



12

Advanced IP Addressing

CERTIFICATION OBJECTIVES

12.01 Variable-Length Subnet Masking

12.02 Route Summarization



Two-Minute Drill

Q&A Self Test

In Chapter 11, you were introduced to two advanced routing protocols: OSPF and EIGRP. Both of these protocols are classless protocols that support advanced IP addressing concepts, including variable-length subnet masking (VLSM) and route summarization. This chapter focuses on these two advanced concepts.

CERTIFICATION OBJECTIVE 12.01

VLSM

VLSM, defined in RFC 1812, allows you to apply different subnet masks to the same class address. For instance, a good mask for point-to-point links is 255.255.255.252, which provides for two host addresses in each network. A good mask for a LAN connection might be 255.255.255.192, which provides for 62 host addresses for each network. Using a 255.255.255.252 mask for a LAN connection will not give you enough host address, and using a 255.255.255.192 mask on a point-to-point connection wastes addresses. One solution would be to divide the mask values in the middle to limit the waste of addresses, but this doesn't scale well. VLSM solves this problem by enabling you to use different subnet mask values on the same class address space. The following sections cover the advantages that VLSM provides as well as how to use VLSM in your own network.

Features of VLSM

VLSM enables you to have more than one mask for a given class of address, albeit a class A, B, or C network number. Classful protocols, such as RIPv1 and IGRP, do not support VLSM. To deploy VLSM requires a routing protocol that is classless—BGP, EIGRP, IS-IS, OSPF, or RIPv2, for instance. VLSM provides two major advantages:

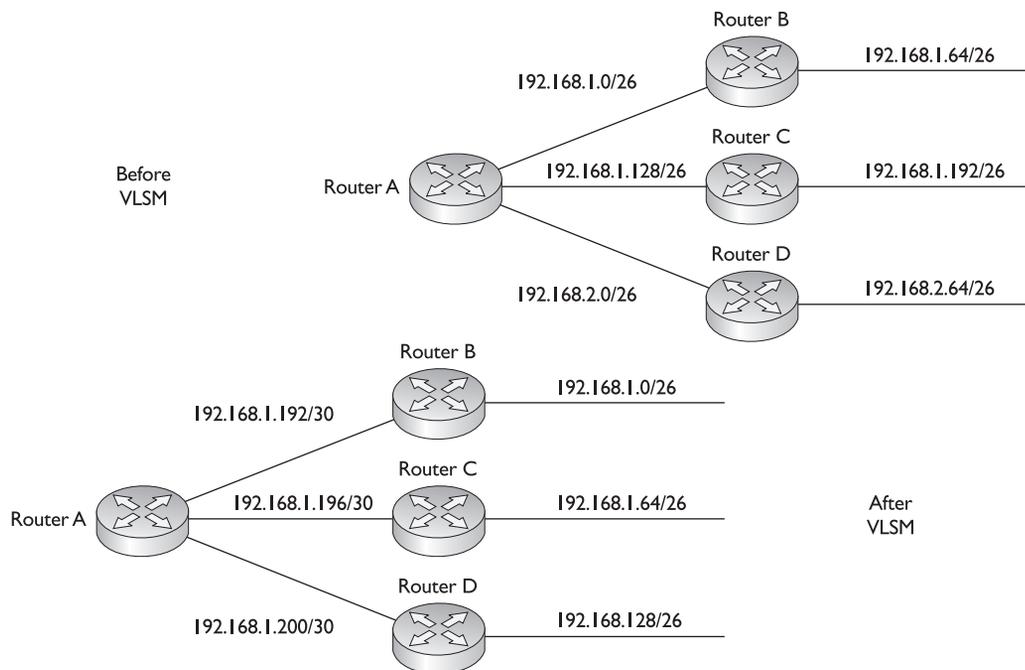
- More efficient use of addressing
- Ability to perform route summarization

As these bullets suggest, one advantage of VLSM is that it allows you to make more efficient use of your IP addressing. Figure 12-1

exam

Watch

VLSM allows you to use more than one subnet mask for a given class address. Remember the two major advantages of VLSM.

FIGURE 12-1 Using VLSM

shows a simple before-and-after example of using VLSM. In this example, there is a router at the corporate site (Router A) with point-to-point WAN connections to the remote office routers (Router B, Router C, and Router D). The LAN segments at these remote sites have about 50 devices (thus the /26) mask. In the *before* design, a single subnet mask is chosen: 255.255.255.192, which allows 62 hosts per subnet. Because of the number of segments, two Class C networks were needed. On the WAN segments, this wastes a lot of addressing space.

The bottom part of Figure 12-1 shows a more efficient use of addressing, using VLSM. In this example, the three remote sites have a 255.255.255.192 mask, but the last subnet, 192.168.1.192/26 was assigned a *different* subnet mask. And these little subnetted subnets were then assigned to the point-to-point links of the WAN connections. Given the VLSM solution, only one Class C network was needed to assign addressing to this network. The second advantage of VLSM, route summarization, is discussed later in this chapter.

Addressing with VLSM

In order to use VLSM, you must be very familiar with IP addressing and normal subnetting. If you have not grasped these concepts yet, then VLSM will be out of your reach. If you are still uncomfortable with IP addressing and subnetting, please review Chapter 3. As I have already mentioned, VLSM basically means taking a subnet (not a network number) and applying a different subnet mask to it. This section covers how to create an efficient addressing scheme using VLSM.

There are certain steps that you should follow when performing VLSM:

1. Find the largest segment in the area—the segment with the largest number of devices connected to it.
2. Find the appropriate subnet mask for the largest network segment.
3. Write down your subnet numbers to fit your subnet mask.
4. For your smaller segments, take one of these newly created subnets and apply a different, more appropriate, subnet mask to it.
5. Write down your newly subnetted subnets.
6. For even smaller segments, go back to step 4.

Actually, you can take a subnetted subnet and subnet it again! With this process, you can come up with a very efficient addressing scheme.

For example, let's assume that you have a class C network (192.168.1.0) and three LAN segments: one with 120 devices, one with 60 devices, and one with 30 devices. In steps 1 and 2, you find the largest segment and an appropriate subnet mask for it. This would be the segment with 120 devices. To accommodate the 120 devices, you would need a subnet mask of 192.168.1.0/25. In step 3, write down the newly created subnets: 192.168.1.0/25 and 192.168.1.128/25. We'll assign the first subnet to this LAN segment. We now have two segments left: 60 and 30 devices. Again, start with the larger segment first. Next we perform step 4. Which subnet mask is appropriate for 60 devices? If you guessed /26 (255.255.255.192), then you guessed correctly—this gives you 62 host addresses. In step 5, you write down your newly created subnetted subnets: 192.168.1.128/26 and 192.168.1.192/26. Let's assign 192.168.1.128/26 to the segment with 60 devices.

