Your Short Cut to Knowledge

The following is an excerpt from a Short Cut published by one of the Pearson Education imprints.

Short Cuts are short, concise, PDF documents designed specifically for busy technical professionals like you.

We've provided this excerpt to help you review the product before you purchase. Please note, the hyperlinks contained within this excerpt have been deactivated.

Tap into learning—NOW!

Visit **www.informit.com/shortcuts** for a complete list of Short Cuts.

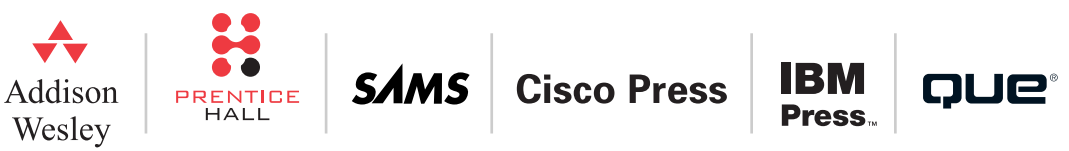

CHAPTER TWO
Typical Zone-Based Firewall Designs

NOTE

Network designers who are already familiar with the zone-based firewall design principles can skip this chapter and continue to the implementation section. All other readers are highly advised to read it, because a well-thought-out and well-documented design almost always results in superior implementation.

In the preceding chapter, you saw how the philosophy of firewall design has evolved from packet-filter-oriented designs toward zone-based designs. In this chapter, we discuss several well-known firewall designs and describe the filtering policies used in them. These designs then serve as the blueprints for the following implementation chapters, in which you'll see how you can translate a zone-based design directly into Cisco IOS configuration commands.

Simple LAN-to-Internet Firewall

The simplest possible design is a protected LAN connected to the public IP network (for example, the Internet) through a firewall. No publicly accessible servers reside in the protected LAN (which means, for example, that the users connected to the LAN have to download their e-mails from somewhere else—for example, from their Yahoo! or Hotmail accounts). Figure 2-1 shows a simple firewall design.

CHAPTER 2

Typical Zone-Based Firewall Designs

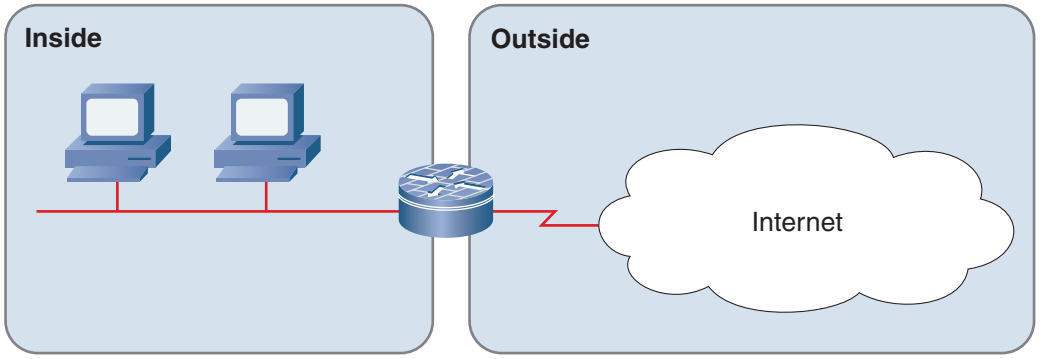

FIGURE 2-1 Simple LAN-to-Internet firewall

Two firewall policies are commonly used in such setups:

- Permissive policy (documented in Table 2-1), where the internal users can access any service on the Internet
- Restrictive policy (documented in Table 2-2), where the internal users can access only a restricted set of services (for example, web and mail services)

NOTE
All the firewall policy tables in this book assume a stateful firewall implementation, so there are no special entries for the return packets.

TABLE 2-1 Permissive Firewall Policy

Client Zone	Server Zone	Sessions Allowed
Inside	Outside	All
Outside	Inside	None

When establishing a restrictive inside-to-outside policy, consider the following caveats:

- Internet Control Message Protocol (ICMP) echo has to be enabled to allow the internal users to use ping for connectivity troubleshooting.

- Domain Name Service (DNS) on UDP port 53 has to be enabled; otherwise, the internal clients will not be able to resolve the hostnames into IP addresses.
- In some environments, additional UDP traffic has to be allowed to support the traceroute program.
- Post Office Protocol version 3 (POP3) or Internet Message Access Protocol (IMAP) as well as Simple Mail Transfer Protocol (SMTP) to some external server (for example, the Internet service provider's servers) have to be enabled to allow users to receive and send e-mail.
- Web servers sometimes use nonstandard port numbers. If your users want to access those servers, you must enable access to specific ports on specific IP addresses.

NOTE

Most versions of Windows use ICMP to trace the path to the destination host. UDP packets as the trace packets are used by some UNIX environments (for example, Linux) as well as Cisco IOS.

NOTE

TCP ports needed for FTP data sessions are identified automatically by the stateful firewall implementations. If you're not using a stateful firewall, you need to allow TCP port 20, too, and advise the clients to use passive FTP.

