Chapter 9

TCP/IP Internetworking II

LEARNING OBJECTIVES

By the end of this chapter, you should be able to:

- Explain IPv4 subnet planning and do calculations needed for working with subnet and host parts and decide on part lengths.
- Explain the purposes of Network Address Translation (NAT), how NAT operates, and problems related to NAT.
- Explain how the Domain Name System (DNS) operates.
- Describe the object model in the Simple Network Management Protocol (SNMP) and describe the enabling value of good security in the use of Set commands.
- Describe IPv6 subnetting and IP protocol stacks.
- Describe how the DNS was modified to deal with IPv6 addresses for host names.

INTRODUCTION

In Chapter 8, we looked at core TCP concepts. In this chapter, we will focus on the security and management of TCP/IP networks.

CORE TCP/IP MANAGEMENT TASKS

If a firm uses TCP/IP as its internetworking protocol, it must do a considerable amount of work to build and maintain the necessary infrastructure of TCP/IP. While switched networks are (generally) capable of operating for long periods without intervention by network managers, TCP/IP internets require constant tuning and support. This results in a need for considerable TCP/IP expertise and management effort.

IP Subnet Planning

As Chapter 8 discussed, IP addresses are 32 bits long. Each organization is assigned a network part. We saw that the University of Hawai`i's network part (128.171) is 16 bits long. A firm has no control over its network part. However, it was up to the university to decide what to do with the remaining 16 bits.

Subnetting at the University of Hawai`i The university, like most organizations, chose to subnet its IP address space. It divided the 16 bits over which it has discretion into an 8-bit subnet part and an 8-bit host part.

The $N = 2^b - 2$ Rule With b bits, you can represent 2^b possibilities. Therefore, with 8 bits, one can represent 2^8 (256) possibilities. This would suggest that the university can have 256 subnets, each with 256 hosts. However, a network, subnet, or host part cannot be all 0s or all 1s.[1] Therefore, the university can have only 254 (256 − 2) subnets, each with only 254 hosts. Figure 9-1 illustrates these calculations.

If a part is b bits long, it can represent $2^b - 2$ networks, subnets, or hosts. For example, if a subnet part is 9 bits long, there can be $2^9 - 2$, or 510, subnets. Alternatively, if a host part is 5 bits long, there can be $2^5 - 2$, or 30, hosts.

If a part is b *bits long, it can represent* $2^b - 2$ *networks, subnets, or hosts.*

Step	Description				
1	Total size of IP address (bits)	32			
2	Size of network part assigned to firm (bits)	16		8	
3	Remaining bits for firm to assign	16		24	
4	Selected subnet/host part sizes (bits)	8/8	6/10	12/12	8/16
5	Possible number of subnets ($2^b - 2$)	254 ($2^8 - 2$)	62 ($2^6 - 2$)	4,094 ($2^{12} - 2$)	254 ($2^8 - 2$)
6	Possible number of hosts per subnet ($2^b - 2$)	254 ($2^8 - 2$)	1,022 ($2^{10} - 2$)	4,094 ($2^{12} - 2$)	65,534 ($2^{16} - 2$)

FIGURE 9-1 IP Subnetting

[1] If you have all 1s in an address part, this indicates that broadcasting should be used. All 0s parts are used by computers when they do not know their own addresses. As we will see later in this chapter, most client PCs get their IP addresses from DHCP servers. All-zero addresses can only be used in the source addresses of DHCP messages sent from a host to a DHCP server.

Balancing Subnet and Host Part Sizes The larger the subnet part, the more subnets there will be. However, the larger the subnet part is made, the smaller the host part will be. This will mean fewer hosts per subnet. There is always a trade-off. More subnets mean fewer hosts, and more hosts mean fewer subnets.

The University of Hawai`i's choice of 8-bit subnet and host parts was acceptable for many years because no college needed more than 254 hosts. Its advantage is that its subnet mask (255.255.255.0) was very simple, breaking at 8-bit boundaries. This made it easy to see which hosts were on which subnets. The host at 128.171.17.5, for instance, was the fifth host on the *17th* subnet. If the subnet mask did not break at an 8-bit boundary, you would not be able to see which subnet a host is on by looking at the address in dotted decimal notation.

However, many colleges in the university now have more than 254 computers, and the limit of 254 hosts required by its subnetting decision has become a serious problem. Several colleges have now been given two subnet numbers. These colleges must connect their two subnets with routers so that hosts on the two subnets can communicate. This is expensive and awkward.

The university would have been better served had it selected a smaller subnet part, say 6 bits. As Figure 9-1 shows, this would have allowed 62 college subnets, which probably would have been sufficient. A 6-bit subnet part would give a 10-bit host part, allowing 1,022 hosts per subnet. This would be ample for several years to come.

A Critical Choice In general, it is critical for corporations to plan their IP subnetting carefully, in order to get the right balance between the sizes of their network and subnet parts.

Test Your Understanding

1. a) Why is IP subnet planning important? b) If a subnet part is X bits long, how many subnets can you have? c) If you have a subnet part of 9 bits, how many subnets can you have? (Check figure: 510 subnets) d) If you have a subnet part of 6 bits, how many subnets can you have? e) Your firm has an 8-bit network part. If you need at least 250 subnets, what must your subnet part size be? (Check figure: 8 bits) f) Continuing the last question part, how many hosts can you have per subnet? (Check figure: 65,534 hosts per subnet) g) Your firm has an 18-bit network part. If you need at least 16 subnets, what must your subnet part size be? h) Continuing the last question part, how many hosts can you have per subnet? i) Your firm has a 22-bit network part. What subnet part would you select to give at least 10 subnets? j) Continuing the last question part, how many hosts can you have per subnet?

Network Address Translation (NAT)

One issue that firms face is whether to allow people outside the corporation to learn their internal addresses. This is a security risk. If attackers know internal IP addresses, this allows them to send attack packets from the outside world. To prevent this, companies can use **network address translation (NAT)**, which presents external IP addresses that are different from internal IP addresses used within the firm.

NAT

- Sends false external IP addresses that are different from internal IP addresses
- To expand the number of IP addresses inside the firm
- For security reasons

NAT Operation (Figure 9-3)

NAT is Transparent to Internal and External Hosts

Security Reason for Using NAT

- External attackers can put sniffers outside the corporation
- Sniffers can learn IP addresses
- Attackers can send attacks to these addresses
- With NAT, attackers only learn false external IP addresses

Expanding the Number of Available IP Addresses

- Companies may receive a limited number of IP addresses from their ISPs
- There are roughly 4,000 possible ephemeral port numbers for each IP address
- So for each IP address, there can be 4,000 external connections
- If a firm is given 254 IP addresses, there can be roughly one million external connections for clients (254 x 4000)
- Even if each internal client averages several simultaneous external connections, there should not be a problem providing as many external IP connections as a firm desires

Private IP Addresses

- Can only be used inside firms
- 10.x.x.x
- 192.168.x.x (most popular)
- 172.16.x.x through 172.31.x.x

Protocol Problems with NAT

- IPsec, VoIP, etc.
- These applications must know the receiving host's true IP address
- Firewall traversal techniques for such applications must be considered carefully.

FIGURE 9-2 Network Address Translation (NAT) (Study Figure)

NAT Operation Figure 9-3 shows how NAT works. An internal client host, 192.168.5.7, sends a packet to an external server host. The source address in this packet is 192.168.5.7, of course. The source port number is 3333. As we saw in Chapter 2, this is an ephemeral port number that the source client host made up for this connection.

When the NAT firewall at the border receives the packet, it makes up a new row in its translation table. It places the internal IP address and port number in the table. It then generates a new external source IP address and external source port number. These are 60.5.9.8 and 4444, respectively.

