



Internet Protocol

This chapter focuses on a number of objectives falling under the CCNP routing principles. Understanding basic Internet Protocol (IP) networking not only applies to the CCNP certification but all Cisco-based certification. A concrete understanding of how IP is used in today's networking environments is one of the most important tools to have before taking on the more advanced chapters in this guide.

This chapter starts by covering basic IP concepts. It then briefly explains how to efficiently configure IP to ensure full use of address space. Next, this chapter covers when and how IP routing tables can be minimized using summarization techniques with various routing protocols.

Five practical scenarios complete your understanding of these topics and ensure you have all the basic IP networking knowledge to complement your knowledge of today's most widely used networking protocol, IP.

Basic Internet Protocol

IP is a term widely used in today's networking world to describe a Network layer protocol that logically defines a distinct host or end systems such as a PC or router with an IP address.

An IP address is configured on end systems to allow communication between hosts that are geographically dispersed. An IP address is 32 bits in length with the network mask or subnet mask (also 32 bits in length) defining the host and subnet portion. A subnet is a network that you, as network administrator, segment to allow a hierarchical routing topology. Routing allows communication between these subnets. The host address is a logical unique address that resides on a subnet.

The Internet Engineering Task Force (IETF) standards body, which is a task force consisting of over 80 working groups responsible for developing Internet standards, defined five classes of addresses and the appropriate address ranges. Table 1-1 displays the five ranges.

Table 1-1 *Class A, B, C, D, and E Ranges*

Class of Address	Starting Bit Pattern	Range	Default Subnet Mask
Class A	0	1–126, 127*	255.0.0.0
Class B	10	128–191	255.255.0.0
Class C	110	192–223	255.255.255.0
Class D	1110	224–239	255.255.255.240
Class E	1111	240–255	Reserved

* 127.0.0.0 is reserved for loopbacks purposes. Other reserved addresses for private use as defined by RFC 1918 are

10.0.0.0-10.255.255.255

172.16.0.0-172.16.255.255

192.168.0.0-192.168.255.255

Soon after these ranges were defined and the Internet's popularity extended beyond the Department of Defense in the United States, it became clear that to ensure that a larger community could connect to the World Wide Web there had to be a way to extend IP address space by using subnetting. Subnetting allows an administrator to extend the boundary for any given subnet.

To best illustrate an IP address and subnet portion, determine how many hosts are available on a particular subnet, or even how to best utilize an IP address space, consider the following example.

You are given the IP address 131.108.1.56 and the subnet mask is 255.255.255.0. This example helps you determine what the subnet is, how many hosts can reside on this subnet, and what the broadcast address is.

You can deduce the subnet for any IP address by performing a logical AND operation along with the subnet mask.

NOTE

A logical AND operation follows two basic rules. One is that positive and positive equal positive, and the second is that negative and positive or negative is negative. So, in binary (positive is 1 and negative is 0), 0 AND 0 is 0, 0 AND 1 is 0, 1 AND 1 is 1, 1 AND 0 is 0, and so forth.

Figure 1-1 displays the logical AND operation used to determine the subnet address.

Figure 1-1 *AND Logic Operation*

IP Address (131.108.1.56)	10000011.01101100.00000001.00111000
IP Subnet MASK (255.255.255.0)	<u>11111111.11111111.11111111.00000000</u>
Logical AND	10000011.01101100.00000001.00000000
In Decimal	131 108 1 0

The result of the logical AND operation reveals the subnet address is 131.108.1.0. The subnet address is reserved and cannot be assigned to end devices.

To determine the number of hosts available in any given subnet, you simply apply the formula $2^n - 2$ where n is the number of borrowed bits. This is best explained with examples. To determine the number of borrowed bits, you must examine the subnet mask in binary. For a default Class C network mask of 255.255.255.0, the last eight bits represent the borrowed bits. So, for a Class C network, the number of hosts that can reside are $2^8 - 2 = 256 - 2 = 254$ hosts. (You subtract two host addresses for the subnet address and the broadcast address, which are not permitted to be used by host devices.) In IP, a broadcast address consists of all binary 1s, so for this example, the broadcast address for the subnet 131.108.1.0 is 131.108.1.255. (255 in binary is 11111111.)

Now consider another example. Given the host address 171.224.10.67 and the subnet mask of 255.255.255.224, this example shows you how to determine the subnet and the number of hosts that can reside on this network.

To determine the subnet, perform a logical AND. Figure 1-2 displays the operation.

Figure 1-2 *Logical AND Operation*

IP Address (171.224.10.67)	10101011.11100000.00001010.01000011
IP Subnet MASK (255.255.255.224)	<u>11111111.11111111.11111111.11100000</u>
Logical AND	10101011.11100000.00001010.01000000
In Decimal or subnet	171 224 10 64

The subnet is 171.224.10.64. The number of hosts that can reside on this network with a subnet mask of 255.255.255.224 (or 11100000, 5 borrow bits) is $2^5 - 2 = 32 - 2 = 30$ hosts. You can apply the technique used in this simple example to any Class A, B, or C address, and applying a subnet mask that is not the default or classful kind enables you to extend IP address space and allow a larger number of devices to connect to the IP network.

