented. This provides a training tool to help you become more familiar with the CLI.

Implementing Firewall Features

One of the first things you need to do is secure your firewall(s) itself. In other words, if a hacker breaks into your firewall, your network is wide open to a multitude of attacks. Therefore, you should be very careful about how it is administered, what services it is running, and how an administrator can access and manage it. Part II, "Managing Access to Routers," covers this information.

Most firewall devices use rule sets to set up security rules and access controls. As you will see with Cisco routers in Part III, "Nonstateful Filtering Technologies," and Part IV, "Stateful and Advanced Filtering Technologies," most of the security rules and access controls are defined by using access control lists (ACLs). As you will see, you can use many methods in the Cisco IOS to create your rule sets. These rule sets should be defined as specifically as possible: If you are permitting traffic between two machines, be specific about the type of traffic, such as TCP on port 80, or UDP on port 69. Basically, the rule-set premise should be to permit certain things and drop everything else.

Other features also should be implemented on your firewalls, if necessary: address translation (Part V, "Address Translation and Firewalls"), access authentication to your network (Part VI, "Managing Access Through Routers"), IDS (Part VII, "Detecting and Preventing Attacks"), VPNs (Part VIII, "Virtual Private Networks"), and other necessary features and services.

Firewall Administration and Management

After you have completed your firewall system implementation, you need to administer and manage it on an ongoing basis. One of the weakest links in your security setup is the people maintaining it: People are prone to making errors. Therefore, your security solution should be simple enough for your administrators to make changes to it and to troubleshoot it, yet still meet the objectives outlined in your company's security policy.

Even in this situation, configuration errors will be made in your firewall system. Therefore, before any changes are made to the firewall system, it is important that you back them up before the changes are made. After changes are made, it is of the utmost importance that you always test changes in your firewall system.

You should use two basic methods when testing the configuration. First, you should print the configuration and compare its rule set to your security policy, to double-check that the configuration it is using follows the security policy. Second, use software tools to test the change. In this situation, you, the administrator, are pretending to be the hacker. Many tools available to you enable you to perform all kinds of tricks, such as IP spoofing, DoS attacks, and others. Some of these I cover in this book as I show you how to use your Cisco IOS router to protect yourself from them. Many vendors also sell software packages geared toward security policy testing. The tools I use are either freeware or shareware.

CAUTION

Any major changes to a firewall system first should be done in a lab environment and should be tested before being configured on your production firewall system. Failure to do this can create serious security problems for your network if you make configuration mistakes.