This leaves us with one extra subnet. You could easily assign it to this segment, but this segment needs only 30 hosts and the mask has 62 hosts, which is not the most efficient mask. If you want, you can go back to step 4 and repeat the process for this subnet. The mask /27 (255.255.255.224) is a subnet mask that results in 30 host addresses, resulting in two more subnets: 192.168.1.192/27 and 192.168.1.224/27. In this example, you have one extra subnet that you could use for future growth!

exam**Watch**

Remember how to perform VLSM. Therefore, know the six steps mentioned previously and also go over the next two VLSM examples thoroughly.

As you can see, with VLSM, you can be very efficient in your IP addressing design. I do recommend that you leave room in each subnet for future growth. For instance, in the last example, using a mask of /27 on the 192.168.1.192 subnet creates two more subnets, each with 30 host addresses. If you use this address scheme and the 30-host segment grows, then you'll have

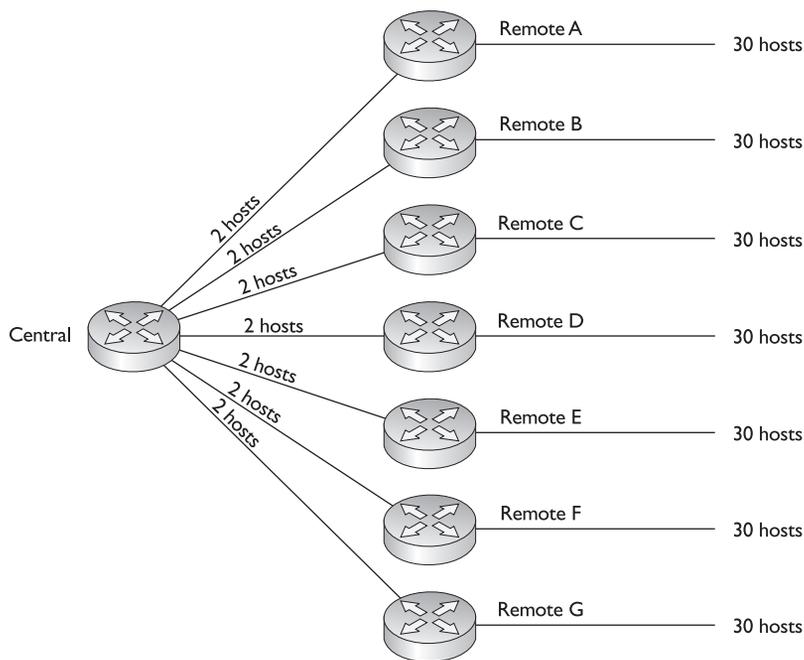
to go back and readdress a portion of your network, which is not fun.

VLSM Example I

Now that you have an understanding of the basics of performing VLSM, let's look at a more difficult example. I'll use the network shown in Figure 12-2. In this example, you are given a Class C network: 192.168.2.0/24. You are tasked to use VLSM to accommodate the following requirements: each remote site (total of 7) has no more than 30 devices. The links between the central and remote routers are point-to-point.

FIGURE 12-2

VLSM example I



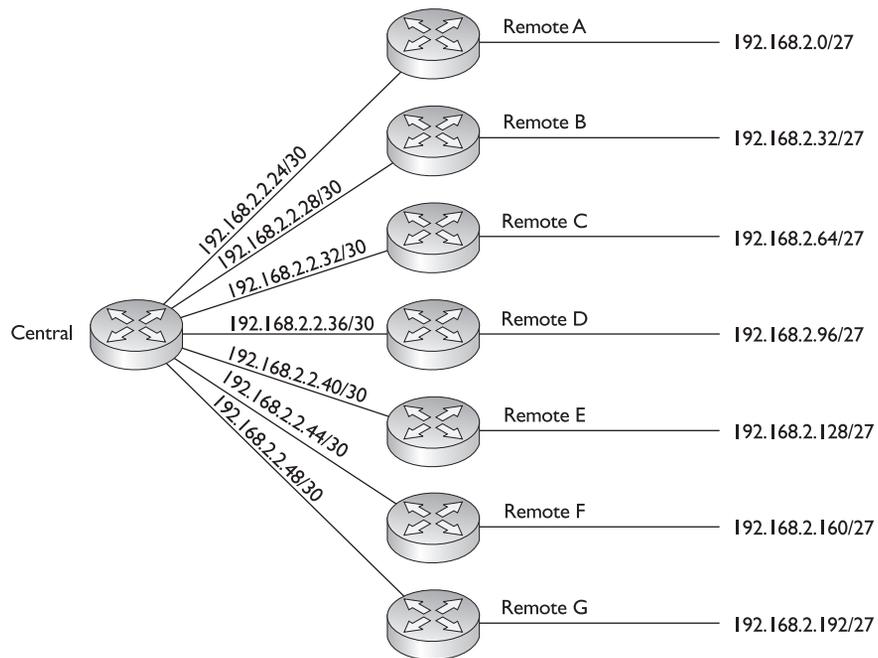
For this example, you first need to worry about handling the largest segments, which are the remote sites with 30 devices. To handle 30 devices, you need a $225.255.255.224$ (/27) subnet mask. This mask results in the following subnets: $192.168.2.0/27$, $192.168.2.32/27$, $192.168.2.64/27$, $192.168.2.96/27$, $192.168.2.128/27$, $192.168.2.160/27$, $192.168.2.192/27$, and $192.168.2.224/27$.

With a /27 mask, you have actually created eight subnets; however, you need only seven for the remote offices. This leaves you one subnet mask, but seven point-to-point links between the central and remote routers. Let's assign the first seven subnets for the remote LAN segments and use the last subnet ($192.168.2.224/27$) for the point-to-point links. To accommodate the point-to-point links, use a $255.255.255.252$ (/30) subnet mask. This results in the following subnetted subnets: $192.168.2.224/30$, $192.168.2.228/30$, $192.168.2.232/30$, $192.168.2.236/30$, $192.168.2.240/30$, $192.168.2.244/30$, $192.168.2.248/30$, and $192.168.2.252/30$.

With a /30 mask on the $192.168.2.224$ subnet, you have created eight little subnets. You need only seven for the point-to-point links, which leaves one small subnet left over. Figure 12-3 shows the actual networking layout based on this example. Notice that this example used two subnet mask values: $255.255.255.224$ and $255.255.255.252$.

