CHAPTER 6

VLSM and CIDR

Objectives

Upon completion of this chapter, you should be able to answer the following questions:

- What are the differences between classful and classless IP addressing?
- What is VLSM, and what are the benefits of classless IP addressing?
- What is the role of the classless interdomain routing (CIDR) standard in making efficient use of scarce IPv4 addresses?

Key Terms

This chapter uses the following key terms. You can find the definitions in the Glossary at the end of the book.

classful IP addressing *page 280*

discontiguous address assignment *page 280*

supernet *page 280*

private addressing *page 281*

high-order bits *page 282*

prefix aggregation *page 285*

network prefix *page 286*

contiguous *page 294*

supernetting *page 294*

Prior to 1981, IP addresses used only the first 8 bits to specify the network portion of the address, limiting the Internet—then known as ARPANET—to 256 networks. Early on, it became obvious that this was not going to be enough address space.

In 1981, RFC 791 modified the IPv4 32-bit address to allow three different classes or sizes of the networks:

- Class A addresses, which used 8 bits for the network portion of the address
- Class B addresses, which used 16 bits for the network portion of the address
- Class C addresses, which used 24 bits for the network portion of the address

This format became known as ***classful IP addressing***.

The initial development of classful addressing solved the 256-network limit problem, for a time. A decade later, it became clear that the IP address space was depleting rapidly. In response, the Internet Engineering Task Force (IETF) introduced classless interdomain routing (CIDR), which used variable-length subnet masking (VLSM) to help conserve address space.

With the introduction of CIDR and VLSM, Internet service providers (ISP) could now assign one part of a classful network to one customer and a different part to another customer. This ***discontiguous address assignment*** by ISPs was paralleled by the development of classless routing protocols. To compare: Classful routing protocols always summarize on the classful boundary and do not include the subnet mask in routing updates. Classless routing protocols do include the subnet mask in routing updates and are not required to perform summarization. The classless routing protocols discussed in this course are Routing Information Protocol version 2 (RIPv2), Enhanced Interior Gateway Routing Protocol (EIGRP), and Open Shortest Path First (OSPF).

With the introduction of VLSM and CIDR, network administrators had to use additional subnetting skills. VLSM is simply subnetting a subnet. Subnets can be further subnetted in multiple levels, as you will learn in this chapter. In addition to subnetting, it became possible to summarize a large collection of classful networks into an aggregate route, or ***supernet***. In this chapter, you will also review route summarization skills.

Classful and Classless Addressing

One of the ways to characterize routing protocols is either as classful or classless. This is a result of the evolution from classful to classless IPv4 addressing. As networks began to use classless addressing, classless routing protocols had to be modified or developed to include the subnet mask in the routing update. The following sections are a review of classful and classless addressing, along with an introduction to classless routing protocols.

Classful IP Addressing

Figure 6-1 shows the exponential growth of hosts in the Internet from 1992 to 2006.

Figure 6-1 Growth of the Internet from 1992 to 2006 (Source: http://en.wikipedia.org/wiki/Image:Number_of_internet_hosts.svg)

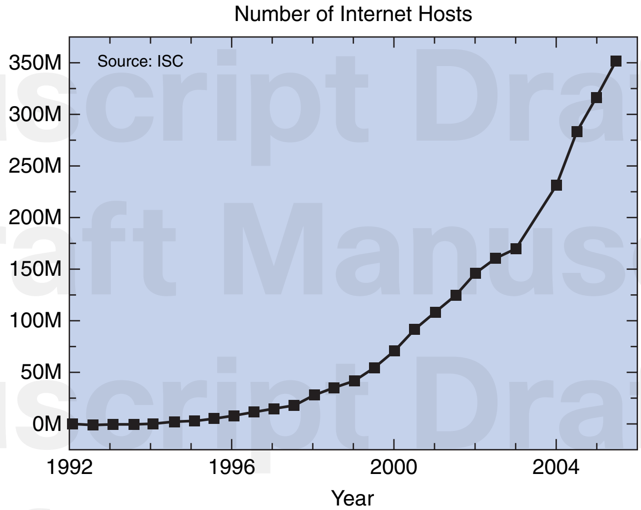

Note

The Internet Software Consortium (ISC) tracks the number of hosts on the Internet. To see more information about ISC's host count, visit "ISC Domain Survey: Number of Internet Hosts," http://www.isc.org/index.pl?/ops/ds/host-count-history.php.

When the ARPANET was commissioned in 1969, no one anticipated that the Internet would explode out of the humble beginnings of this research project. By 1989, ARPANET had been transformed into what we now call the Internet. Over the next decade, the number of hosts on the Internet grew exponentially, from 159,000 in October 1989 to over 72 million by the end of the millennium. As of January 2007, there were over 433 million hosts on the Internet.

Without the introduction of VLSM and CIDR notation in 1993 (RFC 1519), Network Address Translation (NAT) in 1994 (RFC 1631), and ***private addressing*** in 1996 (RFC 1918), the IPv4 32-bit address space would now be exhausted.

High-Order Bits

IPv4 addresses were initially allocated based on class, as shown in Table 6-1.

Table 6-1 High-Order Bits

Class	High-Order Bits	Start	End
Class A	0	0.0.0.0	127.255.255.255
Class B	10	128.0.0.0	191.255.255.255
Class C	110	192.0.0.0	223.255.255.255
Multicast	1110	224.0.0.0	239.255.255.255
Experimental	1111	240.0.0.0	255.255.255.255

In the original specification of IPv4 (RFC 791, http://www.ietf.org/rfc/rfc791.txt) released in 1981, the authors established the classes to provide three different sizes of networks for large, medium, and small organizations. As a result, Class A, B, and C addresses were defined with a specific format for the ***high-order bits***. High-order bits are the leftmost bits in a 32-bit address.

Table 6-1 shows the following details:

- Class A addresses begin with a 0 bit. Therefore, all addresses from 0.0.0.0 to 127.255.255.255 belong to Class A. The 0.0.0.0 address is reserved for default routing, and the 127.0.0.0 address is reserved for loopback testing.
- Class B addresses begin with a 1 bit and a 0 bit. Therefore, all addresses from 128.0.0.0 to 191.255.255.255 belong to Class B.
- Class C addresses begin with two 1 bits and a 0 bit. Class C addresses range from 192.0.0.0 to 223.255.255.255.

The remaining addresses were reserved for multicasting and future uses. Multicast addresses begin with three 1s and a 0 bit. Multicast addresses are used to identify a group of hosts that are part of a multicast group. This helps reduce the amount of packet processing that is done by hosts, particularly on broadcast media. In this course, you will see that the routing protocols RIPv2, EIGRP, and OSPF use designated multicast addresses.

IP addresses that begin with four 1 bits were reserved for future use.

Note

To see the assignment of individual multicast addresses, visit "Internet Multicast Addresses," http://www.iana.org/assignments/multicast-addresses.

IPv4 Classful Addressing Structure

The designations of network bits and host bits were established in RFC 790 (released with RFC 791). Figure 6-2 shows how the subnet mask for a network is determined based on its class.