Table 1-2 displays some common subnets used in today's network and the number of hosts available on those subnets.

Table 1-2 *Common Subnets in Today's Networks*

Decimal	Subnets	Hosts
252 (1111 1100)	64 subnets	2 hosts*
248 (1111 1000)	32 subnets	6 hosts
240 (1111 0000)	16 subnets	14 hosts
224 (1110 0000)	8 subnets	30 hosts
192 (1100 0000)	4 subnets	62 hosts
128 (1000 0000)	2 subnets	126 hosts
64 (0100 0000)		

* Used commonly for WAN circuits when no more than 2 hosts reside.

Variable-Length Subnet Masks (VLSM)

A variable-length subnet mask (VLSM) is designed to allow more efficient use of IP address space by borrowing bits from the subnet mask and allocating them to host devices. To allow a greater number of devices to connect to the Internet and intranets, the standards body of various routing protocols designed an IP routing algorithm to cater to IP networks with a different subnet mask than the default used in classful networks.

NOTE The following routing algorithms support VLSM: RIP Version 2, OSPF, IS-IS, EIGRP, and BGP4.

To demonstrate the use of VLSM, consider the example of connecting two Cisco routers through a wide-area link. Only two devices host systems are needed.

To use any IP address space effectively, it would be wise to use the lowest possible number of subnet bits and lowest possible number of host bits. You could use a Class C mask or a mask that allows for 254 hosts. For a link that never uses more than two hosts, this wastes a vast amount of space, 252 addresses in fact.

Apply the formula to determine the best subnet to use to cater to two hosts on any given subnet and class of address. Remember that you must subtract two host addresses for the subnet address and broadcast address.

Applying the formula, you get $2^n - 2 = 2$, or $2^n = 4$, or $n = 2$ borrowed bits. You need to borrow only two bits from the subnet mask to allow for two host addresses. The subnet mask is 30 bits in length or 255.255.255.252 in binary, which is represented as 11111111.11111111.11111111.11111100. The last two bits (00) are available for host addresses; the subnet is 00; the first host address is 01, the second is 10, and the broadcast address is 11.

NOTE

Loopback interfaces configured on Cisco routers are typically configured with a host address using a 32-bit subnet mask, which allows, for example, a Class C network with 255 hosts among 255 different routers and conserves valuable IP address space.

Summarization and How to Configure Summarization

Summarization, put simply, enables a given routing protocol to minimize IP routing tables by taking steps to advertise a smaller or lesser IP route destination for a large set of subnets or networks. IP routing entries consume bandwidth of expensive links between different geographic locations, take CPU cycles on routers, and, most importantly, require memory.

To give network designers the ability to manage large networks, summarization is important for limiting or reducing IP routing tables. The most important consideration to make when summarizing any IP address space is to ensure a hierarchical design.

In a hierarchical design, IP address space is configured across any given router so that it can be easily summarized. To illustrate the capabilities of summarization consider the following IP address ranges in Table 1-3.

Table 1-3 *IP Address Range*

IP Subnet	Binary Last Third Octet
131.108.1.0/24	0000 0001
131.108.2.0/24	0000 0010
131.108.3.0/24	0000 0011
131.108.4.0/24	0000 0100
131.108.5.0/24	0000 0101
131.108.6.0/24	0000 0110
131.107.7.0/24	0000 0111

A router would normally advertise each of the seven IP address ranges, from 131.108.1–7, as seven different IP route entries.

The binary examination of the subnets 1 to 7 in Table 1-3 displays that the first five bits (shaded) are unchanged. The most important fact is that these seven networks are contiguous or in a range that you can easily summarize. Because the high-order bits are common in Table 1-3 (0000 0) and all seven routes are contiguous (binary 001 to 111), you can perform summarization. Because the first five bits are the same, you can apply the mask 248 (11111 000) on the third octet and send an advertisement encompassing all seven routes. Before looking at how to complete this summarization using RIP, EIGRP, or OSPF, the following is a list of benefits when using summarization:

- Reduces routing table sizes
- Allows for network growth
- Simplifies routing algorithm recalculation when changes occur
- Reduces requirements for memory and CPU usage on routers significantly

The alternatives to network summarization are not easy to accomplish, and this includes renumbering an IP network or using secondary addressing on Cisco routers, which is not an ideal solution for management purposes and also provides extra overhead on a router. Also, it is important to understand that if a range of addresses is not contiguous (that is, they do not start from a range that can be easily summarized, such as the range of addresses 131.108.1.0/24 and 131.108.10.0/24), summarization is impossible. You could still summarize the first seven networks, for example, but they might reside in other parts of your network and cause IP routing problems. The best practice is to assign a group of addresses to a geographic area so that the distribution layer of any network enables summarization to be relatively easy to complete.