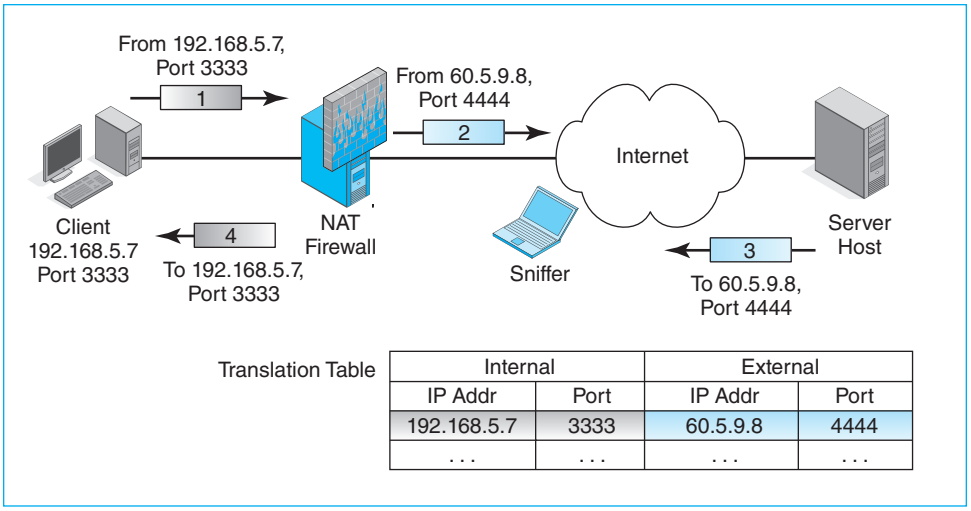

FIGURE 9-3 Network Address Translation (NAT) Operation

When packets arrive from the external host, they have 60.5.9.8 in their destination IP address fields and 4444 in their destination port number fields. The NAT firewall looks these values up in its translation table, replaces the external values with the internal values, and sends them on to the client PC.

Transparency NAT is transparent to both internal and external hosts. Hosts do not even know that NAT is happening. Consequently, there is no need to change the ways in which they operate.

NAT and Security Figure 9-3 shows how NAT brings security. An attacker may be able to install a **sniffer program** beyond the corporation's NAT firewall. This sniffer will be able to read all packets coming out of the firm. With NAT, an eavesdropper only learns false (external) IP addresses and port numbers. In theory, if an attacker can attack immediately, it can send packets to the external IP addresses and port numbers, and the NAT firewall will pass them on to the internal host. However, it is rarely possible for an attacker to act immediately, and NAT rows are only kept active for a few minutes at most. NAT provides a surprising amount of security despite its simple operation.

Expanding the Effective Number of IP Addresses An equally important potential reason for using NAT is to permit a firm to have many more internal IP addresses than its ISP gives it. Suppose that an ISP only gives a firm 254 IP addresses because it has a network part of 24 bits. In this case, the firm would not do subnetting. It would use all of the remaining 8 bits for the host part. Without NAT, the firm can only have 254 internal clients simultaneously using the Internet.

However, there are approximately 4,000 ephemeral port numbers, and therefore 4,000 possible external connections for each of the 254 public IP addresses. This gives a million external connections (4,000 times 254). NAT can map these millions of

connections into any combination of hosts and connections per host that it wishes. For example, it could have connections for 100,000 internal clients, each with 10 external connections. This shows how NAT can give a company far more internal clients than the number of external IP addresses it has.

Using Private IP Addresses To support NAT, the Internet Assigned Numbers Authority (IANA) has created three sets of **private IP address ranges** that can only be used *within* firms. These are the three ranges:

- 10.x.x.x
- 192.168.x.x
- 172.16.x.x through 172.31.x.x

The 192.168.x.x private IP address range is the most popular because it allows companies to use 255.255.0.0 and 255.255.255.0 network and subnet masks, respectively. These break at convenient 8-bit boundaries. However, the other two private IP address ranges are also widely used.

Protocol Problems with NAT In terms of security and expanding IP effective address ranges, NAT is a simple and effective tool. However, some protocols cannot work across a NAT firewall or can work only with considerable difficulty. These include the popular IPsec cryptographic system in transport mode and several voice over IP (VoIP) protocols. The problem is that these applications must know the true IP address of the other party. A number of firewall traversal techniques occur for such applications, but each application requires considerable attention.

Test Your Understanding

2. a) What is NAT? (Do not just spell it out.) b) Describe NAT operation. c) What are the two benefits of NAT? d) How does NAT enhance security? e) How does NAT allow a firm to deal with a shortage of IP addresses given to it by its ISP? f) How are private IP address ranges used? g) What are the three ranges of private IP addresses? h) What problems may firms encounter when using NAT?

The Domain Name System (DNS)

As we saw in Chapter 1, if a user types in a target host's host name, the user's PC will contact its local Domain Name System (DNS) server. The DNS server will return the IP address for the target host or will contact other DNS servers to get this information. The user's PC can then send IP packets to the target host. In this chapter, we will look at DNS and its management in more detail.

Figure 9-4 looks at how a DNS provides an IP address when a host sends a DNS request message specifying a host name. In many cases, as we saw in Chapter 1, the local DNS server will know the IP address and send it back. In other cases, the local DNS host will not know the host's IP address. It must then find the **authoritative DNS server** for the domain in the host name. In the figure, dakine.pukanui.com's authoritative DNS server is authoritative for the pukanui.com domain. This DNS server will send the IP address to the local DNS server, which will pass the address on to the host that sent the DNS request.

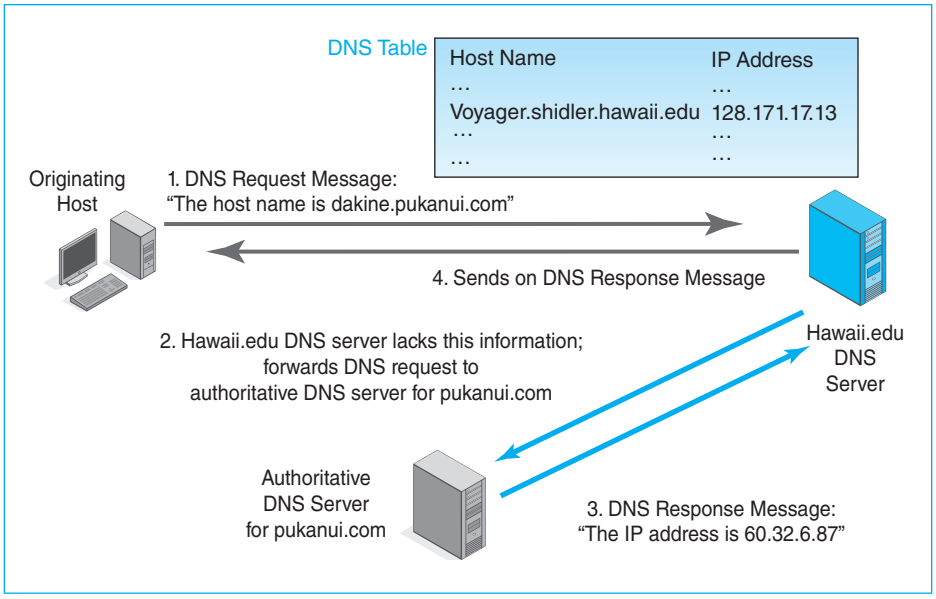

FIGURE 9-4 Domain Name System (DNS) Lookup

What is a Domain? Figure 9-5 shows that the **Domain Name System (DNS)** and its servers are not limited to providing IP addresses for host names. More generally, DNS is a general system for naming domains. A **domain** is any group of resources (routers, single networks, and hosts) under the control of an organization. The figure shows that domains are hierarchical, with host names being at the bottom of the hierarchy.