Figure 6-2 Subnet Mask Based on Class

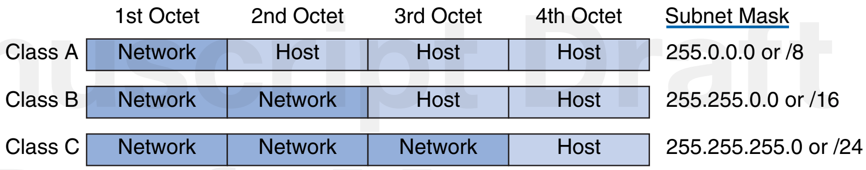

Table 6-2 shows the number of networks available per class as well as the number of hosts per network.

Table 6-2 Number of Networks and Hosts per Network for Each Class

Address Class	First Octet Range	Number of Possible Networks	Number of Hosts per Network
Class A	0 to 127	128 (2 are reserved)	16,777,214
Class B	128 to 191	16,344	65,534
Class C	192 to 223	2,097,152	254

Class A networks used the first octet for network assignment, which translated to a 255.0.0.0 classful subnet mask. Because only 7 bits were left in the first octet (remember, the first bit is always 0), this made 2^7, or 128, networks.

With 24 bits in the host portion, each Class A address had the potential for over 16 million individual host addresses. Before CIDR and VLSM, organizations were assigned an entire classful network address. What was one organization going to do with 16 million addresses? Now you can understand the tremendous waste of address space that occurred in the beginning days of the Internet, when companies received Class A addresses. Some companies and governmental organizations still have Class A addresses. For example, General Electric owns 3.0.0.0/8, Apple Computer owns 17.0.0.0/8, and the U.S. Postal Service owns 56.0.0.0/8. (The note at the end of this section provides a link to a listing of all the IANA assignments.)

Class B was not much better. RFC 790 specified the first two octets as the network portion of the address. With the first 2 bits already established as 1 and 0, 14 bits remained in the first two octets for assigning networks, which resulted in 16,384 Class B network addresses. Because each Class B network address contained 16 bits in the host portion, it controlled 65,534 addresses. (Remember, two addresses were reserved for the network and broadcast addresses.) Only the largest organizations and governments could ever hope to use all 65,000 addresses. Like Class A, Class B address space was wasted.

To make things worse, Class C addresses were often too small! RFC 790 specified the first three octets as the network portion of the address. With the first 3 bits established as 1, 1, and 0, 21 bits remained for assigning networks for over 2 million Class C networks. But, each Class C network only had 8 bits in the host portion, or 254 possible host addresses.

Note

The following links provide some background to the discussion of the IPv4 classful structure and the depletion of the IPv4 address space:

- "A Brief History of the Internet," http://www.isoc.org/internet/history/brief.shtml
- "Internet Protocol v4 Address Space," http://www.iana.org/assignments/ipv4-address-space

Classful Routing Protocol

Now that we have reviewed classful addressing, we take another look at classful routing protocol updates. Remember, a classful routing protocol does not include the subnet mask in the routing update.

Using classful IP addresses meant that the subnet mask of a network address could be determined by the value of the first octet, or more accurately, the first 3 bits of the address. Routing protocols, such as RIPv1, only needed to propagate the network address of known routes and did not need to include the subnet mask in the routing update. This is because the router receiving the routing update could determine the subnet mask simply by examining the value of the first octet in the network address, or by applying its ingress interface mask for subnetted routes. The subnet mask was directly related to the network address.

Figure 6-3 illustrates how classful routing determines the subnet mask for a given network.

In the figure, R1 knows that subnet 172.16.1.0 belongs to the same major classful network as the outgoing interface. Therefore, it sends a RIP update to R2 containing subnet 172.16.1.0. When R2 receives the update, it applies the receiving interface subnet mask (/24) to the update because the update belongs to the same major network as the interface. R2 adds 172.16.1.0 to the routing table with the /24 mask.

However, when sending updates to R3, R2 summarizes subnets 172.16.1.0/24, 172.16.2.0/24, and 172.16.3.0/24 into the major classful network 172.16.0.0. Because R3 receives this update on a major network address other than 172.16.0.0, it will apply the classful mask for a Class B network, /16.

Figure 6-3 Classful Routing Updates

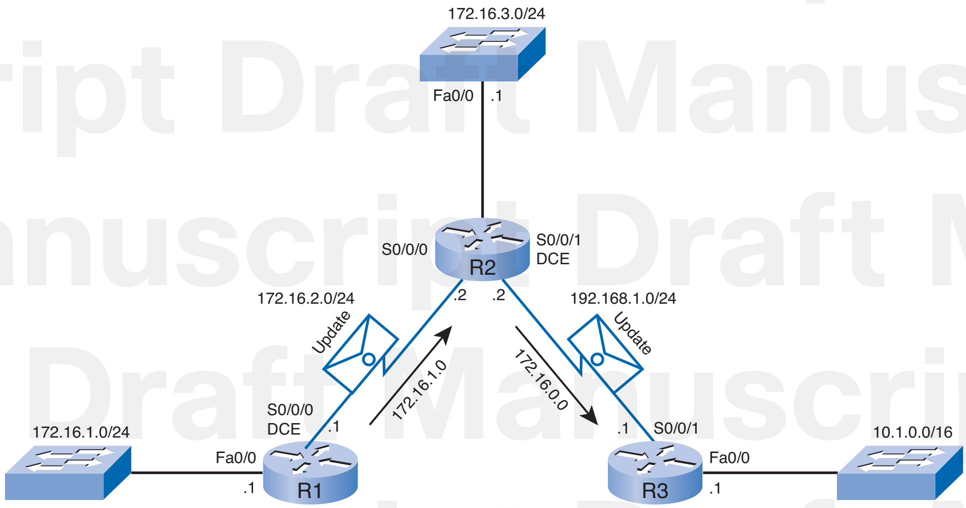

Classless IP Addressing

The previous sections discussed classful addressing and classful routing protocols. The following sections now take a look at the evolution of classless addressing and classless routing protocols.

Moving Toward Classless Addressing

By 1992, members of the IETF had serious concerns about the exponential growth of the Internet and the limited scalability of Internet routing tables. They were also concerned with the eventual exhaustion of 32-bit IPv4 address space. The depletion of the Class B address space was occurring so fast that within two years there would be no more Class B addresses available (RFC 1519). This depletion was occurring because every organization that requested and obtained approval for IP address space received an entire classful network address—either a Class B with 65,534 host addresses or a Class C with 254 host addresses. One fundamental cause of this problem was the lack of flexibility. No class existed to serve a midsized organization that needed thousands of IP addresses but not 65,000.

In 1993, IETF introduced classless interdomain routing (CIDR) (RFC 1517). CIDR allowed the following:

- More efficient use of IPv4 address space
- ***Prefix aggregation***, which reduced the size of routing tables

To CIDR-compliant routers, address class is meaningless. The network portion of the address is determined by the network subnet mask, also known as the ***network prefix***, or prefix length (/8, /19, and so on). The network address is no longer determined by the class of the address.