Depending on the routing protocols in use, summarization may be enabled by default. Automatic summarization simply announces a Class A network with an 8-bit mask, 255.0.0.0, Class B with 16-bit mask, and a Class C mask with a 24-bit mask, 255.255.255.0. With RIPv2, automatic summarization occurs. In other words, you must disable automatic summarization to allow the more specific routes to be advertised; otherwise a default mask is assumed.

To disable automatic summaries with RIPv2, use the following command:

```
router rip
version 2
no auto-summary
```

The command **no auto-summary** disables automatic summaries and allows subnets to be advertised.

EIGRP also applies automatic summaries but it also enables the manual configuration of summary addresses. The following example shows you how to summarize the networks in Table 1-3 using EIGRP.

To configure summarization with EIGRP, you must first disable automatic summarization with the following command:

```
router eigrp 1
no auto-summary
```

Then, you apply the manual summarization on the interface to which you want to send the advertised summary. Example 1-1 displays the command you use to summarize the seven networks in Table 1-3.

Example 1-1 *Summary with EIGRP*

```
interface serial 0
ip summary-address eigrp 1 131.108.1.0 255.255.248.0
```

Example 1-1 applies a summary on the serial interface. Also note that the EIGRP autonomous system number is 1, matching the configuration on the router because you can have more than one EIGRP process running. The actual summary is 131.108.1.0 255.255.248.0, which replaces the seven individual routers numbered 131.108.1-7.0/24 with one simple route.

OSPF allows summarization manually under the OSPF process ID. Now look at how to configure the seven networks in Table 1-3 with an OSPF summary. You use the following command in OSPF to summarize internal OSPF routes:

```
area area-id range address mask
```

Example 1-2 displays the configuration required to summarize the seven networks in Table 1-3. Assume the *area-id* for now is 1.

NOTE With OSPF, you can correctly configure summarization only on area border routers (ABRs). An ABR resides in more than one OSPF area. For this example, assume the Cisco router is an ABR.

Example 1-2 *OSPF summary*

```
router ospf 1
area 1 range 131.108.1.0 255.255.248.0
```

NOTE OSPF also enables you to summarize external OSPF routes redistributed from such protocols as IGRP or RIP. BGP and IS-IS, covered in Chapters 4, “Advanced OSPF and Integrated Intermediate System-to-Intermediate System,” 6, “Basic Border Gateway Protocol” and 7, “Advanced BGP,” also provide complex summarization techniques.

IP Helper Address

As in any network, broadcasts are used to find and discover end systems. In a Layer 2 environment, you use broadcasts to find an end system's MAC address. Layer 3 of the TCP/IP model, IP also uses broadcasts for such services as sending IP datagrams to all hosts on a particular network. Broadcasts on any network consume CPU and bandwidth to reduce this even more. In an IP network, you use the IP helper address to change a broadcast into a more specific destination address so not all devices must view the IP data, which conserves bandwidth.

To save on bandwidth, all Cisco routers installed with Cisco Internet Operating System (IOS) software by default have an algorithm that dictates that not all broadcast packets be forwarded. So to allow the ability to forward packets wisely, you can use the IP helper address command to convert a broadcast into a more specific destination address. The command to enable an IP help address is as follows:

```
ip helper-address address
```

You can configure more than one helper address per interface on a Cisco router. The IP helper address forwards packets that are normally discarded by default to the following services:

- Trivial File Transfer Protocol (TFTP)
- Domain Name System (DNS)
- BOOTP server
- BOOTP client
- NetBIOS Name Server
- Dynamic Host Configuration Protocol (DHCP)

NOTE The most common use for the helper address is for clients running DHCP, which remote servers assign IP addresses and subnet masks usually performed locally through a broadcast to be served remotely with a unicast (one) packet.

Scenarios

The following scenarios are designed to draw together some of the content described in this chapter and some of the content you have seen in your own networks or practice labs. There is no one right way to accomplish many of the tasks presented, and using good practice and defining your end goal are important in any real-life design or solution. The five scenarios


presented in this chapter are based on simple IP technologies to introduce you to the configuration of IP on Cisco routers and give you the basic foundation required to complete the more advanced topics and scenarios found later in this book. Readers who are familiar with these basics may want to skip this chapter and move on to Chapter 2, “Routing Principles.”

Scenario 1-1: Configuring a Cisco Router for IP

In this scenario, you see how to configure one Cisco router for IP routing using a Class B (/16) network 161.108.1.0 with a Class C subnet mask (255.255.255.0 or /24 mask).

Figure 1-3 displays the one router, named R1, with one Ethernet interface.

Figure 1-3 *IP Routing on Cisco Routers*



Example 1-3 displays the IP configuration performed on R1's Ethernet interface.

Example 1-3 *IP Configuration on R1*

```
R1(config)#int e 0/0
R1(config-if)#ip address 161.108.1.1 255.255.255.0
R1(config-if)#no shutdown
4w1d: %LINK-3-UPDOWN: Interface Ethernet0/0, changed state to up
4w1d: %LINEPROTO-5-UPDOWN: Line protocol on Interface Ethernet0/0,
      changed state to up
```

NOTE

When you enable the Ethernet interface with the command **[no] shutdown**, the IOS message tells you the Ethernet interface and the line protocol are up. To see these messages remotely, enable **terminal monitor** on any VTY lines. Also, by default, all Cisco routers are enabled for IP routing with the command **ip routing**. You can disable IP routing with the command **[no] ip routing**.