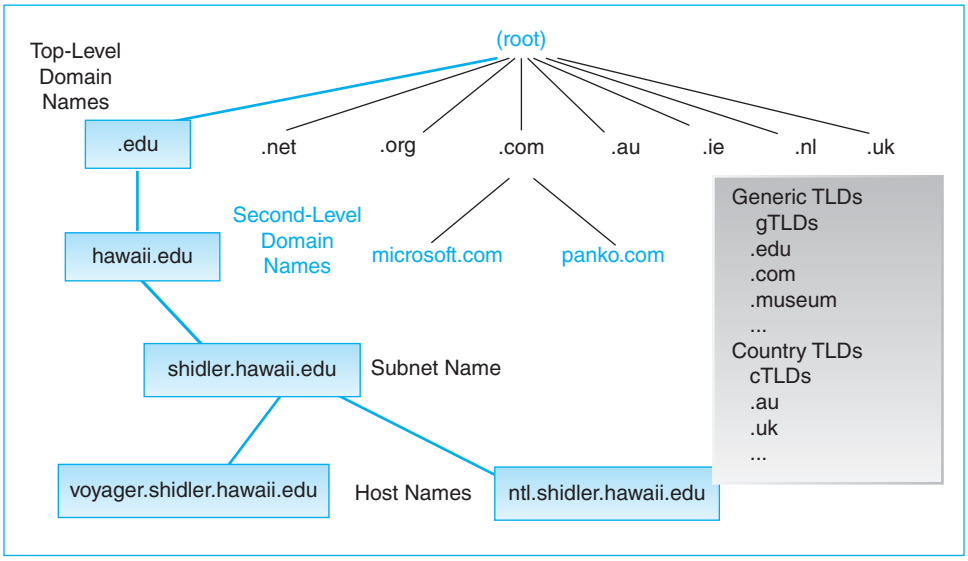

FIGURE 9-5 Domain Name System (DNS) Hierarchy

A domain is any group of resources (routers, single networks, and hosts) under the control of an organization.

Root The domain name system is a hierarchy. At the top of the DNS hierarchy is the root, which consists of all domain names. The overall control of the root, and therefore of the entire directory tree, is managed by the Internet Corporation for Assigned Names and Numbers (ICANN). Thirteen **root DNS servers** keep overview information for the system.

Top-Level Domains Under the root are **top-level domains** that categorize the domain in one of two ways.

- **Country top-level domains (cTLDs)** specify the country of the domain owner. Examples are .uk, .ca, .ie, .au, .jp, .nl, .tv, .md, and .ch.
- **Generic top-level domains (gTLDs)** specify that the organization owning the name is a particular type of organization. The first gTLDs included .com, .edu, .net, .info, .gov, and .org. Later, the IANA added several more gTLDs, such as .name and .museum. In 2012, ICANN opened the naming system widely, permitting any organization to propose new generic top-level domains.

Note the distinction between the root and top-level domains. The root consists of all domains. It is not named as a level, however. If you are familiar with the UNIX operating system, the root directory concept is similar.

Also, note that it is possible for a domain to have two top-level designations, for instance, AAAA.com.ie. Most organizations, however, tend to use either a country TLD or a generic TLD.

Second-Level Domains Under top-level domains are **second-level domains**, which usually specify a particular organization (microsoft.com, hawaii.edu, cnn.com, etc.). Sometimes, however, specific products, such as movies, get their own second-level domain names. Competition for good second-level domain names is fierce. Organizations and individuals compete fiercely to get second-level domains because this is how the public will reach them.

Organizations and individuals compete fiercely to get second-level domains because this is how the public will reach them.

Companies get second-level domain names from domain registrars for nominal fees. However, getting a second-level domain name is only the beginning. Each organization that receives a second-level domain name must have a DNS server to host its domain name information. Large organizations have their own internal DNS servers that contain information on all subnet and host names. Individuals and small businesses that use webhosting services depend on the webhosting company to provide this DNS service.

In addition, a second-level domain name does nothing for the firm until the firm buys or rents a webserver, builds a website, and pays an ISP to connect the website to the Internet.

Lower-Level Domains Domains can be further qualified. For instance, within hawaii.edu, which is the University of Hawai`i's second-level domain, there is a *shidler.hawaii.edu domain*. This is the Shidler College of Business. Within shidler.hawaii.edu is *voyager.shidler.hawaii.edu*, which is a specific host within the college.

Test Your Understanding

3. a) Is the Domain Name System only used to send back IP addresses for given host names? Explain. b) What is a domain? c) Distinguish between the DNS root and top-level domains. d) What are the two types of top-level domains? e) Which level of domain name do corporations most wish to have? f) What are DNS root servers? g) How does a company or individual obtain a second-level domain name? h) After you get a second-level domain name, what more must you do to have a working website for your company?

Simple Network Management Protocol (SNMP)

We saw the Simple Network Management Protocol (SNMP) in Chapter 4. We will now look at SNMP in more detail, focusing on the management information base (MIB) and the security implications of the Set command.

The Management Information Base (MIB) When the manager retrieves information from agents on managed devices, it stores this information in a database called the **management information base (MIB)**. As in databases in general, "MIB" refers both to the physical database and to the schema (organization) of the information in the database. We will focus on the latter.

The MIB schema is not relational. Instead, the SNMP MIB schema is organized as a hierarchy of objects. This term is a little confusing at first. An **object** is a piece of information about a managed device. The managed device itself is not an object. Figure 9-7 shows the basic schema for organizing SNMP objects.

SNMP Objects (see Figure 9-7)

- Not the managed devices themselves
- Objects are specific pieces of information about a managed device
- Information is kept in the management information base (MIB)

Set Commands

- Dangerous if used by attackers
- Many firms disable set to thwart such attacks
- However, they give up the ability to manage remote resources without travel

FIGURE 9-6 Simple Network Management Protocol (SNMP) (Study Figure)

System Objects

- System name
- System description
- System contact person
- System uptime (since last reboot)

IP Objects

- Forwarding (for routers). Yes if forwarding (routing), No if not
- Subnet mask
- Default time to live
- Traffic statistics
- Number of discards because of resource limitations
- Number of discards because could not find route
- Number of rows in routing table
- Rows discarded because of lack of memory
- Individual row data

TCP Objects

- Maximum/minimum retransmission time
- Maximum number of TCP connections allowed
- Opens/failed connections/resets
- Segments sent
- Segments retransmitted
- Errors in incoming segments
- No open port errors
- Data on individual connections (sockets, states)

UDP Objects

- Error: no application on requested port
- Traffic statistics

ICMP Objects

- Number of errors of various types

Interface Objects (One per Interface)

- Type (e.g., 69 is 100Base-FX; 71 is 802.11)
- Status: up/down/testing
- Speed
- MTU (maximum transmission unit—the maximum packet size)
- Traffic statistics: octets, unicast/broadcast/multicast packets
- Errors: discards, unknown protocols, etc.

FIGURE 9-7 SNMP MIB Hierarchical Object Model

- There is one set of objects for the system (switch, router, host, etc.) as a whole. For example, the manager may ask a router its system uptime—how long it has operated since its last reboot. If this is only a few minutes, the router may be suffering intermittent failures that cause it to crash and reboot frequently.
- There is also one set each of IP objects, TCP or UDP objects, and ICMP objects. For example, the manager can ask the agent for a router' routing object's value. If the value is "no," the router does not route packets. Rows discarded because of lack of memory is another useful object value to know. If a router is discarding more than a tiny number of packets because its memory is full, it is time to add more memory.
- A router may have multiple interfaces, and so will a switch (although switch *interfaces* are called ports). Each interface will have its own set of objects, including its interface speed and numbers of errors it has experienced. If an interface has too many errors, it may have problems that need attention.