ISPs could now more efficiently allocate address space using any prefix length, starting with /8 and larger (/8, /9, /10, and so on). ISPs were no longer limited to a /8, /16, or /24 subnet mask. Blocks of IP addresses could be assigned to a network based on the requirements of the customer, ranging from a few hosts to hundreds or thousands of hosts.

CIDR and Route Summarization

CIDR uses variable-length subnet masking (VLSM) to allocate IP addresses to subnets according to individual need rather than by class. This type of allocation allows the network/host boundary to occur at any bit in the address. Networks can be further divided or subnetted into smaller and smaller subnets.

Just as the Internet was growing at an exponential rate in the early 1990s, so were the size of routing tables that were maintained by Internet routers under classful IP addressing. CIDR allowed prefix aggregation, which you already know as route summarization. Recall from Chapter 2, "Static Routing," that you can create one static route for multiple networks. Internet routing tables were now able to benefit from the same type of aggregation of routes. The capability for routes to be summarized as a single route helped reduce the size of Internet routing tables. Figure 6-4 illustrates a simple example of this capability.

Figure 6-4 Example of CIDR Used to Summarize Routes

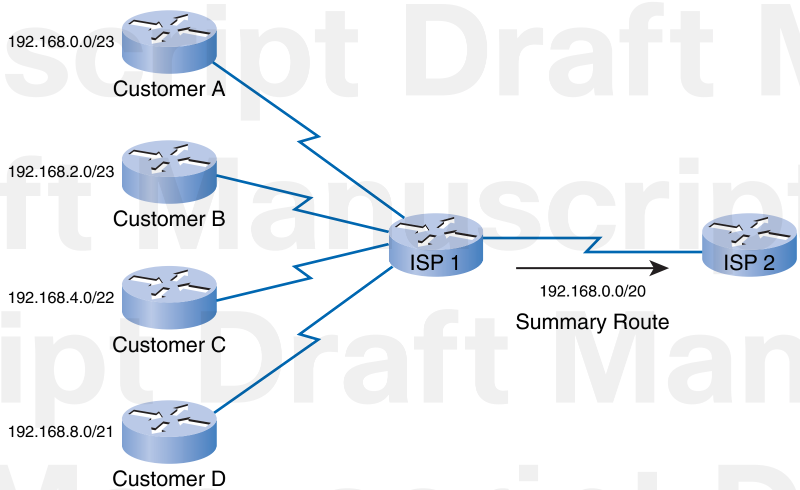

In the figure, notice that ISP1 has four customers, each with a variable amount of IP address space. However, all the customer address space can be summarized into one advertisement to

ISP2. The 192.168.0.0/20 summarized or aggregated route includes all the networks belonging to customers A, B, C, and D. This type of route is known as a supernet route. A supernet summarizes multiple network addresses with a mask less than the classful mask.

Propagating VLSM and supernet routes requires a classless routing protocol, because the subnet mask can no longer be determined by the value of the first octet. The subnet mask now needs to be included with the network address. Classless routing protocols include the subnet mask with the network address in the routing update.

Classless Routing Protocol

Classless routing protocols include RIPv2, EIGRP, OSPF, Intermediate System–to–Intermediate System (IS-IS), and Border Gateway Protocol (BGP). These routing protocols include the subnet mask with the network address in their routing updates. Classless routing protocols are necessary when the mask cannot be assumed or determined by the value of the first octet.

For example, in Figure 6-5, the networks 172.16.0.0/16, 172.17.0.0/16, 172.18.0.0/16, and 172.19.0.0/16 can be summarized as 172.16.0.0/14, known as a *supernet*.

Figure 6-5 Classless Routing

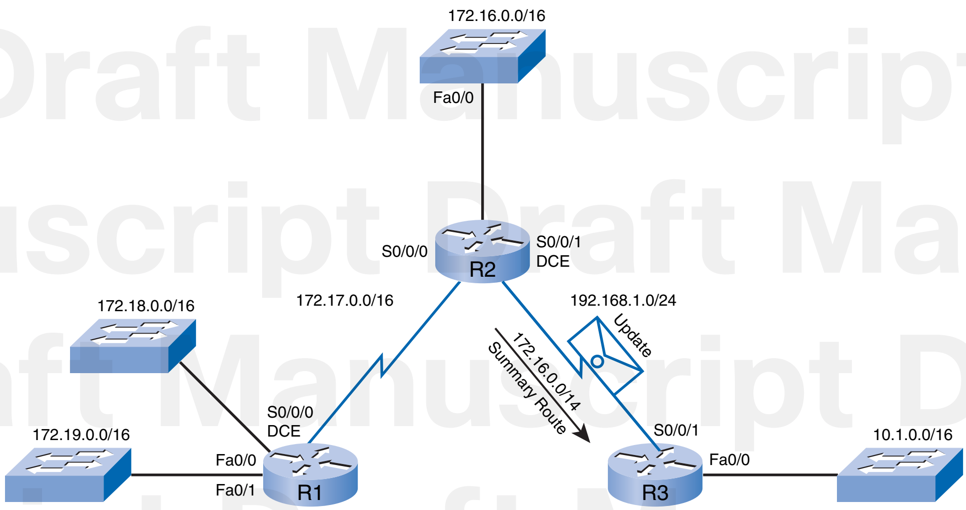

If R2 sends the 172.16.0.0 summary route without the /14 mask, R3 only knows to apply the default classful mask of /16. In a classful routing protocol scenario, R3 is unaware of the 172.17.0.0/16, 172.18.0.0/16, and 172.19.0.0/16 networks.

Note

Using a classful routing protocol, R2 can send these individual networks without summarization, but the benefits of summarization are lost.

Classful routing protocols cannot send supernet routes because the receiving router will apply the default classful mask to the network address in the routing update. If our topology contained a classful routing protocol, R3 would install only 172.16.0.0/16 in the routing table.

Note

When a supernet route is in a routing table, for example, as a static route, a classful routing protocol will not include that route in its updates.

With a classless routing protocol, R2 will advertise the 172.16.0.0 network along with the /14 mask to R3. R3 will then be able to install the supernet route 172.16.0.0/14 in its routing table, giving it reachability to the 172.16.0.0/16, 172.17.0.0/16, 172.18.0.0/16, and 172.19.0.0/16 networks.

VLSM

The Network Fundamentals course described how VLSM allows the use of different masks for each subnet. After a network address is subnetted, those subnets can be further subnetted. As you most likely recall, VLSM is simply subnetting a subnet. VLSM can be thought of as sub-subnetting.

VLSM in Action

Figure 6-6 shows that the network 10.0.0.0/8 has been subnetted using the subnet mask of /16, which gives the potential of 256 subnets:

10.0.0.0/16

10.1.0.0/16

10.2.0.0/16

.

.

.

10.255.0.0/16

In this example, 10.0.0.0/8 has been subnetted into four subnets, 10.0.0.0/16, 10.1.0.0/16, 10.2.0.0/16, and 10.3.0.0/16.