Example 1-4 displays the active Ethernet interface up and the current IP address configuration.

Example 1-4 *show interface ethernet e0/0 on R1*

```
R1#show interfaces ethernet 0/0
Ethernet0/0 is up, line protocol is up → Interface is up and active
  Hardware is AmdP2, address is 0001.9645.ff40 (bia 0001.9645.ff40)
  Internet address is 161.108.1.1/24 →configure IP address
  MTU 1500 bytes, BW 10000 Kbit, DLY 1000 usec, rely 255/255, load 1/255
  Encapsulation ARPA, loopback not set, keepalive set (10 sec)
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:21, output 00:00:02, output hang never
  Last clearing of "show interface" counters never
  Queueing strategy: fifo
  Output queue 0/40, 0 drops; input queue 0/75, 0 drops
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
    315871 packets input, 30894958 bytes, 0 no buffer
    Received 315628 broadcasts, 0 runts, 0 giants, 0 throttles
    0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
    0 input packets with dribble condition detected
  470705 packets output, 43588385 bytes, 0 underruns
    0 output errors, 3 collisions, 45 interface resets
    0 babbles, 0 late collision, 22 deferred
    0 lost carrier, 0 no carrier
    0 output buffer failures, 0 output buffers swapped out
```

Next, you see how to configure a secondary address on R1 using the IP address 131.108.1.1/24. Example 1-5 displays the secondary IP address assignment.

Example 1-5 *Secondary Address Configuration on R1*

```
R1(config)#interface ethernet 0/0
R1(config-if)#ip address 131.108.1.1 255.255.255.0 secondary
```

R1 now has two IP address assignments: 161.108.1.1/24 and 131.108.1.1/24. Confirm the IP address assignment by displaying the interface statistics with the command **show interfaces Ethernet 0/0**. Example 1-6 displays the Ethernet statistics on R1 and is truncated for clarity.

Example 1-6 *show interfaces ethernet 0/0*

```
R1#show interfaces ethernet 0/0
Ethernet0/0 is up, line protocol is up
  Hardware is AmdP2, address is 0001.9645.ff40 (bia 0001.9645.ff40)
  Internet address is 161.108.1.1/24
  MTU 1500 bytes, BW 10000 Kbit, DLY 1000 usec, rely 255/255, load 1/255
  Encapsulation ARPA, loopback not set, keepalive set (10 sec)
  ...truncated
```

Example 1-6 does not show the secondary addressing on R1. Unfortunately, the Cisco IOS does not display IP secondary addressing, and the only way to view any secondary addressing is to view the configuration. Example 1-7 displays the full working configuration on R1 along with the secondary IP address, 131.108.1.1.

Example 1-7 *Full working configuration on R1*

```
hostname R1
!
interface Ethernet0/0
 ip address 131.108.1.1 255.255.255.0 secondary
 ip address 161.108.1.1 255.255.255.0
!
interface Serial0/0
 shutdown
!
interface Serial0/1
 shutdown
!
line con 0
line aux 0
line vty 0 4
!
end
```

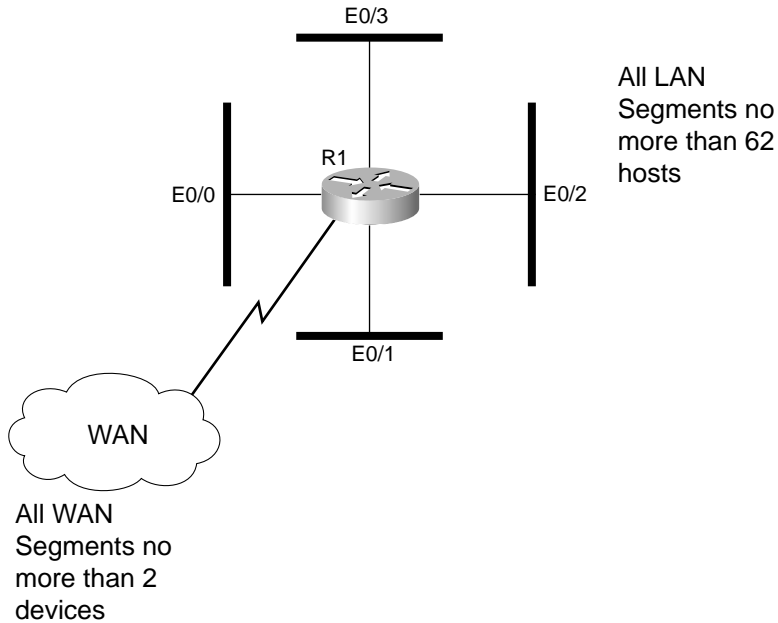
Scenario 1-2: Efficiently Configuring a Network for IP

Suppose you have been asked by a network architect to break up the Class B address 131.108.1.0/24 into four equal subnets that can be used to allow at most 62 hosts per subnet. In addition to this, you must use the address space 131.108.2.0/24 for all wide-area network (WAN) connections that use no more than two hosts per subnet. The network architect has also asked you to document all WAN addresses for future use.

