ExamLabs

Designing Cisco Enterprise Networks (ENSLD)

Study Guide

Exam 300-420

Content

Chapter 1. Internet Protocol Version 4 (IPv4) Design

This chapter covers the following subjects:

- **IPv4 Header:** This section describes the fields in the IPv4 header.
- **IPv4 Addressing:** This section covers the format of IP addresses.
- **IP Address Subnets:** This section discusses how to divide IP address space into subnets.
- **IP Addressing Design:** This section shows how to fully design the IP addressing for regions, a campus, buildings, floors, and VLANs.
- **Address Assignment and Name Resolution:** This section describes how to assign IPv4 addresses to devices and groups and fully qualified domain names.

This chapter covers concepts and terminology related to Internet Protocol Version 4 (IPv4) address design. It provides an overview of IPv4 address structures and IPv4 address types. IPv4 is the version of the protocol that the Internet has used since the initial allocation of IPv4 addresses in 1981. In those days, the size of the enterprise determined the address class that was allocated. This chapter covers the IPv4 header to give you an understanding of IPv4 characteristics. The mid-1990s saw the implementation of classless interdomain routing (CIDR), Network Address Translation (NAT), and private address space to prevent the apparent exhaustion of the IPv4 address space. Companies implement variable-length subnet masking (VLSM) in their networks to provide intelligent address assignment and summarization. Separate IP subnets are used for IP phones and wireless LANs to segregate this traffic from wired data traffic. In 2011, the Internet Assigned Numbers Authority (IANA) allocated the last remaining address blocks of the IPv4 address space, thus depleting the free pool of IPv4 address space. Furthermore, in 2015, the American Registry for Internet Numbers (ARIN) issued the final IPv4 addresses in its free pool. Careful allocation of available IPv4 address space must be part of network design. A CCNP enterprise designer needs to have a deep understanding of all these concepts in order to create structured addressing plans for IPv4 addressing for a network.

This chapter covers the following objective from the ENSLD 300-420 exam:

- Create structured addressing plans for IPv4

“Do I Know This Already?” Quiz

The “Do I Know This Already?” quiz helps you identify your strengths and deficiencies in this chapter’s topics. This quiz, derived from the major sections in the “Foundation Topics” portion of the chapter, helps you determine how to spend your limited study time. Table 1-1 outlines the major topics discussed in this chapter and the “Do I Know This Already?” quiz questions that correspond to those topics. You can find the answers in Appendix A, “Answers to the ‘Do I Know This Already?’ Quiz Questions Q&A Questions.”

Table 1-1 “Do I Know This Already?” Foundation Topics Section-to-Question Mapping

Foundation Topics Section	Questions
IPv4 Header	4, 10
IPv4 Addressing	1, 5, 9
IPv4 Address Subnets	2–3, 6–7
Address Assignment and Name Resolution	8

1. Which of the following addresses is an IPv4 private address?

- **a.** 198.176.1.1
- **b.** 172.31.16.1
- **c.** 191.168.1.1
- **d.** 224.130.1.1

2. How many IP addresses are available for hosts in the subnet 198.10.100.64/27?

- **a.** 14
- **b.** 30
- **c.** 62
- **d.** 126

3. What subnet mask should you use in loopback addresses?

- **a.** 255.255.255.252
- **b.** 255.255.255.254
- **c.** 255.255.255.0
- **d.** 255.255.255.255

4. In what IPv4 field are the precedence bits located?

- **a.** Priority field
- **b.** IP Protocol field
- **c.** Type of Service field
- **d.** IP Options field

5. What type of address is 225.10.1.1?

- **a.** Unicast
- **b.** Multicast
- **c.** Broadcast
- **d.** Anycast

6. Which subnetworks are summarized by the summary route 150.10.192.0/21?

- **a.** 150.10.192.0/24, 150.10.193.0/24
- **b.** 150.10.192.0/22, 150.10.196.0/23, 150.10.197.0/24
- **c.** 150.10.192.0/22, 150.10.199.0/22
- **d.** 150.10.192.0/23, 150.10.194.0/23, 150.10.196.0/23, 150.10.199.0/24, 150.10.198.0/24

7. What type of network and subnet mask would you use to save address space in a point-to-point WAN link?

   a. 100.100.10.16/26
   b. 100.100.10.16/28
   c. 100.100.10.16/29
   d. 100.100.10.16/30

8. What protocol is used to automatically assign IP addresses?

   a. Dynamic Host Control Protocol
   b. Dedicated Host Configuration Protocol
   c. Dynamic Host Configuration Protocol
   d. Automatic Host Configuration Protocol

9. A company needs to use public IP addresses so that four network servers are accessible from the Internet. What technology is used to meet this requirement?

   a. DNS
   b. IPsec
   c. Static NAT
   d. Dynamic NAT

10. The DS field of DSCP is capable of how many codepoints?

   a. 8
   b. 32
   c. 64
   d. 128

Foundation Topics

This chapter reviews IPv4 headers, address classes, and assignment methods.