FIGURE 12-3

VLSM example 1
address design



VLSM Example 2

Let's take a look at another example, shown in Figure 12-4. You have been given a class C network: 192.168.3.0, with the addressing requirements shown in the figure. You need to come up with a VLSM solution to accommodate all of the network segments with the single class C network.

For this example, the first network you should worry about is the backbone router segment, since it is the largest. It requires 126 host addresses. Therefore, you will need to use a subnet mask that accommodates 126 hosts: 255.255.255.128 (/25). This results in two subnets: 192.168.3.0/25 and 192.168.3.128/25. Let's assign the first subnet to the backbone router, leaving us the second subnet for further subnetting.

You next need to be concerned about the second largest subnet: the smaller router LAN segments. Each of these sites needs networks that will accommodate 30 host addresses. Take the remaining subnet (192.168.3.128/25) and apply a mask to it that will give you your remote site's addresses. The mask of 255.255.255.224 (/27) will do this for you. This results in the following subnets: 192.168.3.128/27, 192.168.3.160/27, 192.168.3.192/27, and 192.168.3.224/27.

You now have four subnets, with 30 host addresses each. Take the first three of these and assign them to your smaller router LAN segments. Use the last subnet for your router-to-router connections. These links need six host addresses each. A 255.255.255.248 (/29) subnet mask will accommodate your addressing needs. Applying this to the fourth subnet results in the following smaller subnets: 192.168.3.224/29, 192.168.3.232/29, 192.168.3.240/29, and 192.168.3.248/29. You need only three of these subnets, leaving one for future growth. In all, this network design, shown in Figure 12-5, used three different subnet masks: 255.255.255.128, 255.255.255.224, and 255.255.255.248.

FIGURE 12-4

VLSM example 2

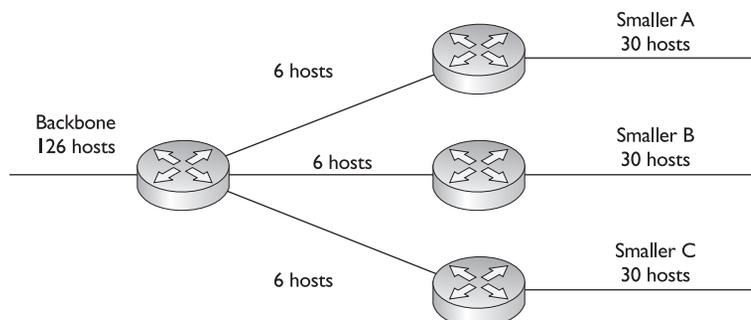
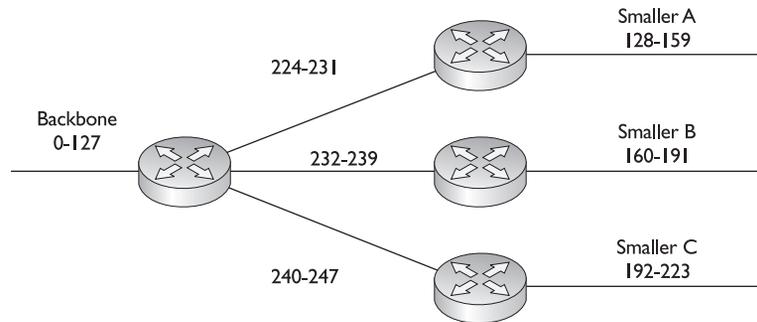


FIGURE 12-5

VLSM example 2
address design



CERTIFICATION OBJECTIVE 12.02

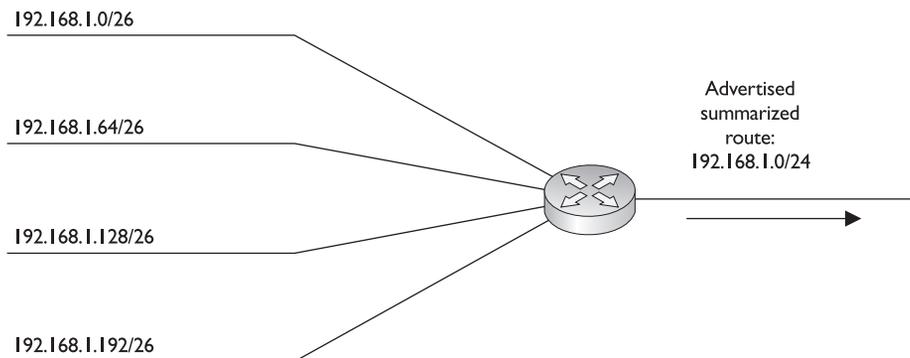
Route Summarization

Route summarization is the ability to take a bunch of contiguous network numbers in your routing table and advertise these contiguous routes as a single summarized route. VLSM allows you to summarize subnetted routes back to the class boundary. For instance, if you have 192.168.1.0/24 and have subnetted it to 192.168.1.0/26, giving you four networks, you could summarize these subnets in your routing table and advertise them as the Class C network number 192.168.1.0/24, as is shown in Figure 12-6.

In this example, you have reduced your routing entries from 4 down to 1 in your routing updates. Summarization is a form of VLSM. Notice in the preceding

FIGURE 12-6

Simple route
summarization
example



example that the same class network, 192.168.1.0, has two masks associated with it: 255.255.255.192 and 255.255.255.0.

Advantages of Summarization

Summarization allows you to create a more efficient routing environment by providing the following advantages:

- It reduces the size of routing tables, requiring less memory and processing.
- It reduces the size of updates, requiring less bandwidth.
- It contains network problems.

As you can see from the design shown in Figure 12-6, the size of the routing table update was reduced from four routes to one route, which requires less processing to any routers receiving this information. Thus, less bandwidth is required to advertise the update and less memory and processing are required on the receiving routers to process the update.

Another advantage of route summarization is that it helps contain certain kinds of network problems. For example, assume that 192.168.1.64/26 was going up and down, up and down (a flapping route). This condition obviously affects the connected router and any router that knows about this specific subnet. However, routers that only know the summarized route are not affected by the subnet that is flapping. In order for these routers to be affected, all four subnets would have to fail, causing the router performing the summarization to stop advertising the summarized route.

This, obviously, is an advantage, but it does have a down side. Route summarization hides the complete picture of the network. This can cause problems with routers making bad assumptions. For instance, assume that 192.168.1.64/26 really is down but that routers in another part of the network are still receiving updates

exam

watch

Route summarization is taking a bunch of contiguous network numbers in a routing table and reducing them to a smaller number of routes.