Figure 1-4 displays the network topology graphically.

Start by breaking up the subnet 131.108.1.0/24 into four equal subnets. To do this, examine the subnet in binary. The last eight bits are used for host addresses, so by default you have 254 IP address available. To allow at most 62 hosts, you use the formula $2^n - 2 = 62$, which becomes $2^n = 64$. n , which is the borrowed amount of bits, becomes six bits. So to allow at most 62 hosts, you must use the subnet mask of 255.255.255.192, where 192 in binary is 11000000. The host devices use the last six bits. This is only half the job; you must also configure the four different subnets on R1 in Figure 1-4. To determine the four subnets you must count in binary.

Figure 1-4 IP Address Configuration Requirements



The first subnet starts from 131.108.1.0. You know the broadcast address ends in all 1s, so count from binary 0 to all 1s. Count only from the last octet. Table 1-4 displays the binary calculation.

Table 1-4 Binary Addition 1

Decimal	Binary	Comment
0	000000	Subnet (all zeros)
1	000001	First host address
2	000010	Second host address
3	000011	Third host address
...		
62	111110	Last host address
63	111111	Broadcast address (all 1s)

Table 1-4 counts in binary from 0 to 3 and so forth until 63, which in binary is 00111111. Notice that the last six bits are all 1s, which indicates the broadcast address, so the first

subnet ranges from 131.108.1.0 to 131.108.1.63. The subnet is 131.108.1.0, and the broadcast address is 131.108.1.63.

Table 1-5 performs the same calculation in binary without the intermediate steps to demonstrate the broadcast address for the second subnet.

Table 1-5 *Binary Addition Subnet 2*

Decimal	Binary	Comment
64	1000000	Subnet all zeros
65	1000001	First host address
66	1000010	Second host address
...		
126	1111110	Last host address
127	1111111	Host address

Table 1-5 displays the second subnet with all zeros as 131.108.1.64 and the broadcast of 131.108.1.127.

Table 1-6 displays the third subnet calculation starting from the next available decimal number of 128.

Table 1-6 *Binary Addition Subnet 3*

Decimal	Binary	Comment
128	10000000	Subnet (all zeros)
129	10000001	First host address
130	10000010	Second host address
131	10000011	Third host address
...		
190	10111110	Last host address
191	10111111	Broadcast address (all 1s)

Table 1-6 displays the subnet as 131.108.1.128, and the broadcast address as 131.108.1.191.

Finally, you can deduce the last subnet available in exactly the same way. Table 1-7 displays the final binary addition.

Table 1-7 *Binary Addition Subnet 4*

Decimal	Binary	Comment
192	11000000	Subnet (all zeros)
193	11000001	First host address
194	11000010	Second host address
195	11000011	Third host address
...		
253	11111110	Last host address
255	11111111	Broadcast address (all 1s)

NOTE

If you are confused about how to convert binary from decimal, simply use a Windows-based calculator to perform the calculation to assist in your first few calculations. It is vital that you can perform these steps without much thought, so you can quickly break up any type of subnet in various design situations or examination scenarios.

Table 1-7 displays the subnet as 131.108.1.192 and the broadcast address for the final subnet as 131.108.1.255.

Now that you have the four broken subnets, configure the Router R1 in Figure 1-4 for IP routing. Example 1-8 displays the IP configuration on the four interfaces on R1.

Example 1-8 *IP Configuration on R1 with Four Subnets*

```
R1(config)#interface ethernet 0/0
R1(config-if)#ip address 131.108.1.1 255.255.255.192
R1(config)#interface ethernet 0/1
R1(config-if)#ip address 131.108.1.65 255.255.255.192
R1(config)#interface ethernet 0/2
R1(config-if)#ip address 131.108.1.129 255.255.255.192
R1(config)#interface ethernet 0/3
R1(config-if)#ip address 131.108.1.193 255.255.255.192
```

The mask is 255.255.255.192 in Example 1-8. The mask or subnet mask is derived from the six bits you borrowed to extend the Class B address 131.108.1.0. Binary 1100000 is 192.

To complete this scenario, you have to break up the network 131.108.2.0/24 into 30-bit sized subnets so that they can be used on WAN circuits that contain no more than two hosts.

Once more, use the simple formula $2^n - 2 = 2$, or $2^n = 4$, where $n = 2$. So, you need two bits per subnet, and you have already discovered that the mask is 255.255.255.252.

Table 1-8 displays the first four subnets available along with the subnet, broadcast address, and binary equivalent.