Any of these /16 subnets can be subnetted further. For example, in Figure 6-7, the 10.1.0.0/16 subnet is subnetted again using the /24 mask.

Figure 6-6 VLSM: First Round of Subnets

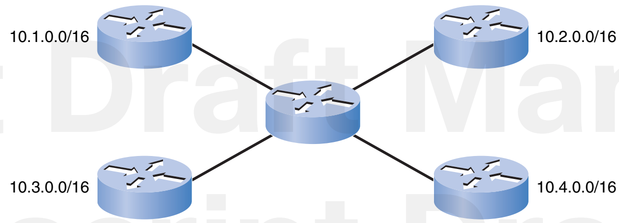

Figure 6-7 VLSM: Subnetting a Subnet

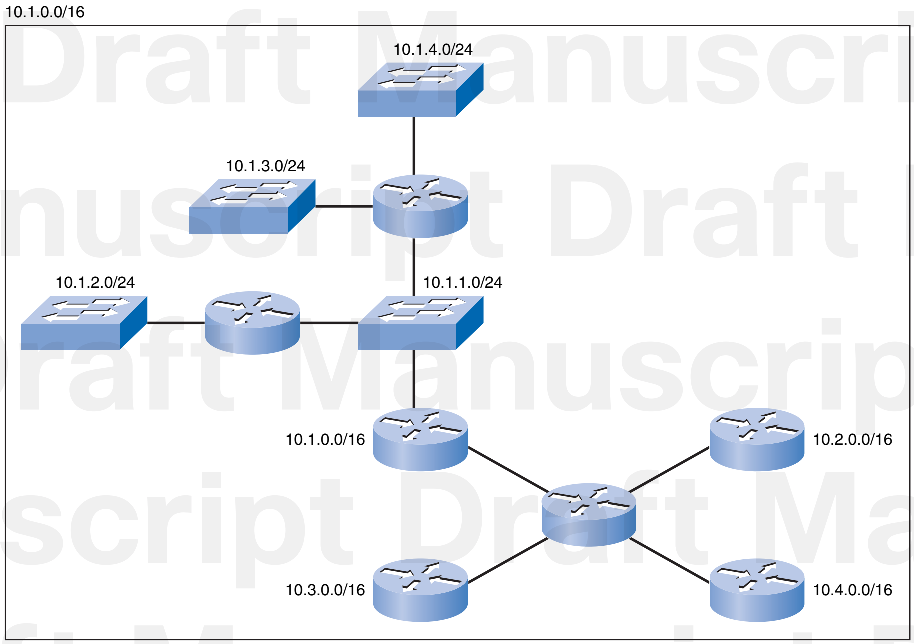

Subnetting 10.1.0.0/16 with a /24 mask results in the following potential subnets:

10.1.1.0/24

10.1.2.0/24

10.1.3.0/24

.

.

.

10.1.255.0/24

Figure 6-8 shows that the 10.2.0.0/16 subnet is also subnetted again with a /24 mask. The 10.3.0.0/16 subnet is subnetted again with the /28 mask, and the 10.4.0.0/16 subnet is subnetted again with the /20 mask.

Figure 6-8 VLSM: Additional Levels of Subnetting

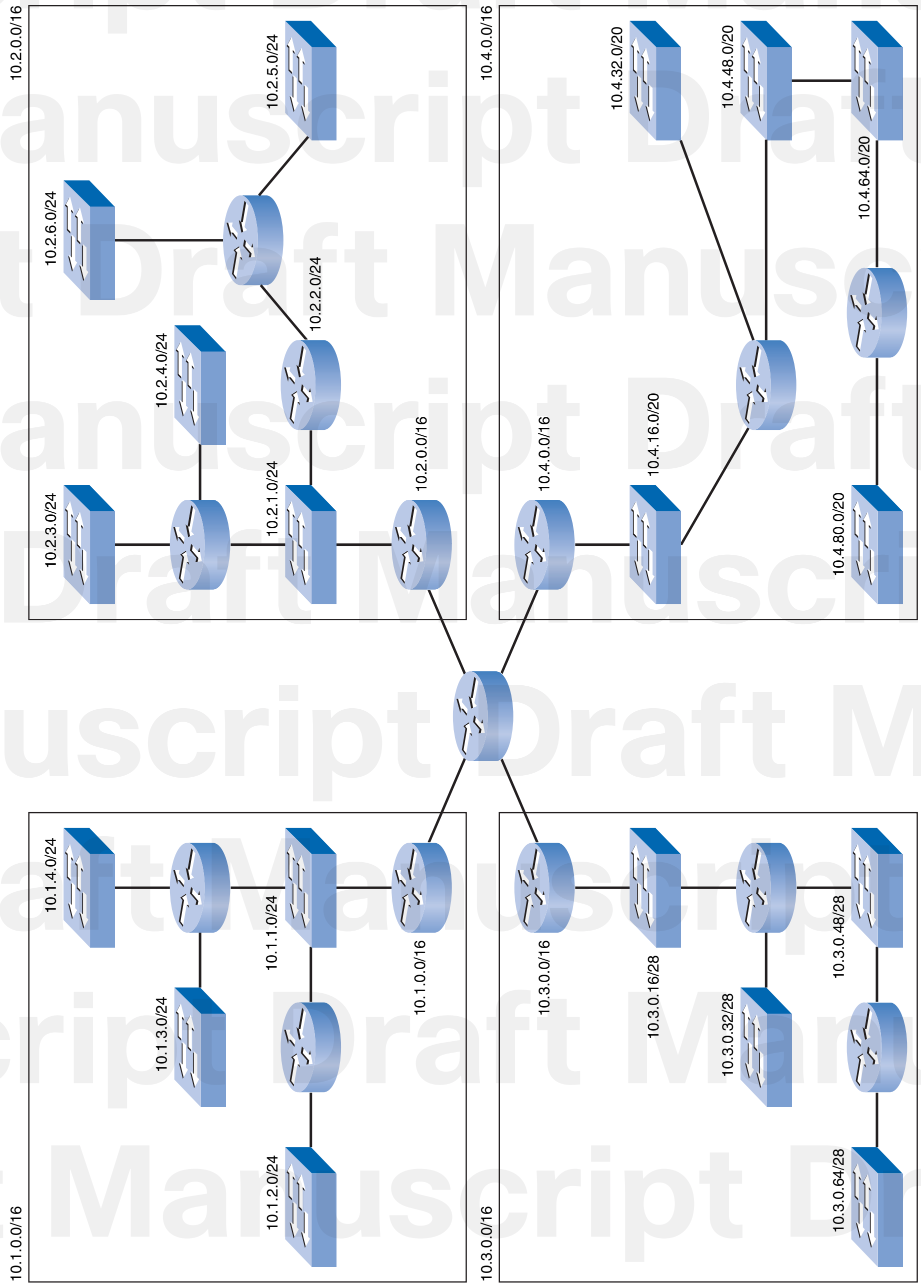

Individual host addresses are assigned from the addresses of "sub-subnets." For example, the figure shows the 10.1.0.0/16 subnet divided into /24 subnets. The 10.1.4.10 address would now be a member of the more specific subnet 10.1.4.0/24.