Route summarization benefits include smaller routing tables and updates and containment of networking problems.

concerning the summarized route (192.168.1.0/24). From their perspective, since the router summarizing the route is still advertising this route, all addresses from 192.168.1.0 through 192.168.1.255 must be available. Obviously, this is not true, and thus other routers will still send traffic to 192.168.1.64/26, since they still think it's reachable.

Classless Interdomain Routing

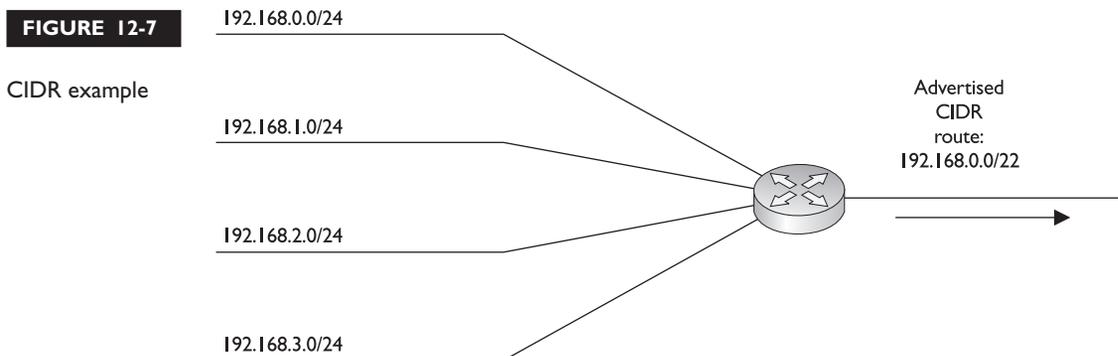
Classless Interdomain Routing (CIDR), specified in RFC 2050, is an extension to VLSM and route summarization. With VLSM, you can summarize subnets back to the Class A, B, or C network boundary. For example, if you have a Class C network 192.168.1.0/24 and subnet it with a 26-bit mask, you have created four subnets. Using VLSM and summarization, you can summarize these four subnets back to 192.168.1.0/24. CIDR takes this one step further and allows you to summarize a block of contiguous class A, B, and C network numbers. This practice is commonly referred to as *supernetting*. Today's

classless protocols support supernetting. However, it is most commonly configured by ISPs on the Internet using BGP.

Figure 12-7 shows an example of CIDR. In this example, a router is connected to four networks: 192.168.0.0/24, 192.168.1.0/24, 192.168.2.0/24, and 192.168.3.0/24. The router is summarizing these routes into a single entry: 192.168.0.0/22. Notice the subnet mask for this summarization:

255.255.252.0. This mask, along with the beginning network, 192.168.0.0, includes addresses from 192.168.0.0 to 192.168.3.255, which are behind this router.

exam
Watch
CIDR is similar to VLSM in that CIDR allows you to summarize multiple contiguous class networks together, like multiple Class C networks. This is also called supernetting.



Hierarchical Addressing

In order to perform route summarization, you will need to set up your addressing in a hierarchical fashion. Hierarchical addressing provides the following benefits:

- It enables more efficient routing.
- It uses route summarization to decrease the size of routing tables.
- It decreases the amount of memory needed to store the smaller routing tables.
- It decreases the impact on the router when needing to rebuild the routing table.
- It provides a design to simplify your troubleshooting process.

exam
Watch *Understand the benefits of hierarchical addressing.*

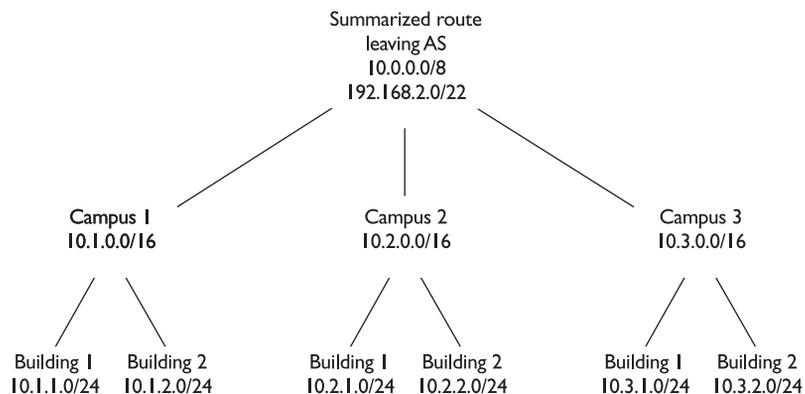
Figure 12-8 shows a simple example of hierarchical addressing. In this example, the network is using 10.0.0.0/8. This is summarized before being sent to another network. This addressing space is broken up into three campuses: 10.1.0.0/16, 10.2.0.0/16, and 10.3.0.0/16. Each of these sets of addresses is summarized

when sharing routes between the campuses. Within each campus, the addressing is further broken up for the two buildings: 10.x.1.0/24 and 10.x.2.0/24.

To implement a hierarchical addressing design and to take advantage of route summarization, you'll need a routing protocol that supports VLSM: BGP, EIGRP,

FIGURE 12-8

Simple hierarchical addressing example



e x a m
W a t c h **Memorize the preceding three bulleted points.**

IS-IS, OSPF, or RIPv2. And when implementing route summarization, you'll need to consider the following items:

- The routing protocol must carry the subnet mask with the corresponding network entries.
- Routing decisions must be made on the entire destination IP address.
- In order to summarized routing entries, they must have the same highest-order matching bits.