Table 1-8 WAN Host Assignment

Decimal	Binary	Comment
131.108.2.0	00000000	First subnet, last two bits all zeros
131.108.2.1	00000001	First host
131.108.2.2	00000010	Second host
131.108.2.3	00000011	Broadcast address, last two bits all 1s
131.108.2.4	00000100	Second subnet, last two bits all zeros
131.108.2.5	00000101	First host
131.108.2.6	00000110	Second Host
131.108.2.7	00000111	Broadcast address, last two bits all 1s
131.108.2.8	00001000	First subnet, last two bits all zeros
131.108.2.9	00001001	First host
131.108.2.10	00001010	Second host
131.108.2.11	00001011	Broadcast address, last two bits all 1s
131.108.2.12	00001100	Second subnet, last two bits all zeros
131.108.2.13	00001101	First host
131.108.2.14	00001110	Second host
131.108.2.15	00001111	Broadcast address, last two bits all 1s

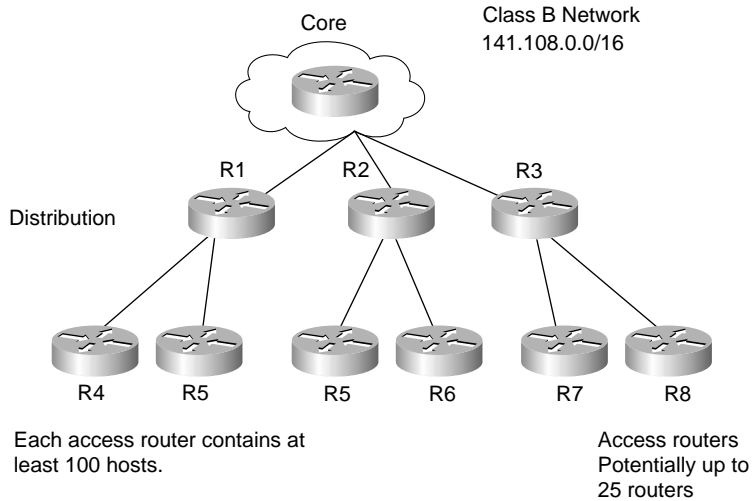
As an exercise, you can try to complete the table on your own. Simply count in binary and the next available subnet is clearly evident to you. Notice that the subnets in decimal count in fours, so the first subnet is 131.108.2.0/30, then 131.108.2.4/30, 131.108.2.8/30, 131.108.2.12/30, and so forth.

Scenario 1-3: Configuring IP VLSM for a Large Network

This scenario is slightly more complex. Figure 1-5 displays a network requiring a core network with a large number of routers (assume around 20), a distribution network with three routers, and an access network initially containing only six routers. The access network should have a potential for at most 25 routers (commonly known as access-level

routers) to be connected through the distribution routers. Figure 1-5 displays the core network surrounded by three distribution routers and the six access-level routers.

Figure 1-5 *VLSM in a Large Network*



The Class B address 141.108.0.0 has been assigned to you for this task. You should ensure this address space is designed so that company growth allows you to use IP address space wisely to conserve it. Ensure summarization is possible with the three distribution routers.

It is important that the IP addressing scheme is correctly laid out in a hierarchical fashion so that you can use summarization IP routing tables to keep them to a minimum. Start with the core of the network with a possible 20 routers. The core network of any large organization typically grows at a slower pace than access routers, so assume that allowing for over 1500 hosts should suffice. Assign seven Class C networks for the core, and reserve another eight for future use. Using 15 subnets allows for easy summarization as well. Assign the range 141.108.1.0–141.108.15.255 to the core network. In binary, this is the range 00000001 to 00001111, so the first four bits are common.

The distribution routers generally perform all the summarization, so you can assign another seven subnets and reserve another eight Class C networks for future use. So now the distribution routers use the range 141.108.16.0–141.108.31.255.

The access-level routers, where the users generally reside, typically grow at a fast rate, and in this scenario, each site has over 100 users; it is also possible that over 30 (90 in total) remote sites will be added in the future. It is vital that the subnets used here are contiguous

so that summarization can take place on the distribution Routers R1, R2, and R3. The following describes a sample solution:

- For access Routers R4 and R5 and possible new routers, use the range 141.108.32.0 to 141.100.63.255; in binary that ranges from 100000 (32) to 63(11111).
- For access Routers R6 and R7 and possible new routers, use the range 141.108.64.0 to 141.100.95.255; in binary that ranges from 1000000(64) to 1011111(95).
- For access Routers R8 and R9 and possible new routers, use the range 141.108.96.0 to 141.108.127.255; in binary that ranges from 1100000(96) to 1111111(127).
- You can reserve the remaining 128 subnets for future use.

This is by no means the only way you can accomplish the tasks in this scenario, but you need to apply these principles in any IP subnet addressing design.

NOTE

Cisco IOS gives you even more IP address space by allowing the use of subnet zero with the IOS command **ip subnet-zero**. Of course non-Cisco devices may not understand subnet zero. A good use for subnet zero would be for WAN links or loopback interfaces and conserving IP address space for real hosts, such as UNIX devices and user PCs. Subnet zero, for example, when using the Class B address 141.108.0.0 is 141.108.0.0, so a host address on a Cisco router could be 141.108.0.1/24.