Each SNMP Get command asks a managed device for the value of an object. By polling managed devices with Get commands wisely, the manager can maintain a database of important information about each managed device. The network visualization program can use this information to provide useful pictures of how well individual devices and sections of the network are functioning.

SNMP Set Security Get is very useful, but the SNMP *Set* command is even more powerful. The manager can use a Set command to tell an agent to change the configuration of a managed device. If a router interface seems to be malfunctioning, for example, the manager can tell the agent to set the value of an interface to "testing." The agent will then put the interface into testing mode. Set commands can also turn off interfaces to avoid using expensive transmission lines when demand is low.

By allowing administrators to manage devices remotely, the Set command can save companies a great deal of money by avoiding travel to fix problems. Unfortunately, many firms are reluctant to use Set commands because of security dangers. If Set is permitted and attackers learn how to send Set commands to managed devices, the results could be catastrophic. Fortunately, SNMP security has improved over time.

Test Your Understanding

4. a) Explain the difference between managed devices and objects. b) What type of SNMP object is the number of rows in the routing table? c) The number of segments sent? d) Speed? e) Why are firms often reluctant to use *Set* commands?

SECURING INTERNET TRANSMISSION

When the Internet was created, little thought was given to security. As Jon Postel, who edited the main Internet RFCs, explained to the first author, "It just wasn't a problem then, and we were stretched thin." Today, however, Internet security is very much a pressing issue. Companies are beginning to address the security of transmissions across the Internet (and within site networks as well) by using virtual private networks.

Virtual Private Networks

Figure 9-8 shows that corporations can cryptographically protect traffic flowing between two sites or between a site and a remote user. In Chapters 3 and 7, we saw host-to-host virtual private networks. Figure 9-8 shows a **remote-site-access VPN** connecting a remote user to a site and a **site-to-site VPN** that connects two corporate sites.

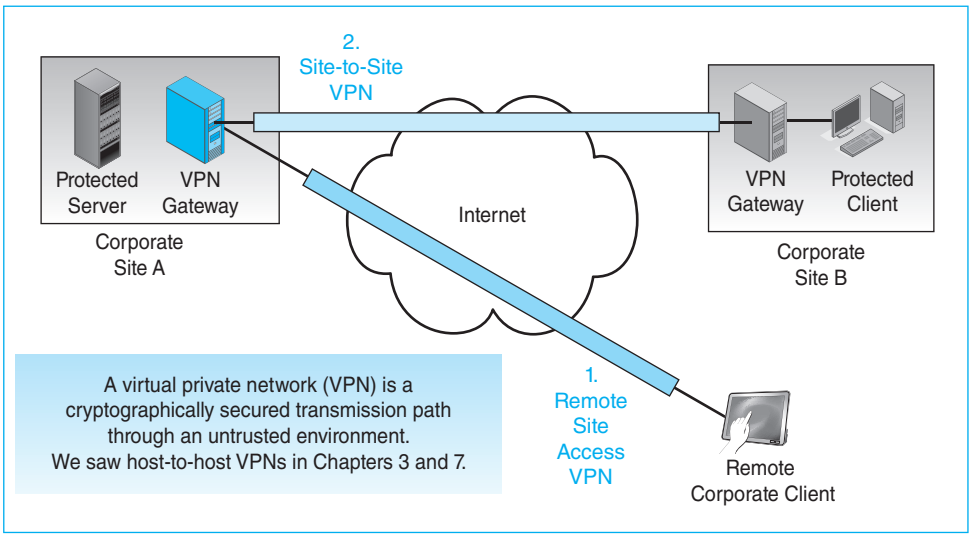

FIGURE 9-8 Remote Access and Site-to-Site Virtual Private Networks (VPNs)

A remote-site-access VPN connects a remote user to a site.

A site-to-site VPN connects two corporate sites.

- Remote-site-access VPNs are attractive because single hosts trying to connect to a corporate site via the Internet are extremely vulnerable. They are even more vulnerable if they connect to their network wirelessly.
- Site-to-site VPNs protect all traffic flowing between a pair of sites. Typically, traffic between sites is heavy, so site-to-site VPNs tend to carry much more traffic than remote-site-access VPNs.
- In Chapters 3 and 7, we also saw host-to-host VPNs, in which a client connects to a particular server using a VPN.

Both remote-site-access VPNs and site-to-site VPNs usually terminate in a **VPN gateway** at the border of each site. This VPN gateway handles cryptographic protections when dealing with remote users or a VPN gateway at another site.

IPsec VPNs

For security over the Internet, the Internet Engineering Task Force created a family of standards collectively called **IP security (IPsec)**.[3] As its name suggests, IPsec operates at the internet layer. It provides protection to at least part of the IP header. More importantly, it also provides protection to all content at the transport and application layers and to part of the IP packet header.[4] This protection is transparent, meaning that nothing has to be done to upper-layer content in order to be protected and that all upper-layer content is protected. By securing packets and their contents, the IETF provided a single mechanism to protect TCP/IP traffic.

IPsec operates at the internet layer. It provides security to all content at the transport and application layers.

IPsec offers the strongest security and should eventually dominate remote-site-access VPN transmission, site-to-site VPN transmission, and internal site IP transmission. However, IPsec is complex to manage and therefore relatively expensive to manage.

I

[3] IPsec is pronounced "eye-pea-sek," with emphasis on the sek.

[4] Actually, the term *all* is a bit too strong. In transport mode, which is discussed later, attackers can read the IP addresses because the packet is addressed to the destination host instead of to the IPsec gateway server. However, on exams, call it all.

Remote-Site-Access and Site-to-Site VPNs

IPsec is a versatile protocol that can be used for both remote-site-access VPNs and site-to-site VPNs. Coupled with its ability to protect all upper-layer content transparently, IPsec is a general solution for a firm's cryptographic protection needs.

IPsec Security Associations and Policy Servers

One advantage of IPsec as a VPN technology is that it can be centrally managed. Figure 9-10 shows that before two IPsec hosts begin to communicate, they negotiate how they will perform security. They negotiate **security associations (SAs)**, which are agreements about what security methods and options the two devices will use when they communicate.

The figure shows that the gateways implement an association in each direction. If security conditions require it, these SAs can use different security options. For example, for a remote user, remote site access might have a stronger security association from the IPsec gateway to the user than from the user to the IPsec gateway.

Some security options are very strong. Others may not be. With IPsec, companies can use central **IPsec policy servers** to manage IPsec security options in order to prevent weak options from being used. These servers specify what SA options are allowable for various gateway pairs and what options must not be used. Policy servers are especially important if a firm has many IPsec gateways.

Test Your Understanding

5. a) At what layer does IPsec operate? b) What layer content does IPsec protect? c) Does IPsec protect upper-layer protocols transparently?