VLSM and IP Addresses

Another way to view the VLSM subnets is to list each subnet and its sub-subnets. In Figure 6-9, the 10.0.0.0/8 network is the starting address space.

Figure 6-9 Subnets of the Subnet: First Round

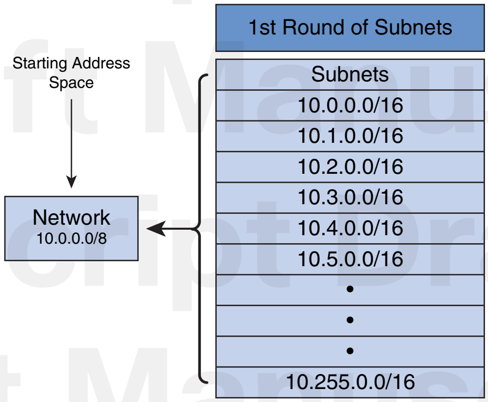

The 10.0.0.0/8 network is subnetted with a /16 mask on the first round of subnetting. You already know that borrowing 8 bits (going from /8 to /16) creates 256 subnets. With classful routing, that is as far as you can go. You can choose only one mask for all your networks. With VLSM and classless routing, you have more flexibility to create additional network addresses and use a mask that fits your needs.

For subnet 10.1.0.0/16 (see Figure 6-10), 8 more bits are borrowed again, to create 256 subnets with a /24 mask.

This mask will allow 254 host addresses per subnet. The subnets ranging from 10.1.0.0/24 to 10.1.255.0/24 are subnets of the subnet 10.1.0.0/16.

Subnet 10.2.0.0/16 is also further subnetted with a /24 mask (see Figure 6-11).

Figure 6-10 Subnets of the Subnet: 10.1.0.0/16

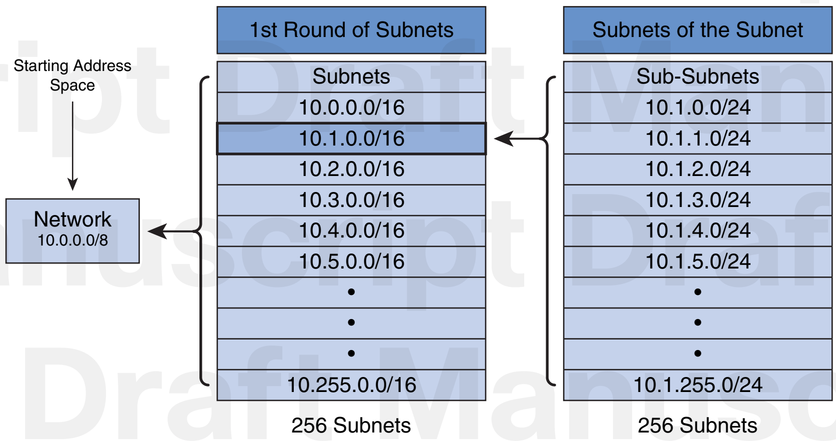

Figure 6-11 Subnets of the Subnet: 10.2.0.0/16

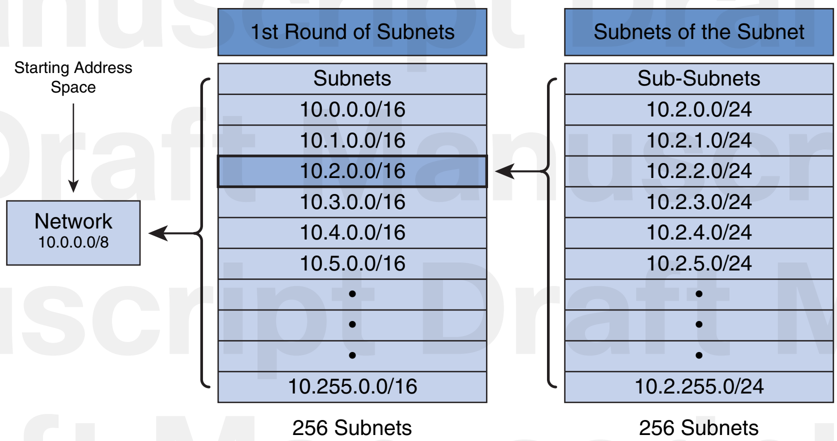

The subnets ranging from 10.2.0.0/24 to 10.2.255.0/24 are subnets of the subnet 10.2.0.0/16.

Subnet 10.3.0.0/16 is further subnetted with a /28 mask (see Figure 6-12).

This mask will allow 14 host addresses per subnet. Twelve bits are borrowed, creating 4096 subnets ranging from 10.3.0.0/28 to 10.3.255.240/28.

Subnet 10.4.0.0/16 is further subnetted with a /20 mask (see Figure 6-13).

Figure 6-12 Subnets of the Subnet: 10.3.0.0/16

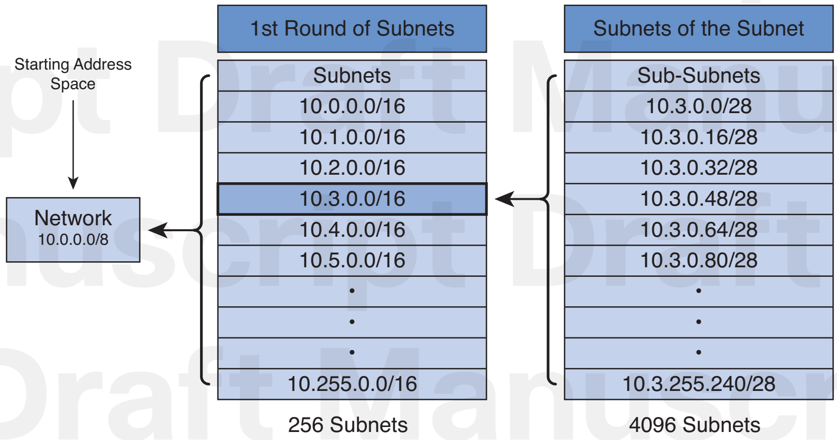

Figure 6-13 Subnets of the Subnet: 10.4.0.0/16

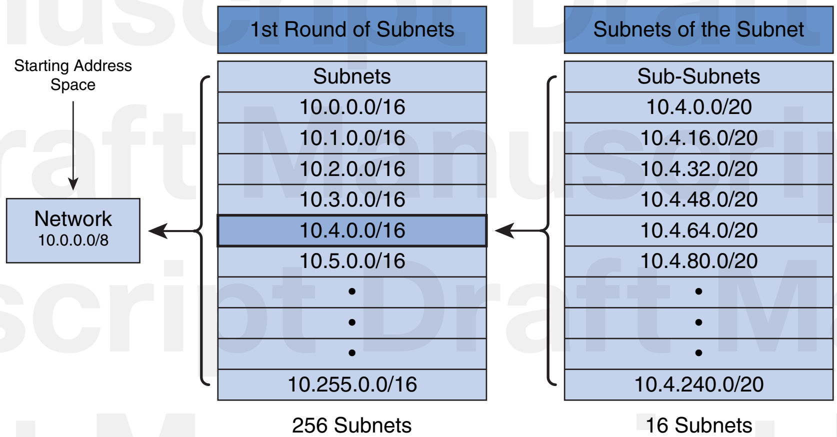

This mask will allow 4094 host addresses per subnet. Four bits are borrowed, creating 16 subnets ranging from 10.4.0.0/20 to 10.4.240.0/20. These /20 subnets are big enough to subnet even further, allowing more networks.