Routing and Subnet Masks

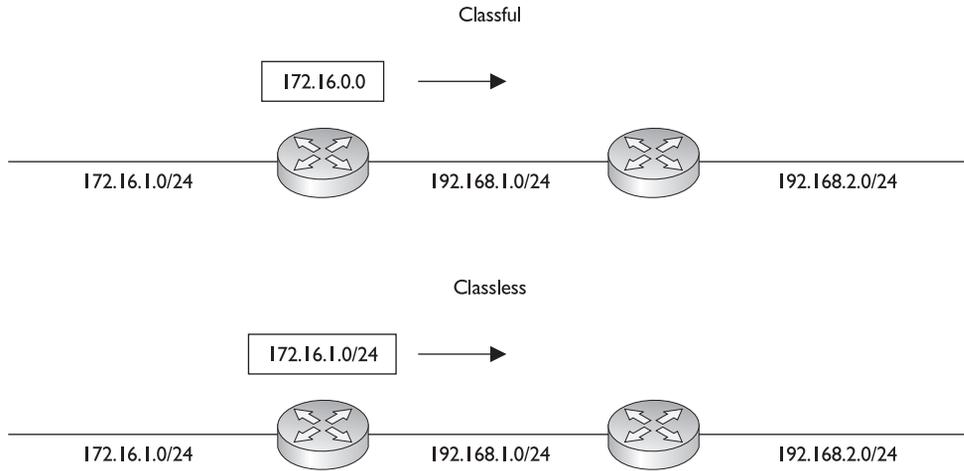
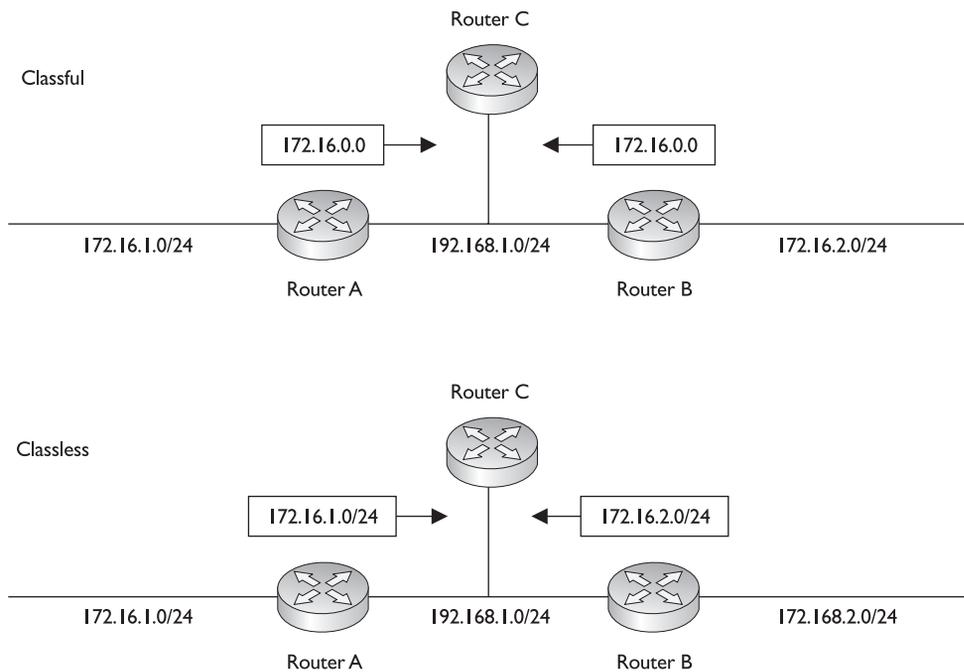
As mentioned in the preceding section, the routing protocol must carry the subnet mask with the corresponding network entries if you want to take advantage of route summarization. Otherwise, if you had more than one subnet mask applied to a class network number, the router wouldn't know which mask to use when routing a packet to a destination.

A good example of this problem is apparent in classful protocols, such as RIPv1 and IGRP, and how you lay out your IP addresses in your network. With classful protocols, routing updates are sent out with only network entries: no subnet masks are included. The assumption is that the routers on other segments are connected to the same class network and thus know about the subnet mask.

If a network number crosses boundaries from one class network to another, the classful protocol will *automatically* summarize it to the class address network number (A, B, or C), as is shown in the top part of Figure 12-9. As you can see, the classful protocol advertises just the network number (172.16.0.0) without any subnet mask. Plus, since the network number crosses a class boundary (172.16.0.0 to 192.168.1.0), the subnet (172.16.1.0) is not advertised, but instead the class address (172.16.0.0) is.

The bottom part of Figure 12-9 shows how classless protocols react (either by default or with configuration) when crossing a class boundary. Notice two things: the subnet mask is included in the routing update, and the routing update is *not* automatically summarized across the class boundary.

Given the routing behavior of classful routing protocols, certain addressing designs will create problems. Let's use the network shown in the top part of Figure 12-10. With a classful protocol, like RIPv1, the routers, when advertising networks across a class boundary, summarize them back to their class boundary. In this example, both RouterA and RouterB advertise 172.16.0.0—they don't advertise their specific subnets for 172.16.0.0. This creates a problem with RouterC, which receives two

FIGURE 12-9 Classful versus classless protocols and routing updates**FIGURE 12-10** Discontiguous subnets

routes for 172.16.0.0. If RouterC wanted to reach 172.16.1.0/24, it really wouldn't know which router (RouterA or RouterB) to send its packets to.

This network design is referred to as a *discontiguous* subnet design—not all of the subnets are connected together. In this network, 172.16.1.0/24 and 172.16.2.0/24 are not connected via another 172.16.0.0 subnet number. This creates routing problems for other routers not connected to the 172.16.0.0 network, and therefore, discontiguous subnet designs are not recommended with classful protocols.

Discontiguous subnets *are*, however, supported by *classless* protocols. As is shown in the bottom part of Figure 12-10, classless protocols include the subnet mask in the routing update. In this example, RouterC knows exactly where 172.16.1.0/24 and 172.16.2.0/24 are located, since the mask is included in the routing updates. However, discontiguous subnets are not recommended even with classless protocols, since they limit your ability to summarize routing information in the most efficient fashion.

exam

Watch

Discontiguous subnets are not supported by classful protocols but are supported by classless protocols.

Classful protocols do not include the subnet mask when advertising network and subnet numbers.

The Routing Table

When implementing route summarization, another thing you'll need to consider is that routing decisions, by a router, must be made on the entire destination IP address in the IP packet header. The router always uses the longest matching prefix in the routing table to perform its routing decision. Let's use the following simplified routing table to illustrate the router's decision-making process:

1. 172.16.17.66/32
2. 172.16.17.64/27
3. 172.16.17.0/24
4. 172.16.0.0/16
5. 0.0.0.0/0

A router receives an inbound packet on one of its interfaces and examines the destination IP address in the packet header: 172.16.17.65. The router then needs to

examine its routing table and find the best match for this packet and then route the packet out the corresponding interface to reach the destination. The router will basically sort the entries in the routing table from the most bits in a mask to the least number of bits.

In the preceding routing table, entry 1 isn't a valid match since the mask for the entry indicates a *host* address (32 bits). When comparing all 32 bits of 172.16.17.66 with 172.16.17.65, there isn't a match. Typically, host address routes are placed in the routing table whenever you have moved a host from its native network segment to another but, for logistical purposes, cannot change the address on the device to correspond to its new segment. In other words, you need this device to retain its old IP address.