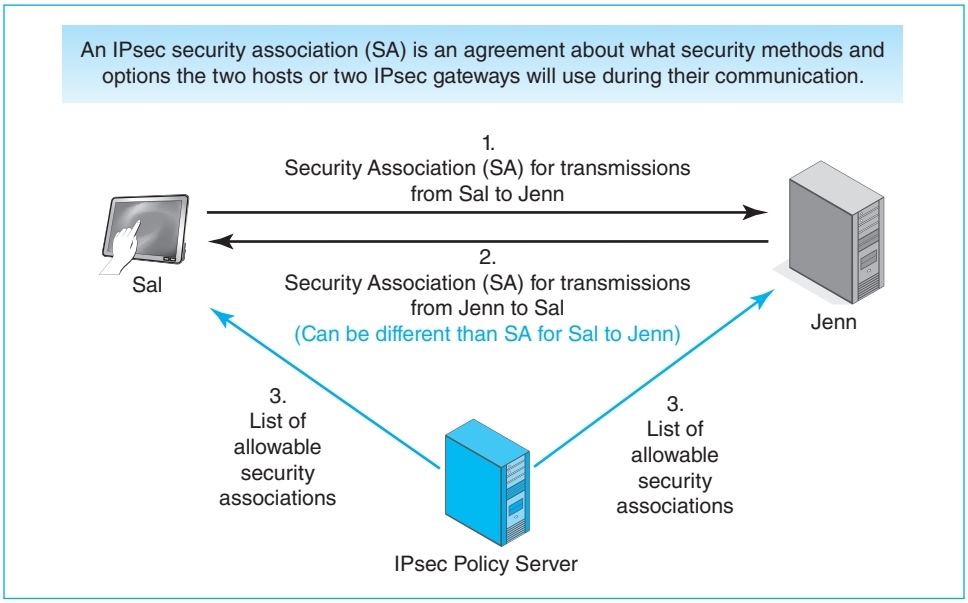

FIGURE 9-10 IPsec Security Associations (SAs) and IPsec Policy Servers

d) Is IPsec used for remote-site-access VPNs, site-to-site VPNs, or both?

6. a) In IPsec, what are security associations (SAs)? b) Must security associations be the same in the two directions? c) Describe how IPsec gateways can be managed centrally to ensure that weak SAs are not permitted.

SSL/TLS VPNs

Although IPsec is an enormously powerful tool for creating highly secure VPNs, it is expensive to implement. For many purposes, companies implement VPNs using the SSL/TLS. SSL/TLS, which we saw briefly in Chapter 3 for VPNs, is far less expensive to implement than IPsec, yet it offers reasonably strong security.

Figure 9-11 compares IPsec VPNs with SSL/TLS VPNs. It shows that the Internet Engineering Task Force is the standards agency for both. The IETF created IPsec directly. The Netscape Corporation created **Secure Sockets Layer (SSL)** standard but later passed it to the IETF. Not recognizing the concept of a "sockets layer," and noting that SSL operates at the transport layer, the IETF renamed the standard **Transport Layer Security (TLS)**. The SSL abbreviation is still widely used. Consequently, we will refer to these standards as SSL/TLS. Note that the IPsec standard operates at the internet layer.

Characteristic of VPN Technology	IPsec	SSL/TLS
Standards Organization	IETF	IETF (created by Netscape as SSL, renamed TLS by the IETF)
Layer	Layer 3	Layer 4
Built into Browsers, Webservers, and Mail Servers, So Protects These Applications at Little or No Cost	No	Yes
Can Protect any Application	Yes (also protects transport-layer header and some of the IP header)	No (only SSL/TLS-aware applications such as web and e-mail)
Type of VPNs Supported in the Standard	Host-to-Host Remote Site Access Site-to-Site	Host-to-Host
Strength of Security	Excellent	Good
Security Can Be Managed Centrally	Yes	No

FIGURE 9-11 SSL/TLS VPNs

The table emphasizes the relative strengths of IPsec and SSL/TLS by shading the box of the superior standard for selected characteristics. Note that SSL/TLS has only one "win." It is built into every browser, webserver, and e-mail program. This makes it free and simple to implement. Free and simple are powerful benefits. However, SSL/TLS can only protect a few "SSL/TLS–aware" applications and is designed only for host-to-host VPNs.

In contrast, IPsec protects everything in the transport header and application message and part of the IP header as well. It provides this protection transparently, meaning that nothing has to be changed at the transport or application layers for protection to work. In addition, IPsec supports all three types of VPNs. It is the gold standard for VPN security, and devices using IPsec can be managed centrally to ensure that they follow company policy for IPsec. Overall, SSL/TLS is a tactical VPN security tool that works for some needs, while IPsec is a strategic tool for protecting all IP and higher communication within a company.

Overall, SSL/TLS is a tactical VPN security tool that works for some needs, while IPsec is a strategic tool for protecting all IP and higher communication within a company.

You have personally used SSL/TLS. In fact, you probably used it today. If you have ever purchased something online, your URL at some point began with https://. The *s* signified that the transaction was secured with SSL/TLS. In addition, if you use a webmail e-mail service, it is likely that all communication between you and the webserver from which you get your mail is secured by SSL/TLS.

Test Your Understanding

7. a) Why are SSL/TLS VPNs attractive? b) Compare the relative advantage of SSL/TLS over IPsec. c) Compare the relative advantages of IPsec over SSL/TLS. d) How can you tell if your connection to a server uses SSL/TLS?

MANAGING IP VERSION 6 (IPV6)

In this chapter, we have looked at some aspects of managing IPv4. IPv6 generally has the same management needs. In this section, we will focus on important differences in managing IPv4 and IPv6.

Internet Layer Protocol Stacks

The internet layer sits between the transport layer and the data link layer. As Figure 9-12 illustrates, there has traditionally been a single internet layer process sitting between a host's transport layer process and its data link layer process. The internet layer process is called an **internet layer protocol stack** because it handles more than the Internet Protocol (IP). For example, it also handles the Internet Control Message Protocol (ICMP) and, in the case of IPv4, the Address Resolution Protocol (ARP). It is usually called, simply, the **IP stack**.

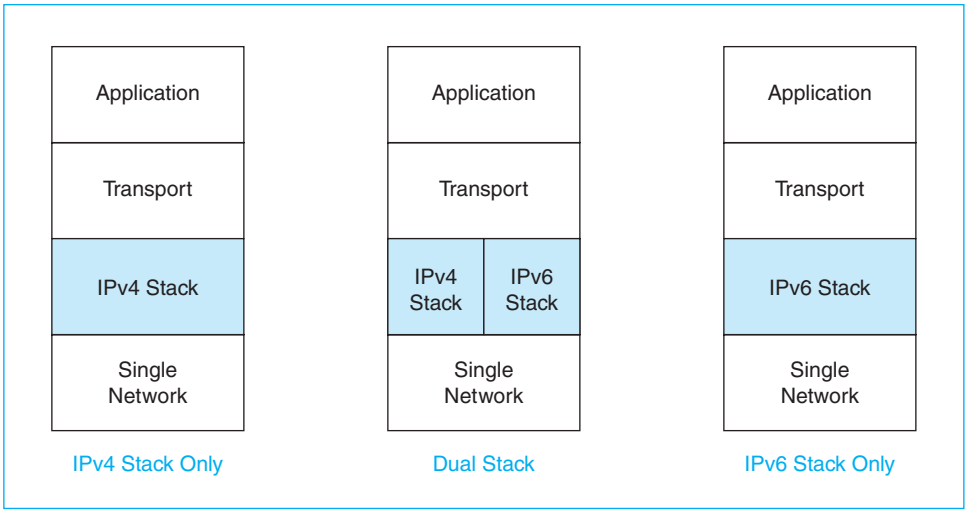

FIGURE 9-12 Internet Layer Protocol Stacks

Figure 9-12 shows that there was only a single IP stack on hosts originally—the **IPv4 stack**. The coming of IPv6 created two more alternatives.

- One is for hosts to have a single **IPv6 stack** *instead of* a single IPv4 stack.
- The other is to have a **dual-stack** host with both IPv4 and IPv6 stacks. This allows the host to transmit either IPv4 or IPv6 packets and to execute IPv4 or IPv6 internet layer protocols as needed.

The Internet Engineering Task Force expected that dual-stack implementation would become the norm. After a while, all hosts and other network devices would have both stacks. Support for IPv4 would be gradually turned off.