CIDR

Classless interdomain routing (CIDR) is a prefix-based standard for the interpretation of IP addresses. CIDR allows routing protocols to summarize multiple networks, a block of addresses, as a single route. With CIDR, IP addresses and their subnet masks are written as four octets, separated by periods, and followed by a forward slash and a number that represents the subnet mask (slash notation). An example is 172.16.1.0/24.

Route Summarization

As you previously learned, route summarization, also known as route aggregation, is the process of advertising a ***contiguous*** set of addresses as a single address with a less-specific, shorter subnet mask. Remember that CIDR is a form of route summarization and is synonymous with the term ***supernetting***.

You should already be familiar with route summarization that is done by classful routing protocols such as RIPv1. RIPv1 summarizes subnets to a single major network classful address when sending the RIPv1 update out an interface that belongs to another major network. For example, RIPv1 will summarize 10.0.0.0/24 subnets (10.0.0.0/24 through 10.255.255.0/24) as 10.0.0.0/8.

CIDR ignores the limitation of classful boundaries and allows summarization with masks that are less than that of the default classful mask. This type of summarization helps reduce the number of entries in routing updates and reduces the number of entries in local routing tables. It also helps reduce bandwidth utilization for routing updates and results in faster routing table lookups.

Only classless routing protocols can propagate supernets. Classless routing protocols include both the network address and the mask in the routing update. Classful routing protocols cannot include supernets in their routing updates because they cannot apply a mask less than the default classful mask. A static route can be used to configure a supernet route because the network address and mask are configured directly on that router.

Figure 6-14 shows a single static route with the address 172.16.0.0 and the mask 255.248.0.0 summarizing all the 172.16.0.0/16 to 172.23.0.0/16 classful networks.

Figure 6-14 Route Summarization

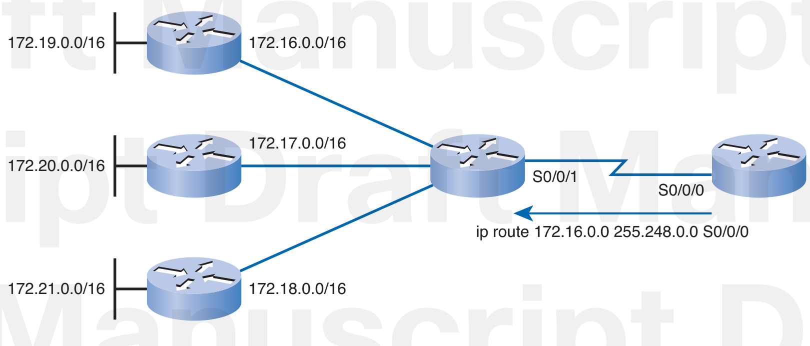

Although 172.22.0.0/16 and 172.23.0.0/16 are not shown in the graphic, these are also included in the summary route. Notice that the /13 mask (255.248.0.0) is less than the default classful mask /16 (255.255.0.0).

Note

You might recall that a supernet is always a route summary, but a route summary is not always a supernet.

A router could have both a specific route entry and a summary route entry covering the same network. Assume that Router X has a specific route for 172.22.0.0/16 using Serial 0/0/1 and a summary route of 172.16.0.0/14 using Serial 0/0/0. Packets with the IP address of 172.22.n.n match both route entries. These packets destined for 172.22.0.0 would be sent out the Serial 0/0/1 interface because there is a more specific match of 16 bits than with the 14 bits of the 172.16.0.0/14 summary route.

Calculating Route Summarization

Calculating route summaries and supernets is identical to the process that you already learned in Chapter 2. Therefore, the following example is presented as a quick review.

Summarizing networks into a single address and mask can be done in three steps, as shown in Figure 6-15.

Figure 6-15 Calculating a Route Summary

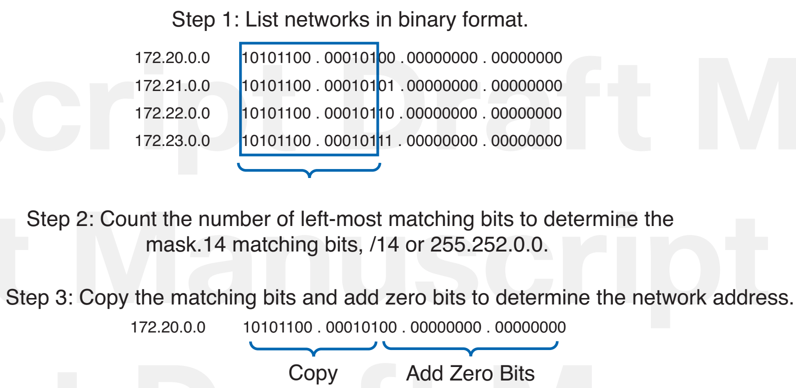

Consider the following four networks:

- 172.20.0.0/16
- 172.21.0.0/16
- 172.22.0.0/16
- 172.23.0.0/16

The steps to summarizing these networks are as follows:

How To

Step 1. List the networks in binary format. The figure shows all four networks in binary.

Step 2. Count the number of leftmost matching bits to determine the mask for the summary route. You can see in the figure that the first 14 leftmost bits match. This is the prefix, or subnet mask, for the summarized route: /14 or 255.252.0.0.

Step 3. Copy the matching bits and then add 0 bits to determine the summarized network address. The figure shows that the matching bits with 0s at the end result in the network address 172.20.0.0. The four networks—172.20.0.0/16, 172.21.0.0/16, 172.22.0.0/16, and 172.23.0.0/16—can be summarized into the single network address and prefix 172.20.0.0/14.

Summary

Classless interdomain routing (CIDR) was introduced in 1993, replacing the previous generation of IP address syntax, classful networks. CIDR allowed more efficient use of IPv4 address space and prefix aggregation, known as route summarization or supernetting.

With CIDR, address classes (Class A, Class B, and Class C) became meaningless. The network address was no longer determined by the value of the first octet, but assigned a prefix length (subnet mask). Address space, the number of hosts on a network, could now be assigned a specific prefix depending on the number of hosts needed for that network.

CIDR allows supernetting. A supernet is a group of major network addresses summarized as a single network address with a mask less than that of the default classful mask.

CIDR uses VLSM (variable-length subnet masks) to allocate IP addresses to subnetworks according to need rather than by class. VLSM allows subnets to be further divided or subnetted into even smaller subnets. Simply put, VLSM is just subnetting a subnet.