When comparing entry 2 in the routing table, the router is comparing the first 27 bits of 172.16.17.64 with the first 27 bits of 172.16.17.65, which *do* match. When comparing entry 3, the router compares the first 24 bits of 172.16.17.0 with the first 24 bits of 172.16.17.64, which *also* match. When comparing entry 4, the router compares the first 16 bits of 172.16.0.0 with the first 16 bits of 172.16.17.65, which also match.

When comparing entry 5, the router finds that the entry is a default route and matches any packet. Given this example, the first entry doesn't match, but the last four do match. The router needs to pick one entry and use it. When picking an entry, it uses the one that best matches—the one with the longest number of

matching bits. Therefore, the router will use entry 2 to route this packet to the corresponding destination.

exam

Watch

Remember how to pick out the correct entry in the routing table given a destination IP address.

Performing Summarization

As was mentioned earlier, in order to summarize routing entries, they must have the same highest-order matching bits. In other words, you can perform summarization when the network numbers in question are a power of 2 and start on a network boundary that is either a power of 2 or a multiple of a power of 2. The network boundary is based on the subnet mask. If you have a subnet mask of 255.255.255.240, you cannot start the summarization on a network number that is not a multiple of 16 (the number of address accommodated by a mask of 240).

For instance, 192.168.1.16/28 is a valid summarization for this mask, while 192.168.1.8 is not (doesn't start on a multiple of 16). If the increment is not a power of 2 or a multiple of a power of 2, you can sometimes take the addresses and summarize them

into a set of smaller summarized routes. Here is a list of power-of-2 numbers: 0, 2, 4, 8, 16, 32, 64, and 128. Also, when performing summarization, you want to make sure that *all* of the routes that are aggregated are associated with the router (or behind the router) that is advertising the summarized route. The rest of this section talks about the basics of summarizing routes.

When summarizing, remember that you can summarize routes only on a bit boundary (power of 2), or a multiple of a power-of-2 boundary. The trick to summarization is to look at your subnet mask options: 0, 128, 192, 224, 240, 248, 252, 254, and 255. Each of these masks cover a range of numbers, as is shown in Table 12-1. For instance, say you have a set of Class C subnets: 192.168.1.0/30 and 192.168.1.4/30. These networks contain a total of eight addresses and start on a power-of-2 boundary: 0. Therefore, you could summarize these as 192.168.1.0/29, which encompasses addresses from 192.168.1.0 through 192.168.1.7.

Let's take a look at another example. Say you have a set of Class C subnets: 192.168.1.64/26 and 192.168.1.128/26. Each of these networks has 64 addresses, totaling 128 addresses. A mask value that accommodates 128 addresses in a Class C network is 255.255.255.128 (25 bits). However, this subnet mask poses a problem, since the bit value must be a power of 2 *and* start on a power-of-2 network boundary. With a 25-bit mask, there are only two network numbers: 192.168.1.0/25 and 192.168.1.128/25. The address 192.168.1.64/26 falls under the first network number, and 192.168.1.128/26 falls under the second one—so even though the two networks are contiguous, they can't be summarized with a 25-bit mask. You could use a 24-bit mask (255.255.255.0); however, this includes a total of 256 addresses, not just the 128

TABLE 12-1

Summarizing
Network
Numbers

Mask Value	Range of Numbers	Number of Bits
0	256 numbers	0
128	128 numbers	1
192	64 numbers	2
224	32 numbers	3
240	16 numbers	4
248	8 numbers	5
252	4 numbers	6
254	2 numbers	7
255	1 number	8

addresses in question. And as was mentioned in the first paragraph, you should summarize only for addresses that are connected to or behind your router. If 192.168.1.0/26 and 192.168.1.192/26 were also behind your router, you could summarize all four of these as 192.168.1.0/24.

These first two summarization examples were pretty simple. Let's look at a more complicated example to illustrate how difficult summarization can be if you don't lay out your addressing correctly in your network. I'll use the network shown in Figure 12-11. In this network, Router A needs to summarize routes that it and Routers B, C, and D are connected to, realizing, though, that there are other networks to the left of Router A. The goal is to have Router A advertise the least number of routes to routers to the networking cloud to the left.

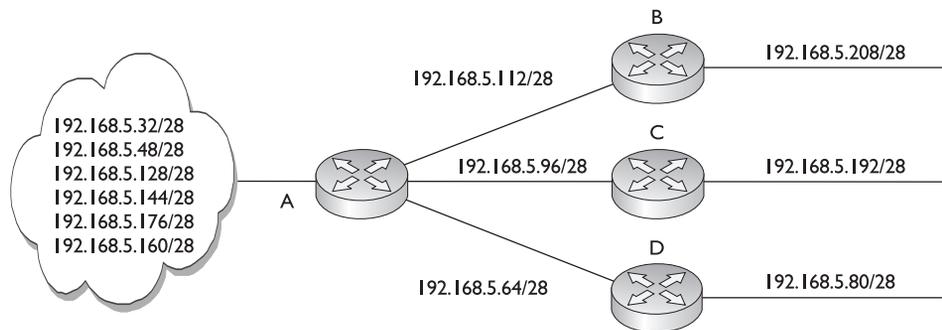
First, remember that Router A should create summarizations only for the routes either that it is connected to or that are behind it (it is connected to Routers B, C, and D). Second, these summarizations should either be a power of 2 or start on a power-of-2 networking boundary. In this example, the first thing you want to do is put the routes that Router A knows about (to its right) in numerical order:

- 192.168.5.64/28
- 192.168.5.80/28
- 192.168.5.96/28
- 192.168.5.112/28
- 192.168.5.192/28
- 192.168.5.208/28

Note that there are other subnets of 192.168.5.0 to the *left* of Router A that should not be included in the summarization. In this example, subnets 64, 80, 96,

FIGURE 12-11

Complex route summarization example



and 112 are contiguous, and if you use a 26-bit summarization mask, this would accommodate addresses from 64 through 127. These addresses are contiguous, and the summarization mask starts on a power-of-2 network boundary (address 64). In order to summarize subnets 192 and 208, you would need a 27-bit mask (255.255.255.224), which would include a block of 32 addresses: from 192 through 223.

Router A can advertise the following summarized routes to the left network cloud:

- **192.168.5.64/26** This covers addresses 64–127, which are to the right of Router A.
- **192.168.5.192/27** This covers addresses 192–223, which are also to the right of Router A.

exam
Watch *Understand the summarization examples in this section and practice summarization in the following exercise.*

As you can see, the number of network entries Router A originally advertised was six network numbers. Through summarization, this was reduced to two summarized routes.