When designing any IP network, you must answer the following core questions:

- How many subnets are available?
- What IP ranges will be used; will private address space be applied to conserve public addresses?
- How many hosts reside on the edge of the network?
- What are the expansion possibilities for the network?
- What are the geographic locations of remote sites?
- Is there a connection to the Internet or WWW?
- Is an IP address space currently being used?
- What are the current sizes of existing IP routing tables?
- Are any non-IP protocols already in use? If so, can you tunnel these non-IP protocols?
- What routing protocols enable the use of VLSM?
- These are just some of the major questions that you need to look at carefully. Cisco Systems provides a comprehensive guide to subnets at the following URL:

www.cisco.com/univercd/cc/td/doc/cisintwk/idg4/nd2003.htm

NOTE Great resources for information on IP addressing and subnet calculators are also available on the Internet.

Scenario 1-4: Summarization with EIGRP and OSPF

In this scenario, given the address ranges in Table 1-9, you see how to configure summarization with EIGRP and OSPF.

Table 1-9 displays the IP address ranges to be summarized, as well as the binary representation of the third octet or the subnet part of the IP address space.

Table 1-9 *IP Address Ranges*

IP Subnet	Subnet Mask	Binary Representation of Third Octet
151.100.1.0	255.255.255.0	00000001
151.100.2.0	255.255.255.0	00000010
151.100.3.0	255.255.255.0	00000011
151.100.4.0	255.255.255.0	00000100
151.100.5.0	255.255.255.0	00000101
151.100.6.0	255.255.255.0	00000110
151.100.7.0	255.255.255.0	00000111
151.100.8.0	255.255.255.0	00001000
151.100.9.0	255.255.255.0	00001001
151.100.10.0	255.255.255.0	00001010
151.100.11.0	255.255.255.0	00001011
151.100.12.0	255.255.255.0	00001100
151.100.13.0	255.255.255.0	00001101
151.100.14.0	255.255.255.0	00001110
151.100.15.0	255.255.255.0	00001111
151.100.16.0	255.255.255.0	00010000

Before configuring EIGRP or OSPF summarization, you first need to decide whether summarization is possible at all. Table 1-9 displays 16 subnets, numbered from 1-16. The first 15 subnets all have one thing in common when viewed in binary: The first four bits or high-order bits are always 0. Therefore, you can summarize the first 15 networks using the subnet mask 255.255.255.240. (240 in binary is **1111000** where the first four bits are common.) The last four bits contains the networks 1 to 15 or in binary encompass all networks from 0000 to 1111.

The last remaining subnet 151.100.16.0 is the odd network out. Although it is contiguous, you cannot summarize it along with the first 15 network, because any summary address range encompasses networks beyond 151.100.16.0, which may reside in other parts of the network.

Configure EIGRP to summarize these routes out of a serial port (serial 0/0 in this example). Example 1-9 displays the configuration required to disable automatic summarization and the two required summary address commands on the serial 0/0 on a router named R1.

Example 1-9 *EIGRP Summary*

```
R1(config)#router eigrp 1
R1(config-router)#no auto-summary
R1(config)#interface serial 0/0
R1(config-if)#ip summary-address eigrp 1 151.100.1.0 255.255.255.240
R1(config-if)#ip summary-address eigrp 1 151.100.16.0 255.255.255.0
```

In Example 1-9, the router R1 sends only two updates: one for the networks ranging from 151.100.1.0 to 151.100.15.0 and another for 151.100.16.0. These two are instead of 16 separate IP route entries. Even in a small scenario like this, you saved 14 IP route entries. Reducing IP routing tables means when a router performs a routing table search, the time it takes to determine the outbound interface is reduced allowing end-user data to be sent faster over a given medium.

With OSPF, you do not need to disable automatic summarization, because OSPF does not automatically summarize IP subnets. Hence, to summarize the same block of addresses of a router (OSPF ABR), you apply two commands under the OSPF process. Example 1-10 displays the summary commands required.

Example 1-10 *OSPF Summary*

```
R1(config)#router ospf 1
R1(config-router)#no area 1 range 151.100.16.0 255.255.255.240
R1(config-router)#area 1 range 151.100.16.0 255.255.255.0
```


Scenario 1-5: Configuring IP Helper Address

The following scenario demonstrates the powerful use of the helper command and how broadcast traffic, which is dropped by default on Cisco routers, can be forwarded in a manageable fashion and enable IP connectivity across a WAN.

In this scenario, you have a group of users on one segment requiring IP address assignment. No local servers reside on the segment with this group of users.

Figure 1-6 displays the network topology.

Figure 1-6 *IP Helper Requirement*



Now, when the users on the local-area network (LAN) segment attached to R1 send out a request for an IP address, this IP packet is sent to the broadcast address, which is dropped by default. Unless a local Dynamic Host Configuration Protocol (DHCP) server exists on this segment, the users' requests for an IP address aren't responded to. To alleviate this problem, you configure a helper address on R1 pointing to the remote file server(s)' address. In this case, two servers are available for redundancy, so you can configure two helper addresses on R1's Ethernet port.