However, the exhaustion of IPv4 addresses has created a growing number of hosts that only have an IPv6 stack and no IPv4 stack. In addition, although IPv6 was defined in the 1990s, many vendors still sell operating systems that have the IPv6 stack turned off, weak IPv6 stacks that cause problems if turned on, and sometimes no IPv6 stack at all. If an IPv6-only client wishes to communicate with an IPv4-only server, this will be impossible. Network elements that cannot handle IPv6 may also impede communication from an IPv6-only device to another device even if it is an IPv6-only device or a dual-stack device.

Test Your Understanding

8. a) What is the advantage of having a dual stack for IP? b) Why is having only an IPv6 stack problematic?

IPv6 Subnetting

Earlier in this chapter, we looked at IPv4 subnetting. Subnetting in IPv6 is similar, but 128-bit addresses change the situation considerably. Figure 9-13 summarizes some of the key changes.

IPv6 Global Unicast Address

Like IPv4 addresses

Terminology

IPv4	IPv6	IPv4 Part Length	IPv6 Part Length
Network Part	Routing Prefix	Variable	Variable
Subnet Part	Subnet ID	Variable	Variable
Host Part	Interface IDs	Variable	64 bits
		Total: 32 bits	Total: 128 bits

Routing Prefix and Subnet ID

Subnet ID is 64 bits

Total length of routing prefix and subnet ID is therefore 64 bits

If the routing prefix is 20 bits, the subnet ID must be 44 bits long

A longer routing prefix means a smaller subnet ID and therefore fewer subnets

A shorter routing prefix means a larger subnet ID and therefore more subnets

FIGURE 9-13 IPv6 Subnetting

IPv6 Global Unicast Addresses Earlier in this chapter, we saw that IPv4 addresses have three parts—a host part, a subnet part, and a network part. Although the total length is always 32 bits, the lengths of the three parts are all variable. We also saw that these IPv4 addresses are normally public, meaning that they can be used on the Internet. In contrast, private IPv4 addresses can be used only within a corporate site or home network.

Global We will now see that **IPv6 global unicast addresses** are organized in the same way. *Global* means that packets with such addresses can be transmitted over the Internet. This is like public IP addresses in IPv4. *Unicast* means that these are addresses to use for one host to transmit to another host. The term *global unicast IPv6 address* is long, so we will call them, simply, *IPv6 addresses*.

The Three Parts Figure 9-13 also illustrates how IPv6 addresses are organized. The figure shows that IPv6 addresses, like IPv4 addresses, are divided into three parts. IPv6 does not use the terms *host part*, *subnet part*, and *network part*, but it does use similar concepts.

- The equivalent of the IPv4 network part is the *routing prefix*. The **routing prefix** lets routers on the Internet route packets to an organization.
- The equivalent of the IPv4 subnet part is the *subnet ID*. The **subnet ID** lets routers within a firm route packets to individual subnets within the firm.
- The equivalent of the IPv4 host part is the *Interface ID*. The **interface ID** identifies an individual host in the firm.

In IPv4, the size of the host part varies. In contrast, the size of the Interface ID in global unicast IPv6 addresses is fixed at 64 bits. It may seem wasteful to "use up" half of all bits in the IPv6 addresses to designate a host. However, with 64 bits left for the routing prefix and the subnet ID, there are still 1.8×10^{19} possibilities for the routing prefix and subnet ID.

The size of the Interface ID in global unicast IPv6 addresses is fixed at 64 bits.

Routing Prefix and Subnet ID Figure 9-13 indicates that the routing prefix and subnet ID are variable in length, although their total must be 64 bits because the interface ID has already consumed 64 of the 128 bits. To give an example, if the routing prefix is 20 bits, the subnet ID must be 44 bits. If an address registrar gives a firm a short routing prefix, then the company can have a large subnet ID and can therefore have many subnets. Smaller firms, needing fewer subnets, are given longer routing prefixes.

Creating the 64-bit Interface ID Returning to the 64-bit Interface ID, it would be nice to be able to use a host's data link layer address as the interface ID. However, the most common type of data link layer address is the EUI-48 address, which is only 48 bits long. Fortunately, the IEEE 802 Committee has defined a way to create a **64-bit modified extended unique identifier (EUI-64)** from a 48-bit EUI-48 address. This modified EUI-64 address can go into the interface ID field.

Creating a modified EUI-64 address from a EUI-48 address requires a series of steps, which Figure 9-14 illustrates. These steps are straightforward, and there are good technical reasons for each step. However, without a lot of in-depth knowledge, which

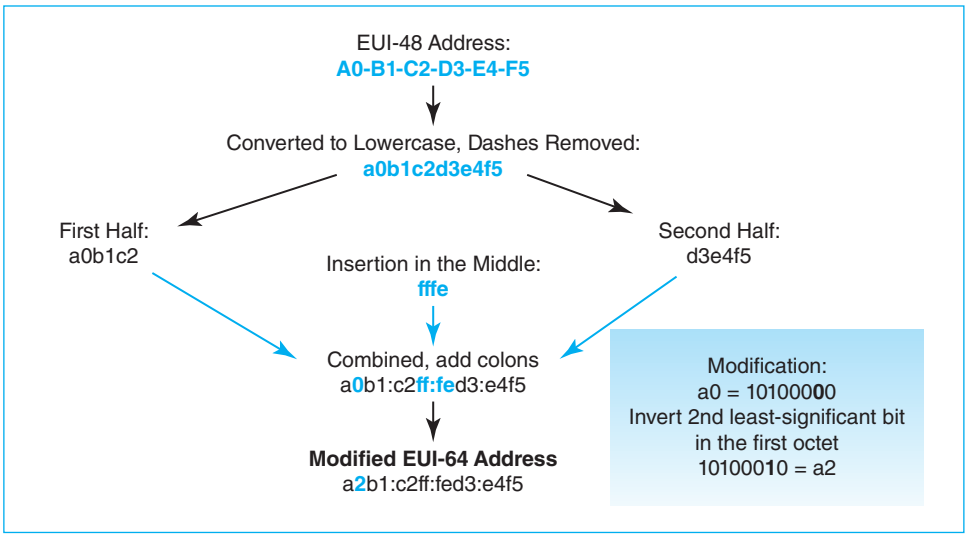

FIGURE 9-14 Converting an EUI-48 Address into an EUI-64 Address

frankly is not worth learning for information systems professionals, these five steps can appear to be illogical and weird. Just think of them as a mystical protocol for joining an obscure secret society.

- First, if the EUI-48 address is displayed in hexadecimal notation, the dashes are removed from the EUI-48 address, and the letters are changed to lowercase. So *A0-B1-C2-D3-E4-F5* becomes *a0b1c2d3e4f5*.
- Second, the 48 bits are divided into two half. Each half has 24 bits. In this case, the first half is *a0b1c2* and the second half is *d3e4f5*.
- Third, the hex symbol *fffe* is inserted between the two halves.[5] This gives a0b1c2**fffe**d3e4f5.
- Fourth, the first half, the new group, and the second half are written together and regrouped into four fields with four hex symbols apiece. Colons separate these fields. This gives the result: *a0b1:c2ff:fed3:e4f5*. Notice that a colon separates the *ff* and the *fe*. Use this as a cross-check to make sure you have done things right.
- Fifth, now we come to the *modified* part of the name. In this final step, the *second least-significant bit* (the second bit from the right end) in the first octet is inverted. For instance, the EUI-48 address in our example begins with a0. These two hex symbols constitute the first octet. In binary, they are *1010 0000*.[6] This must be changed to *1010 0010* by inverting the *second* least-significant bit—the bit that is the second from the right. Inverting a bit means changing it to 1 if it is 0 and changing it to 0 if it is 1. The inversion gives *a2* instead of *a0*. So the final modified EUI-64 is a2*b1:c2ff:fed3:e4f5*.