The CCNA exam covers only the concepts of VLSM and summarization. You will not be required to configure routing protocols, such as EIGRP, OSPF, and RIPv2, to use summarization.

The BSCI exam for the CCNP and CCDP certifications, though, does cover this topic thoroughly.

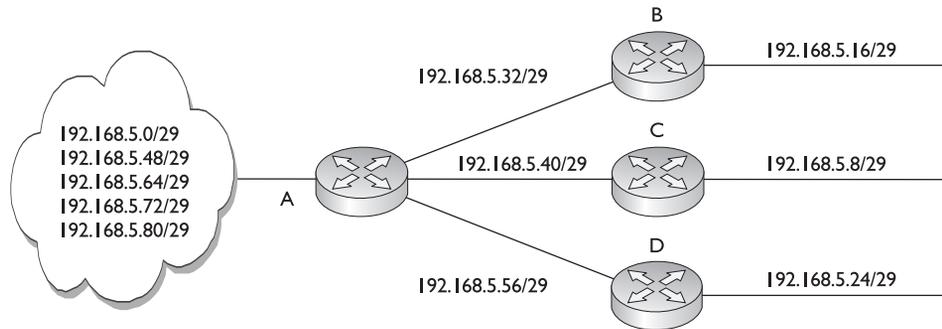
EXERCISE 12-1



Performing Route Summarization

The preceding few sections dealt with route summarization, its advantages and disadvantages. This exercise will help you reinforce this material by looking at an example network and come up with summarized routes for a router. You'll use the network shown in Figure 12-12. In this example, you need to summarize the routes to the right of Router A, making sure that these summarizations don't overlap any of the addresses in the network to the left of Router A.

1. Write down your networks (to the right of Router A) in numerical order:
Here are the networks that you want to summarize: 192.168.5.8/29, 192.168.5.16/29, 192.168.5.24/29, 192.168.5.32/29, 192.168.5.40/29, and 192.168.5.56/29.

FIGURE 12-12Summarization
exercise

- Break up the networks into contiguous blocks of addresses, starting on a power-of-2 network boundary

Given that the subnet mask is 255.255.255.248 (29 bits), here are the blocks of addresses:

- 192.168.5.8/29
- 192.168.5.16/29, 192.168.5.24/29
- 192.168.5.32/29, 192.168.5.40/29
- 192.168.5.56/29

Notice that even though subnets 8 and 16 are contiguous, a summarized mask would have to include subnet 0, which is to the left of Router A. Remember that the summarization must begin on a power-of-2 boundary and must correspond to valid network numbers for this mask value.

- Assign an appropriate *summarized* mask to each of these contiguous blocks:

For the given subnets, here is a list of the ones that can and can't be summarized, as well as the summarized masks.

- 192.168.5.8/29 Can't be summarized
- 192.168.5.16/29, 192.168.5.24/29 Can be summarized: 192.168.5.16/28
- 192.168.5.32/29, 192.168.5.40/29 Can be summarized: 192.168.5.32/28
- 192.168.5.56/29 Can't be summarized

The subnet 192.168.5.8/29 can't be summarized, since, when you shift one bit to the left in the subnet mask, this would include the network 192.168.5.0/29, which is to the left of Router A. Remember that on a 28-bit mask, the networks

increase in multiples of 16, starting at 0: 0, 16, 32, 48, 64, and so on. This is also true with 192.168.5.56. Shifting one bit to the left in the summarization would require the summarized route to start at 192.168.5.48.

CERTIFICATION SUMMARY

VLSM allows you to have more than one subnet mask applied to the same class address. VLSM's advantages include more efficient use of addressing and route summarization. Only classless protocols such as RIPv2, EIGRP, OSPF, IS-IS, and BGP support VLSM.

In order to perform VLSM, find the segment with the largest number of devices. Find an appropriate mask for the segment and write down all of your network numbers using this mask. Take one of these network numbers and apply a different mask to it to create more, yet smaller, subnets.

Route summarization is the ability to take a group of contiguous entries in your routing table and advertise these entries as a single summarized entry. Through proper configuration of summarization, your routing table sizes will decrease, the number of advertised network numbers will decrease, and you'll be able to contain certain networking problems, especially flapping routes. CIDR is a special type of route summarization. VLSM allows you to summarize back only to the class boundary of the network: the Class A, B, or C network number. CIDR allows you to summarize a group of contiguous class network numbers.

Summarization can be achieved only by laying out hierarchical addresses in your network. Used with a proper address design, hierarchical addressing allows for more efficient routing: it decreases routing table sizes, the amount of memory for routing, the number of processing cycles required, and difficulties with troubleshooting.

When implementing route summarization, note that the routing protocol must carry the subnet mask along with the routing entry: Only classless protocols allow this process. And since the mask is carried with the network number, discontinuous subnets are supported with classless protocols, but not classful ones. When the router makes routing decisions, it will use the entire destination IP address to make them.

When creating summarized entries, note that the network numbers being summarized must have the same highest-order matching bits. Remember that you can summarize routes only on a bit boundary, which is a power of 2, or a multiple of a power of 2. When summarizing, you need to know the ranges of addresses a mask value in an octet covers; for example, a subnet mask value of 192 covers 64 numbers.



TWO-MINUTE DRILL

Variable-Length Subnet Masking

- VLSM allows you to have different subnet masks applied to the same class address.
- Classless protocols, such as BGP, IS-IS, OSPF, and RIPv2, support VLSM.
- VLSM uses addressing more efficiently and allows you to configure route summarization.
- When setting up a network with VLSM, first find the largest segment. Then find an appropriate subnet mask for this network. Write down the subnet numbers according to this mask. For smaller segments, take one of the subnets and subnet it further, writing down your newly subnetted subnets.

Route Summarization

- Route summarization is the ability to take a bunch of contiguous network numbers in your routing table and advertise these contiguous routes as a single summarized route. The summarization must begin on a power-of-2 boundary based on the subnet mask value.
- Summarization reduces the routing table size, reduces the bandwidth required for routing updates, and contains network problems. Proper summarization requires a hierarchical addressing design in your network.
- CIDR, commonly called supernetting, allows you to summarize routes to the left of the class boundary, such as a group of Class C networks.
- Routing protocols must carry the subnet mask with the network entry to perform route summarization. Routing decisions must be made on the entire destination IP address. Summarization requires that the routing entries have the same highest-order matching bits.
- Classful protocols have problems with discontinuous subnet masks; classless protocols don't.