Propagating CIDR supernets or VLSM subnets require a classless routing protocol. A classless routing protocol includes the subnet mask along with the network address in the routing update. Classless routing protocols include RIPv2, EIGRP, OSPF, IS-IS, and BGP.

Determining the summary route and subnet mask for a group of networks can be done in three easy steps. The first step is to list the networks in binary format. The second step is to count the number of leftmost matching bits. This will give you the prefix length or subnet mask for the summarized route. The third step is to copy the matching bits and then add 0 bits to the rest of the address to determine the summarized network address. The summarized network address and subnet mask can now be used as the summary route for this group of networks. Summary routes can be used by both static routes and classless routing protocols. Classful routing protocols can only summarize routes to the default classful mask.

Classless routing protocols and their ability to support CIDR supernet, VLSM, and discontiguous networks are described in the following chapters.

Activities and Labs

The activities and labs available in the companion *Routing Protocols and Concepts, CCNA Exploration Labs and Study Guide* [ISBN 1-58713-204-4] provide hands-on practice with the following topics introduced in this chapter:

Activity 6-1: Basic VLSM Calculation and Addressing Design Activity (6.4.1)

In this activity, you will use the network address 192.168.1.0/24 to subnet and provide the IP addressing for a given topology. VLSM will be used so that the addressing requirements can be met using the 192.168.1.0/24 network.

Activity 6-2: Challenge VLSM Calculation and Addressing Design Activity (6.4.2)

In this activity, you will use the network address 172.16.0.0/16 to subnet and provide the IP addressing for a given topology. VLSM will be used so that the addressing requirements can be met using the 172.16.0.0/16 network.

Activity 6-3: Troubleshooting a VLSM Addressing Design Activity (6.4.3)

In this activity, the network address 172.16.128.0/17 was used to provide the IP addressing for a network. VLSM has been used to subnet the address space incorrectly. You will need to troubleshoot the addressing that was assigned to each subnet to determine where errors are present and determine the correct addressing assignments where needed.

Activity 6-4: Basic Route Summarization Activity (6.4.4)

In this activity, you are given a network with subnetting and address assignments already completed. Your task is to determine summarized routes that can be used to reduce the number of entries in routing tables.

Activity 6-5: Challenge Route Summarization Activity (6.4.5)

In this activity, you are given a network with subnetting and address assignments already completed. Your task is to determine summarized routes that can be used to reduce the number of entries in routing tables.

Activity 6-6: Troubleshooting Route Summarization Activity (6.4.6)

In this activity, the LAN IP addressing is already completed for the network. VLSM was used to subnet the address space. The summary routes are incorrect. You will need to troubleshoot the summary routes that have been assigned to determine where errors are present and determine the correct summary routes.

Packet Tracer Companion

Many of the hands-on labs include Packet Tracer Companion Activities, where you can use Packet Tracer to complete a simulation of the lab. Look for this icon in *Routing Protocols and Concepts, CCNA Exploration Labs and Study Guide* (ISBN 1-58713-204-4) for hands-on labs that have a Packet Tracer Companion.

Check Your Understanding

Complete all the review questions listed here to test your understanding of the topics and concepts in this chapter. The section, "Check Your Understanding and Challenge Questions Answer Key" at the end of this chapter lists the answers.

1. For each of the following routing protocols, indicate whether it supports VLSM (VLSM or non-VLSM).

   RIPv1:

   EIGRP:

   IGRP:

   IS-IS:

   OSPF:

   RIPv2:

2. For each of the following definitions, indicate whether it is describing VLSM or route summarization.

   Combining several IP network addresses in one IP address:

   Ability to specify a different subnet mask for the same network number and different subnets:

   Used in supernetting:

   Conserves address space:

   Used to reduce the number of entries in a routing table:

3. What two methods allowed the continued use of IPv4 addressing and helped delay the need to implement IPv6?

   A. Variable-length subnetting.
   B. The IPv4 address range was expanded.
   C. Private addresses were used with address translation.
   D. Classful routing was implemented.
   E. IPv4 was abandoned in favor of IPv6 for all hosts.
   F. Supernetting was implemented.

4. The following subnet masks have been chosen for use with the 192.168.16.0 network:

   255.255.255.252

   255.255.255.240

   255.255.255.192

   Which of the following identify the most efficient use for each of these masks? (Choose three.)

   A. Use the /30 mask for point-to-point links, such as WAN connections.
   B. Use the /30 mask for subnetworks of four or more hosts.
   C. Use the /28 mask for small subnetworks with up to 14 hosts.
   D. Use the /26 mask for larger subnetworks with up to 62 hosts.
   E. Use the /25 mask for subnetworks with up to 30 hosts.
   F. Use the /24 mask for point-to-point links, such as WAN connections.

5. When using a classful Class A IP address scheme, how many octets are used to designate the network portion of the address?

   A. 1
   B. 2
   C. 3
   D. 4

6. Match the VLSM subnet with the network size it is most appropriate for. Each answer can be used only once.

   VLSM subnets:

   172.16.64.0/18

   172.16.16.64/30

   172.16.128.0/19

   172.16.18.0/24

   172.16.5.128/26

   Number of hosts:

   A. 2
   B. 60
   C. 250
   D. 8000
   E. 16,000

7. A network engineer is summarizing the two groups of routes, group A and group B, on Router R1. Which summarization will work for all subnets?

Group A:

192.168.0.0/30

192.168.0.4/30

192.168.0.8/30

192.168.0.16/29

Group B:

192.168.4.0/30

192.168.5.0/30

192.168.6.0/30

192.168.7.0/29

A. 192.168.0.0/23
B. 192.168.0.0/22
C. 192.168.0.0/21
D. 192.168.0.0/28

8. How many bits are used in the IPv4 address space?

A. 8
B. 12
C. 16
D. 30
E. 32
F. 64

9. For the following classful network addresses, indicate whether the address is a Class A address or a Class B address.

191.254.45.0: Class

123.90.78.45: Class

128.44.0.23: Class

129.68.11.45: Class

126.0.0.0: Class

125.33.23.56: Class

10. Refer to Figure 6-16. The network administrator wants to minimize the number of entries in Router1's routing table. What should the administrator implement on the network?

Figure 6-16 Check Your Understanding: Question #10

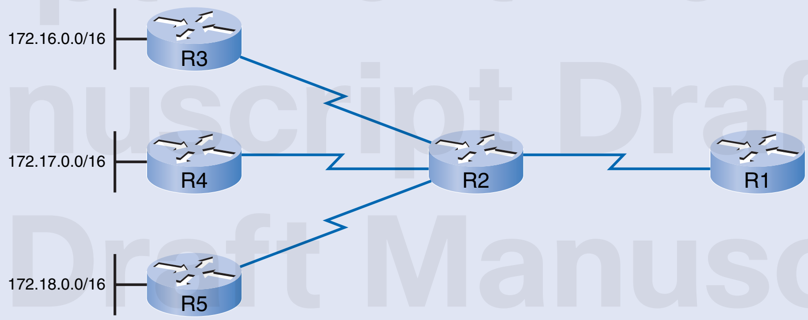

A. VLSM

B. CIDR

C. Private IP addresses

D. Classful routing

11. What distinguishes a classless routing protocol from a classful routing protocol?
12. What are the advantages of using a classless routing protocol?
13. How does a classful routing protocol determine the subnet mask of a routing update?
14. Why did the IETF introduce classless IP addressing, CIDR?
15. What term is used to define the process of subnetting a subnet?