Test Your Understanding

9. a) What field in an IPv6 global unicast address corresponds to the network part of an IPv4 address? b) What field in an IPv6 global unicast address corresponds to the subnet part of an IPv4 address? c) If the subnet ID is 16 bits long, how long is the routing prefix? d) If you are a large company, do you want a large routing prefix or a small routing prefix?

10. a) What field in a global unicast IP address corresponds to the host part of an IPv4 address? b) How long is this field in an IPv6 global unicast address? c) Convert the following EUI-48 address to a modified EUI-64 address: AA-00-00-FF-FF-00. (Check figure: a**e**00:00ff:feff:ff00) d) Repeat for this EUI-48 address: 9B-E5-33-21-FF-0D.

The Domain Name System for IPv6

In order to make IPv6 work effectively, the Internet Engineering Task Force also had to upgrade a number of support standards. One of these was DNS. For each host name, a DNS server contains multiple records giving information about that particular host. For converting a host name to an IP address, there must be two records. One will be for the named host's IPv4 address. The other will be for the named host's IPv6 address.

[5] Don't ask why.

[6] See previous footnote.

- **DNS A Record.** The A record contains the IPv4 address for a target host. When your computer sends a DNS message to request the IPv4 address for a particular host name, the DNS server replies with information in the target host's A record.
- **DNS AAAA Record.** For IPv6 addresses, a new address field had to be added. IPv6 addresses are four times as long as IPv4 addresses, so the added record is called the AAAA record.

Test Your Understanding

11. In the Domain Name System, distinguish between the information contained in the A and AAAA records for a host name.

OTHER TCP/IP STANDARDS

In this section, we will look briefly at several other important TCP/IP standards that network administrators need to master.

Dynamic Routing Protocols

How does a router get the information in its routing table? One possibility is to enter routes manually. However, that approach does not scale to large internets. Instead, as Figure 9-15 shows, routers constantly exchange routing table information with one another using **dynamic routing protocols**.[7]

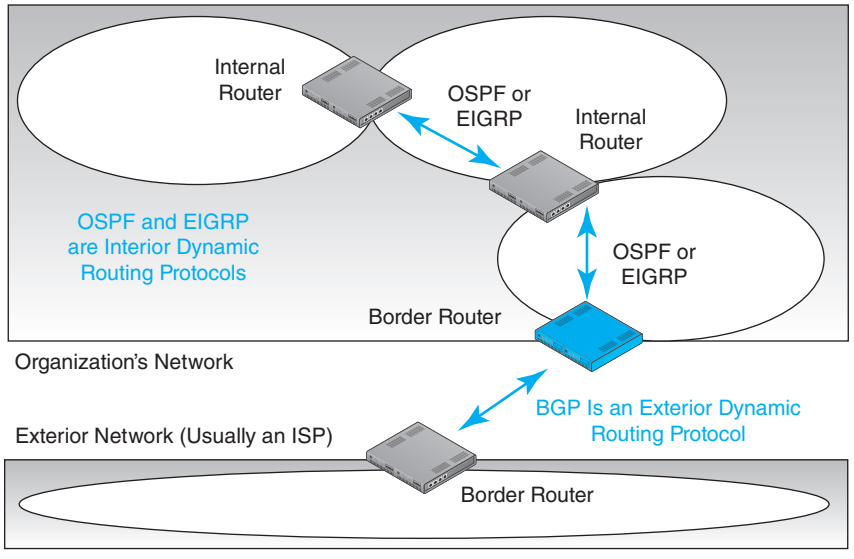

FIGURE 9-15 Dynamic Routing Protocols

[7] Note that TCP/IP uses the term *routing* in two different but related ways. First, we saw that the process of forwarding arriving packets is called routing. Second, the process of exchanging information for building routing tables is also called routing.

Interior Dynamic Protocols: OSPF and EIGRP Recall from Chapter 1 that the Internet consists of many networks owned by different organizations. Within an individual organization's network or internet, the organization decides which **interior dynamic routing protocol** to use for its internal routers, as shown in Figure 9-15. There are two[8] popular interior dynamic routing protocols.[9] Each has relative strengths and weaknesses.

- **Open Shortest Path First (OSPF).** For interior routing, the IETF created the **Open Shortest Path First (OSPF)** dynamic routing protocol. OSPF is very efficient, having a complex metric based on a mixture of cost, throughput, and traffic delays. It also offers strong security. However, it only does TCP/IP routing. Although TCP/IP is dominant today, many corporations still have legacy protocols from other standards architectures, such as IBM's SNA architecture and Novel's SPX/IPX. Corporations cannot use OSPF for routing in these other architectures.
- **EIGRP.** Cisco Systems is the dominant manufacturer of routers. Cisco has its own proprietary interior dynamic routing protocol for large internets—**Enhanced Interior Gateway Routing Protocol (EIGRP)**. The term **gateway** is another term for *router*. EIGRP's metric is very efficient because it is based on a mixture of interface bandwidth, load on the interface (0% to 100% of capacity), delay, and reliability (percentage of packets lost). EIGRP is comparable to OSPF, but unlike OSPF, it can route SNA and IPX/SPX traffic as well as TCP/IP traffic.

Exterior Dynamic Protocol: BGP For communication outside the organization's network, the organization is no longer in control. It must use the **exterior dynamic routing protocol** required by the external network to which it is connected. (This exterior network is usually an ISP.) The almost-universal exterior dynamic routing protocol is the **Border Gateway Protocol (BGP)**.

Dynamic Routing Protocol	Interior or Exterior Routing Protocol?	Remarks
OSPF (Open Shortest Path First)	Interior	For large autonomous systems that only use TCP/IP
EIGRP (Enhanced Interior Gateway Routing Protocol)	Interior	Proprietary Cisco Systems protocol. Not limited to TCP/IP routing. Also handles IPX/SPX, SNA, and so forth
BGP (Border Gateway Protocol)	Exterior	Used almost universally as the exterior routing protocol

FIGURE 9-16 Dynamic Routing Protocols (Study Figure)

[8] A third interior dynamic routing protocol of historical note is the Routing Information Protocol (RIP). RIP is very simple, making it attractive economically. However, its poor security excludes it from organizations today.

[9] A third interior dynamic routing protocol is RIP, the Routing Information Protocol. RIP is simpler than OSPF or EIGRP and was once popular. However, its almost complete lack of security features make it an unacceptable choice today.

Test Your Understanding

12. a) What is the purpose of dynamic routing protocols? b) For its own network, can an organization choose its interior dynamic routing protocol? c) What is the IETF interior dynamic routing protocol? d) When might you use EIGRP as your interior dynamic routing protocol? e) May a company select the routing protocol its border router uses to communicate with the outside world? f) What is the almost-universal exterior dynamic routing protocol?

Internet Control Message Protocol (ICMP) for Supervisory Messages at the Internet Layer

Supervisory Messages at the Internet Layer IP is only concerned with packet delivery. For supervisory messages at the internet layer, the IETF created the **Internet Control Message Protocol (ICMP)**. IP and ICMP work closely together. As Figure 9-17 shows, IP encapsulates ICMP messages in the IP data field, delivering them to their target host or router. There are no higher-layer headers or messages.

Error Advisement IP is an unreliable protocol. It offers no error correction. If the router or the destination host finds an error, it discards the packet. Although there is no retransmission, the router or host that finds the error may send an ICMP error message to the source device to inform it that an error has occurred, as in Figure 9-17. The ICMP error message contains type and code values indicating what the problem is. For example, a host unreachable message is Type 3/Code 1.

Note that this is error advisement (notification) rather than error correction. There is no mechanism within IP or ICMP for the retransmission of lost or damaged packets. ICMP error messages are only sent to help the sending process or its human user diagnose problems. They do not make IP reliable.