NOTE

Remember, a helper address can forward many UDP-based protocols such as DNS and BOOTP requests. You can further restrict which protocols are sent by using the IOS command **ip forward-protocol {udp [port]}** or you can forward a packet based on a particular port number used by a certain application.

Example 1-11 displays the helper address configuration on R1.

Example 1-11 *IP Helper Address Configuration on R1*

```
R1(config)#interface ethernet 0/0
R1(config-if)#ip helper-address 131.108.1.99
R1(config-if)#ip helper-address 131.108.1.100
```

The five basic scenarios in this first chapter are aimed at addressing your basic knowledge or re-enforcing what you already know. The Practical Exercise that follows gives you an opportunity to test yourself on these concepts.

Practical Exercise: IP

NOTE

Practical Exercises are designed to test your knowledge of the topics covered in this chapter. The Practical Exercise begins by giving you some information about a situation and then asks you to work through the solution on your own. The solution can be found at the end.

Given the IP address ranges in Table 1-10 and using EIGRP as your routing algorithm, ensure that the least number of IP routing entries are sent out the Ethernet 0/0 port on a Cisco IOS-based router. Table 1-10 displays the IP subnet ranges.

Table 1-10 *IP Subnet Ranges*

IP Subnet	Subnet Mask	Binary Value of Third Octet
171.100.1.0	255.255.255.0	00000001
171.100.2.0	255.255.255.0	00000010
171.100.3.0	255.255.255.0	00000011
171.100.4.0	255.255.255.0	00000100
171.100.5.0	255.255.255.0	00000101
171.100.6.0	255.255.255.0	00000110
171.100.7.0	255.255.255.0	00000111

Practical Exercise Solution

You should notice that the first five bits are the same and the last three encompass the range 1-7, so you can apply the following summary command:

```
ip summary address eigrp 1 171.100.1.0 255.255.255.248
```

Example 1-12 displays the configuration required to summarize the networks from Table 1-10 on an Ethernet 0/0 port using the ? tool to demonstrate the available options required by Cisco IOS.

Example 1-12 *Sample Configuration*

```
R1(config)#interface ethernet 0/0
R1(config-if)#ip summary-address ?
    eigrp Enhanced Interior Gateway Routing Protocol (EIGRP)

R1(config-if)#ip summary-address eigrp 1 171.100.1.0 255.255.255.248
R1(config-if)#ip summary-address ?
    eigrp Enhanced Interior Gateway Routing Protocol (EIGRP)
R1(config-if)#ip summary-address eigrp ?
    <1-65535> Autonomous system number
R1(config-if)#ip summary-address eigrp 1 ?
    A.B.C.D IP address

R1(config-if)#ip summary-address eigrp 1 171.100.1.0 255.255.255.248
```

NOTE

Example 1-12 displays the Cisco IOS prompts that appear when the user enters the question mark (?) to display the options or parameters the Cisco IOS requires next. They are illustrated here for your reference.

Review Questions

You can find the answers to these questions in Appendix C, “Answers to Review Questions.”

- 1 Given the following host address and subnet mask combinations, determine the subnet address and broadcast addresses:
 - 131.108.1.24 255.255.255.0
 - 151.108.100.67 255.255.255.128
 - 171.199.100.10 255.255.255.224
 - 161.88.40.54 255.255.255.192
- 2 Given the network 141.56.80.0 and a subnet mask of 255.255.254.0, how many hosts are available on this subnet?
- 3 What is the broadcast address for the subnet 131.45.1.0/24?
- 4 What is the purpose of the broadcast address in any given subnet?
- 5 Given the subnet in binary notation 1111111.11111111.00000000.00000000, what is the decimal equivalent?

- 6 Which routing protocols support VLSM and why?
- 7 Which routing protocols do not support VLSM?
- 8 Which subnet mask provides approximately 1022 hosts?
- 9 What is the equivalent subnet mask for the notation 131.108.1.0/24?
- 10 Identify the private address ranges defined in RFC 1918.

Summary

You have successfully worked through five scenarios using common techniques in today's large IP networks. You can now begin to apply this knowledge to the chapters ahead and work through more complex scenarios. The basic information described in this chapter can be applied to any networking scenario you come across when designing and implementing a Cisco-powered network or any network for that matter.

Table 1-11 summarizes the commands used in this chapter.

Table 1-11 *Summary of IOS Commands Used in This Chapter*

Command	Purpose
area <i>area-id</i> range <i>network mask</i>	Summarizes OSPF network ranges between area border routers.
router ospf <i>process id</i>	Enables OSPF routing. The process ID is local to the router. You can have more than one OSPF running.
router eigrp <i>autonomous domain ID</i>	Enables EIGRP routing under a common administrative control known as the autonomous domain or AD.
no auto-summary	Disables automatic summarization.
show interfaces ethernet 0/0	Displays Ethernet statistics on port 0/0.
version 2	Enables RIPv2.
[no] shutdown	Enables or disables an interface. All hardware interfaces are shut down by default.