Challenge Questions and Activities

These questions require a deeper application of the concepts covered in this chapter and are similar to the style of questions you might see on a CCNA certification exam. You can find the answers in the appendix.

1. The 172.16.0.0/16 network is subnetted using a /24 subnet mask. What could you do if you needed to divide the 172.16.10.0/24 subnet into three equal subnets with the maximum number of hosts in each subnet?

2. 172.16.10.0/24 is using the following /28 subnets for LANs:

   172.16.10.16/28

   172.16.10.32/28

   172.16.10.48/28

   172.16.10.64/28

   172.16.10.80/28

   172.16.10.96/28

   172.16.10.112/28

   172.16.10.128/28

   172.16.10.144/28

   172.16.10.160/28

   172.16.10.176/28

   172.16.10.192/28

   172.16.10.240/28

   The network administrator wants to allocate a /28 subnet and subnet it further with a /30 mask for all point-to-point serial links in the network. What /28 subnets are available to be used?

3. What is supernetting? What is required to propagate a supernet route?

4. Summarize the following networks:

   192.168.68.0/24

   192.168.96.0/24

   192.168.80.0/24

Packet Tracer Challenge

Look for this icon in *Routing Protocols and Concepts, CCNA Exploration Labs and Study Guide* (ISBN 1-58713-204-4) for instructions on how to perform the Packet Tracer Skills Integration Challenge for this chapter.

To Learn More

RFC 1519, Classless Inter-Domain Routing (CIDR)

Requests For Comments (RFC) are a series of documents submitted to the IETF (Internet Engineering Task Force) to propose an Internet standard or convey new concepts, information, or occasionally even humor.

RFCs can be accessed from several websites, including http://www.ietf.org. Read all or parts of RFC 1519 to learn more about the introduction of CIDR to the Internet community.

Internet Core Routers

In the "To Learn More" section of Chapter 3, "Introduction to Dynamic Routing Protocols," you accessed route servers to display BGP routes on the Internet. One such site is http://www.traceroute.org.

Access one of the route servers, and using the **show ip route** command, view the actual routing table of an Internet router. Notice how many routes there are on an Internet core router. As of March 2007, there were over 200,000 routes. Many of these are summarized routes and supernets. Use the **show ip route 207.62.187.0** command to view one such supernet.

CAIDA

An interesting website is CAIDA, the Cooperative Association for Internet Data Analysis, http://www.caida.org. CAIDA "provides tools and analyses promoting the engineering and maintenance of a robust, scalable global Internet infrastructure." There are several sponsors for CAIDA including Cisco Systems. Although much of this information might seem beyond your understanding, you will begin to recognize many of these terms and concepts.

Check Your Understanding and Challenge Questions Answer Key

Check Your Understanding

1. Answer:

   RIPv1: non-VLSM

   EIGRP: VLSM

   IGRP: non-VLSM

   IS-IS: VLSM

   OSPF: VLSM

   RIPv2: VLSM

   Classful routing protocols such as RIPv1 and IGRP do not support VLSM. Classless routing protocols such as RIPv2, EIGRP, OSPF, and IS-IS do support VLSM. To support VLSM, a routing protocol must include the subnet mask in the routing update. Therefore, only classless routing protocols support VLSM.

2. Combining several IP network addresses in one IP address: route summarization

   Ability to specify a different subnet mask for the same network number and different subnets: VLSM

   Used in supernetting: route summarization

   Conserves address space: VLSM

   Used to reduce the number of entries in a routing table: route summarization

3. A, C. Variable-length subnetting allowed networks to be subnetted with various numbers of hosts, therefore making better use of network address space. Private addressing with Network Address Translation/Port Address Translation (NAT/PAT) allowed networks to have a larger number of hosts than the public network address they were allocated.

4. A, C, D. 255.255.255.252 is equivalent to /30, 255.255.255.240 is equivalent to /28, and 255.255.255.192 is equivalent to /26.

5. A. A Class A network has a default subnet mask of 255.0.0.0. The first octet represents the network portion of the address and the last three octets represent the host portion.

6. 172.16.64.0/18: E

   172.16.16.64/30: A

   172.16.128.0/19: D

   172.16.18.0/24: C

   172.16.5.128/26: B

   By examining each of the subnet masks, you can determine how many bits are available for the hosts. For example, a /30 subnet mask has 30 bits for the network address, leaving 2 bits for the hosts. Two bits gives a total of four hosts, but the first host address is used for the subnet or network address and the last host address is used for the broadcast address, leaving you with two usable hosts.

7. C. If you write out all the addresses in binary, you will notice that the first 21 bits match, or a /21 subnet mask. Copy the matching bits and add all 0s to the end, and this will give you the summary address of 192.168.0.0.

8. E. IPv4 uses 32-bit addresses (four octets of 8 bits each).

9. 191.254.45.0: Class B

   123.90.78.45: Class A

   128.44.0.23: Class B

   129.68.11.45: Class B

   126.0.0.0: Class A

   125.33.23.56: Class A

   Class A addresses range from 0.0.0.0 to 127.255.255.255. Class B addresses range from 128.0.0.0 to 191.255.255.255.

10. A. CIDR allows route aggregation or route summarization. The 172.16.0.0/16, 172.17.0.0/16, and 172.18.0.0/16 networks can be summarized to 172.16.0.0/14. The /14 is a mask that is less than the classful mask, so this would be considered a supernet.

11. A classless routing protocol includes the subnet mask in the routing update.

12. Classless routing protocols allow the use of VLSM and CIDR within the routing domain.

13. If the network address in the routing update is on the same major classful network as the receiving interface, the classful routing protocol will use the same mask as the interface; otherwise, it will use the default classful mask.

14. The Internet routers' routing tables were experiencing very fast growth. There needed to be a way to summarize the routes in the routing table.

    Because classful addressing only provided /8, /16, or 24 masks, the IPv4 address space was becoming depleted.

15. VLSM

Challenge Questions and Activities

1. 172.16.10.0/24 could be subnetted using the /26 subnet mask.

   <-Network-> | <-Host->

   172.16.10. 0 0|0 0 0 0 0

   172.16.10. 0 1|0 0 0 0 0

   172.16.10. 1 0|0 0 0 0 0

   172.16.10. 1 1|0 0 0 0 0

   This would give four subnets with the maximum number of hosts. You cannot do three subnets; the next lower number of subnets would be two.

2. 172.16.10.0/28

   172.16.10.208/28

   172.16.10.224/28

3. Supernetting refers to the ability to summarize networks less than the classful default mask. Classless routing protocols are able to propagate a supernet route because they include the subnet mask with the summarized network address.

4. 192.168.64.0/18