One important subtlety is that sending error advisement messages is not mandatory. For security reasons, many firms do not allow error advisement messages to leave their

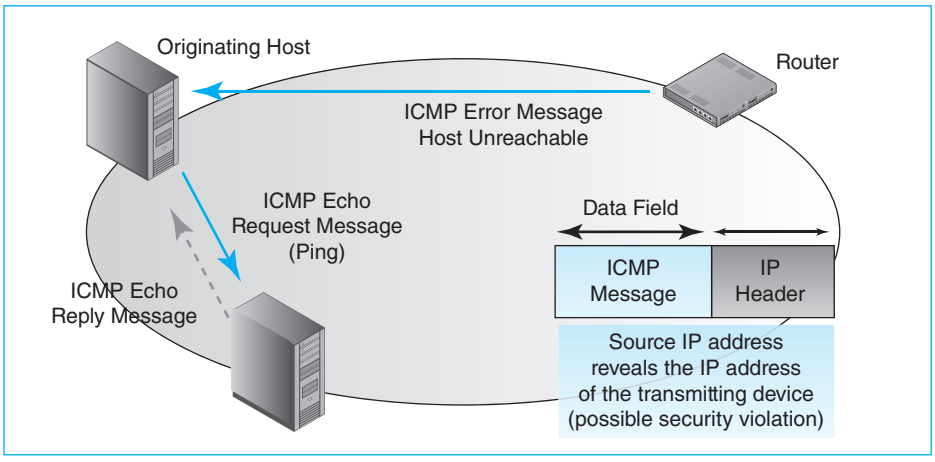

FIGURE 9-17 Internet Control Message Protocol (ICMP) for Supervisory Messages at the Internet Layer

internal internets because hackers can exploit the information contained in them. Most obviously, the ICMP message will be carried in a packet that contains the IP address of the sending router or other device. If adversaries have an exploit to use against routers, they have a target IP address for their attacks.

Echo (Ping) Perhaps the most useful ICMP messages are the echo request and response messages. As we saw in Chapters 1 and 4, one host can use these message to "ping" another host. As in the case of error response messages, the IP header for the echo response message reveals the presence of a potential target at the source IP address.

Test Your Understanding

13. a) For what general class of messages and at what layer is ICMP used? b) Explain error advisement in ICMP. c) What two ICMP message types are used in *ping*? d) What security concern do ICMP error advisement messages and echo response messages create?

CONCLUSION

Synopsis

Chapter 8 dealt with TCP/IP concepts. This chapter focuses on TCP/IP management. The TCP/IP standards that dominate internetworking require a great deal of management attention, both initially and on an ongoing basis. The first step is to develop an IP subnet schema for the firm. This creates a basic trade-off between the number of subnets and the number of hosts per subnet. The firm also has to decide whether or not to use network address translation (NAT). NAT has several benefits, including adding security and increasing the effective number of public IP addresses a firm has; but NAT causes problems for certain protocols.

In this chapter, we looked more closely at the Domain Name System (DNS). We saw that DNS is a hierarchical system of named domains (collections of resources under the control of an organization). Corporations want second-level domain names, such as pearsonhighered.com. After they get one, they must maintain DNS servers for their second-level domain. We also saw that if a local DNS server does not know the IP address for a host name, it contacts the authoritative DNS server for the second-level domain in the IP address.

This chapter looked at SNMP operation in more detail, focusing on the concept of objects and the types of objects specified in MIB schemas. We also looked at security concerns regarding the Set command.

Companies are concerned about security on the Internet. To obtain better security, they use remote-site-access and site-to-site VPNs. There are two main VPN protocols. IPsec offers the strongest security. Most importantly, IPsec offers central manageability.

SSL/TLS can be used for remote-site-access VPNs. SSL/TLS is attractive because all browsers know how to create a secure SSL/TLS connection with host computers. This means that there is no need to add anything to the client computer. However, there are limitations on the services that SSL/TLS can provide easily.

We ended the chapter with a long discussion on the management of IPv6. Managing IPv6 is similar to managing IPv4, but it is not the same. We began with a discussion of single-stack IPv4-only hosts, single-stack IPv6-only hosts, and dual-stack IPv4/IPv6 hosts. Originally, the IETF believed that dual-stack hosts would make the introduction of IPv6 painless. However, there is a growing number of IPv6-only devices, which may have trouble communicating with the large number of existing IPv4-only hosts.

In IPv4, addresses have three parts: network, subnet, and host. In IPv6 global unicast addresses, the comparable parts are called the routing prefix, the subnet ID, and the interface ID. The interface ID is always 64 bits long. It is created from the host's EUI-48 address by a somewhat complex process that we reviewed. This process results in a 64-bit modified extended unique identifier (EUI-64) that is used as the interface ID. The company is assigned a routing prefix. This sets the size of its subnet ID and determines how many subnets it can have.

IPv6 required the extension of a number of existing TCP/IP standards. The major change to DNS is the introduction of an AAAA record for host names. This record contains the IPv6 address of the named host. (IPv4 addresses are contained in traditional A records.)

Routers build their routing tables by communicating with other routers. Routers frequently exchange messages, giving information stored in their routing tables. These messages are governed by dynamic routing protocols.

IP itself does not have supervisory messages. For internet layer supervisory messages, hosts and routers use the Internet Control Message Protocol (ICMP). We looked at two types of ICMP messages—error advisement messages and echo messages (ping). ICMP messages are carried in the data fields of IP packets.

END-OF-CHAPTER QUESTIONS

Thought Questions

9-1. Assume that an average SNMP response message is 100 bytes long. Assume that a manager sends 400 SNMP *Get* commands each second. a) What percentage of a 100 Mbps LAN link's capacity would the resulting response traffic represent? b) What percentage of a 1 Mbps WAN link would the response messages represent? c) What are the management implications of your answers?

9-2. A firm is assigned the network part 128.171. It selects an 8-bit subnet part. a) Write the bits for the four octets of the IP address of the first host on the first subnet. b) Convert this answer into dotted decimal notation. (If you have forgotten how to do this, it was covered in Chapter 1.) c) Write the bits for the second host on the third subnet. (In binary, 2 is 10, while 3 is 11.) d) Convert this into dotted decimal notation. e) Write the bits for the last host on the third subnet. f) Convert this answer into dotted decimal notation.

9-3. A firm is assigned the network part 128.171. It selects a 10-bit subnet part. a) Draw the bits for the four octets of the IP address of the first host on the first subnet. (Hint: Use Windows Calculator.) b) Convert this answer into dotted decimal notation. c) Draw the bits for the second host on the third subnet. (In binary, 2 is 10, while 3 is 11.) d) Convert this into dotted decimal notation. e) Draw the bits for the last host on the third subnet. f) Convert this answer into dotted decimal notation.

Troubleshooting Question

9-4. In your browser, you enter the URL of a website you use daily. After an unusually long delay, you receive a DNS error message that the host does not exist. a) List the five troubleshooting steps discussed in Chapter 1. b) Apply them to this situation.

Hands-On Project

9-5. After Sal Aurigemma received his PhD from the University of Hawaii, he became an assistant professor at the University of Tulsa. There, he introduced the school to Aloha Friday, when people come to work in their colorful Aloha shirts. He got the idea of creating Aloha shirts with Tulsa's school colors and an emblem of the university on the shirt pocket. Suppose that he wants to create a company to sell school-specific Aloha shirts. He will need a company name and a second-level domain name. Got to an Internet domain name registrar. Thoughtfully come up with three appropriate and available domain names. Explain why each is good. Select one and explain why it is best.

Perspective Questions

9-6. What was the most surprising thing to you about the material in this chapter?

9-7. What was the most difficult thing for you in the chapter?