SELF TEST

The following Self Test questions will help you measure your understanding of the material presented in this chapter. Read all the choices carefully, as there may be more than one correct answer. Choose all correct answers for each question.

Variable-Length Subnet Masking

1. VLSM allows you to summarize 192.168.2.0/24 and 192.168.3.0/24 as 192.168.2.0/25.
 - A. True
 - B. False
2. VLSM allows a network segment to have more than one subnet mask.
 - A. True
 - B. False
3. Which protocol supports VLSM?
 - A. RIP
 - B. IGRP
 - C. RIP and IGRP
 - D. None of these
4. You are given a Class C network, 192.168.1.0/24. You need one network with 120 hosts and two networks with 60 hosts. How many subnet masks do you need?
 - A. 1
 - B. 2
 - C. 3
 - D. 4
5. You are given a class C network, 192.168.1.0/24. You need one network with 120 hosts and three networks with 60 hosts. What subnet mask values would you use?
 - A. 255.255.255.128 and 255.255.255.192
 - B. 255.255.255.128
 - C. 255.255.255.192
 - D. None of these

6. You are given a class C network, 192.168.1.0/24. You need three networks with 60 hosts and two networks with 30 hosts. What are the subnet mask values you could use?
 - A. 255.255.255.128 and 255.255.255.192
 - B. 255.255.255.224 and 255.255.255.240
 - C. 255.255.255.192 and 255.255.255.224
 - D. None of these

7. You are given this address space: 172.16.5.0/25. You need one network with 64 hosts and two with 30 hosts. What are the most specific subnet mask values to use?
 - A. /25 and /26
 - B. /26 and /27
 - C. /27 and /28
 - D. None of these

8. You are given a class C network and you have four LAN segments with the following numbers of devices: 120, 60, and two with 30. What subnet mask values would you use to accommodate these segments?
 - A. /24, /25, and /26
 - B. /25, /26, and /27
 - C. /26, /27, and /28
 - D. None of these

Route Summarization

9. VLSM allows you to summarize _____ back to the class boundary.
 - A. Subnets
 - B. Networks

10. Which of the following is not an advantage of route summarization?
 - A. It requires less memory and processing.
 - B. It supports smaller routing update sizes.
 - C. It contains network problems.
 - D. It supports discontinuous subnets.

24 Chapter 12: Advanced IP Addressing

11. _____ allows you to create this summarization: 10.0.0.0/7.
 - A. Subnetting
 - B. CDR
 - C. Supernetting
 - D. VLSM
12. Which of the following are classless protocols?
 - A. IGRP
 - B. EIGRP
 - C. IGRP and EIGRP
 - D. Neither IGRP or EIGRP
13. A routing protocol that supports route summarization must perform all except which of the following?
 - A. Carry the subnet mask with the network entry.
 - B. Make routing decisions based on the entire destination IP address.
 - C. Summarize entries so that the same lowest-order bits match.
 - D. None of these.
14. You have the following two routes: 192.168.1.64/27 and 192.168.1.96/27. Enter the most specific summarized route for these two subnets: _____.
15. You have the following four routes: 192.168.1.32/30, 192.168.1.36/30, 192.168.1.40/30, and 192.168.1.44/30. Enter the most specific summarized route for these four subnets: _____.

SELF TEST ANSWERS

Variable-Length Subnet Masking

1. **B.** CIDR allows you to summarize class networks together; VLSM allows you to summarize subnets only back to the class network boundary.
2. **B.** Each segment has a single network number and mask. VLSM allows a class address, not a network segment, to have more than one subnet mask.
3. **A.** RIPv2 supports VLSM (RIPv1 doesn't).
 B is classful and doesn't support VLSM. **C** includes a classful protocol. There is a correct answer, so **D** is incorrect.
4. **B.** You need two subnet masks: 255.255.255.128 (/25) and 255.255.255.192 (/26). This creates three networks, for instance, 192.168.1.0/25, 192.168.128/26, and 192.168.1.192/26.
5. **D.** This is impossible with a single Class C network. One hundred twenty hosts require a 255.255.255.128 mask, which is half a Class C network. Sixty hosts require a 255.255.255.192 mask; however, you need three of these, which is 3/4 of a Class C network. Therefore it is impossible.
 A is incorrect because it accommodates only the 120-host and two 60-host segments. **B** is incorrect because it accommodates only two subnets. **C** is incorrect because it accommodates the three 60-host segments, but not the 120-host segment.
6. **A** and **C.** **A** creates one 126-host segment and two 62-host segments. **C** creates three 62-host segments and two 30-host segments.
 B's second mask supports only 14 hosts. **D** is incorrect because there is a correct answer.
7. **D.** Sixty-four hosts require a 25-bit mask, and you are only given this to begin with—62 hosts would work with a 26-bit mask.
 A, B, and **C** don't support enough addresses.
8. **B.** A bit mask of 25 creates two networks: 0 and 128. If you take one of these subnets and apply a 26-bit mask, you have two more networks, such as 128 and 192. Taking one of these two subnets, applying a 27-bit mask creates two more subnets, such as 192 and 224.
 A, C, and **D** don't support enough addresses to accommodate all four LAN segments.

Route Summarization

9. **A.** VLSM allows you to summarize subnets back to the class A, B, or C network boundary.
 B is a non-subnetted address space and therefore is a Class A, B or C network number and can't be summarized with VLSM, but can be with CIDR.
10. **D.** Discontiguous subnets are supported by classless protocols, but they are not an advantage of summarization. Actually, summarization is more difficult if you have discontiguous subnets.
 A, B, and C are advantages of route summarization.
11. **C.** Supernetting, or CIDR, supports summarization of contiguous blocks of class A, B, or C networks.
 A is the opposite of summarization. **B** should be CIDR, not CDR. **D** allows you to summarize subnets, not networks.
12. **B.** EIGRP, as well as IS-IS, BGP, OSPF, and RIPv2, is a classless protocol.
 A is not a classless protocol. **C** includes a classful protocol (IGRP). **D** is incorrect because there is a correct answer.
13. **C.** Summarized entries must have the same *highest-order* matching bits, not lowest.
 A and B are things a routing protocol supporting route summarization must perform. **D** is incorrect because there is a correct answer.
14. 192.168.1.64/26: this includes addresses from 192.168.1.64 through 192.168.1.127.
15. 192.168.1.32/28: this includes addresses from 192.168.1.32 through 192.168.1.47.