TABLE 2-2 Restrictive Firewall Policy

Client Zone	Server Zone	Sessions Allowed
Inside	Outside	DNS (UDP port 53) to ISP's name servers ICMP echo HTTP (TCP port 80) to everywhere HTTPS (TCP port 443) to everywhere FTP (TCP port 21) to everywhere POP3 (TCP port 110) to ISP's e-mail servers SMTP (TCP port 25) to ISP's e-mail servers
Outside	Inside	None

The simple LAN-to-Internet firewall design successfully addresses the needs of small home offices and small companies (SOHO environment) that do not own e-mail or web servers. If customers want to run their own publicly accessible servers, however, the firewall policies have to be changed to allow incoming sessions (for example, incoming SMTP and HTTP sessions). Because of the many times hackers have successfully exploited publicly accessible servers to penetrate private networks, it's almost a necessity to put such servers into a separate zone, as described in the next section.

Firewall with Public Servers

There are two common ways of designing a firewall with a perimeter network (also known as demilitarized zone [DMZ]) containing public servers:

- A simple setup in which all three zones are connected to the same networking device, as shown in Figure 2-2.
- A more complex setup with two firewalls and a transit perimeter network, as shown in Figure 2-3. This setup is slightly more secure because the two networking devices are configured independently, thus reducing the risk of a fatal configuration error (defense-in-depth principle).

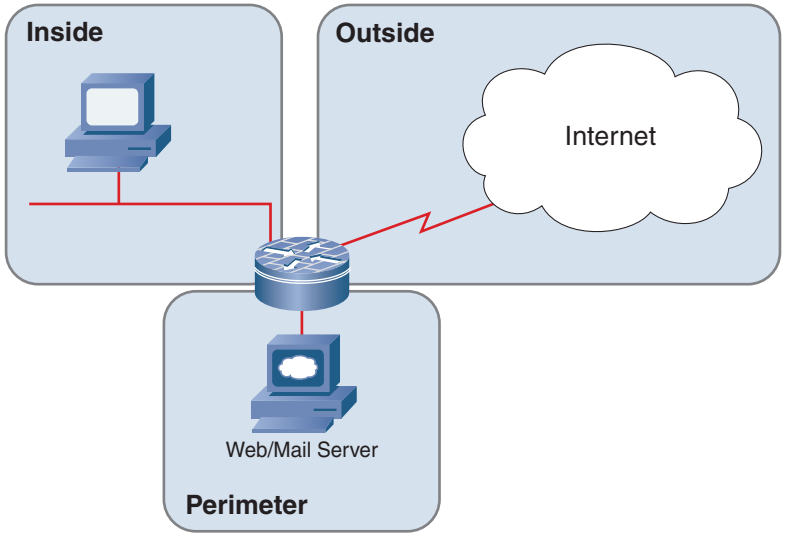

FIGURE 2-2
Simple firewall with a perimeter network

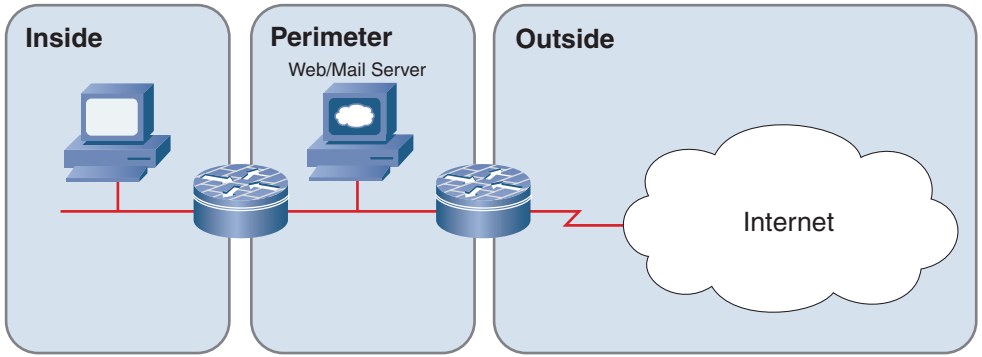

FIGURE 2-3
Transit perimeter network

As in the preceding section, a firewall with a perimeter network can implement a permissive firewall policy, where the internal users can access any service on the public network, or a more restrictive one, as documented in Table 2-3.

When designing a restrictive firewall policy, you must consider the implementation details of the DNS and e-mail services. DNS service can be implemented in the following ways:

- The customer does not run a DNS server. In this case, DNS requests from *inside* to an *outside* (ISP-owned) DNS server must be permitted.
- The customer runs a caching DNS server in the *perimeter* network but does not control its own domain. In this case, DNS requests flow from *inside* to the *perimeter* and from the *perimeter* to *outside*.
- The customer runs an authoritative DNS server for its domain in the *perimeter* network. This server usually acts as a caching DNS server for the inside clients. In this setup, the DNS requests to the perimeter DNS server arrive from inside and outside zones, and the perimeter DNS server (when acting on behalf of the inside clients) sends DNS requests to the outside. Table 2-3 describes this scenario.
- The customer runs a caching name server on the *inside* network and an authoritative name server on the *perimeter* network. DNS requests thus flow from the *inside* server to the *outside.* There is also a bidirectional flow of requests between the *perimeter* server and the *outside*.

NOTE
Please refer to Bind 9 documentation for an in-depth explanation of name server nomenclature.