IP is the network layer protocol in TCP/IP. It contains logical addressing and information for routing packets throughout the internetwork. IP is described in RFC 791, which was prepared for the Defense Advanced Research Projects Agency (DARPA) in September 1981.

IP provides for the transmission of blocks of data, called *datagrams* or *packets*, from a source to a destination. The sources and destinations are identified by 32-bit IP addresses. The source and destination devices are workstations, servers, printers, sensors, cameras, IP phones, firewalls, and routers. A CCNP candidate must understand IPv4 logical address classes and assignment. The IPv4 protocol also provides for the fragmentation and reassembly of large packets for transport over networks with small maximum transmission units (MTUs). A CCNP candidate must have a good understanding of this packet fragmentation and reassembly.

Appendix C, "OSI Model, TCP/IP Architecture, and Numeric Conversion," provides an overview of the TCP/IP architecture and how it compares with the OSI model. It also reviews binary numbers and numeric conversion (to decimal), which is a skill needed to understand IP addresses and subnetting.

IPv4 Header

The best way to understand IPv4 is to know the IPv4 header and all its fields. Segments from Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) are passed on to IP for processing. The IP header is appended to the TCP or UDP segment. The TCP or UDP segment then becomes the IP data. The IPv4 header is 20 bytes in length when it uses no optional fields. The IP header includes the addresses of the sending host and the destination host. It also includes the upper-layer protocol, a field for prioritization, and a field for fragmentation. Figure 1-1 shows the IP header format.

Key Topic

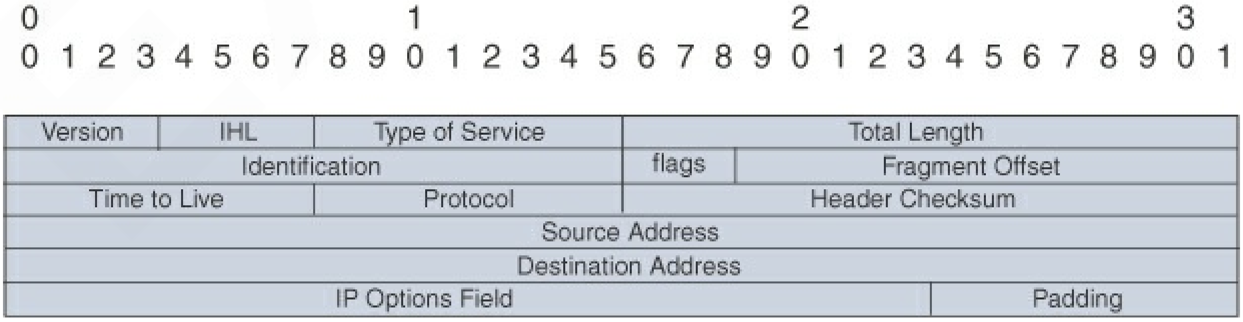

Figure 1-1 *IP Header*

The following is a description of each field in the IP header:

- **Version:** This field is 4 bits in length. It indicates the IP header's format, based on the version number. Version 4 is the current version, and this field is set to 0100 (4 in binary) for IPv4 packets. The Version field is set to 0110 (6 in binary) in IPv6 networks.
- **IHL (Internet Header Length):** This field is 4 bits in length. It indicates the length of the header in 32-bit words (4 bytes) so that the beginning of the data can be found in the IP header. The minimum value for a valid header is 5 (0101) for five 32-bit words.
- **ToS (Type of Service):** This field is 8 bits in length. Quality of service (QoS) parameters such as IP precedence and DSCP are found in this field. (These concepts are explained later in this chapter.)
- **Total Length:** This field is 16 bits in length. It represents the length of the datagram, or packet, in bytes, including the header and data. The maximum length of an IP packet can be $2^{16} - 1 = 65{,}535$ bytes. Routers use this field to determine whether fragmentation is necessary by comparing the total length with the outgoing MTU.
- **Identification:** This field is 16 bits in length. It is a unique identifier that denotes fragments for reassembly into an original IP packet.
- **Flags:** This field is 3 bits in length. It indicates whether the packet can be fragmented and whether more fragments follow. Bit 0 is reserved and set to 0. Bit 1 indicates May Fragment (0) or Do Not Fragment (1). Bit 2 indicates Last Fragment (0) or More Fragments to Follow (1).
- **Fragment Offset:** This field is 13 bits in length. It indicates (in bytes) where in the packet this fragment belongs. The first fragment has an offset of 0.
- **Time to Live:** This field is 8 bits in length. It indicates the maximum time the packet is to remain on the network. Each router decrements this field by 1 for loop avoidance. If this field is 0, the packet must be discarded. This scheme permits routers to discard undeliverable packets.
- **Protocol:** This field is 8 bits in length. It indicates the upper-layer protocol. The Internet Assigned Numbers Authority (IANA) is

responsible for assigning IP protocol values. Table 1-2 shows some key protocol numbers. You can find a full list at www.iana.org/assignments/protocol-numbers.

Key Topic

Table 1-2 IP Protocol Numbers

Protocol Number	IP Protocol
1	Internet Control Message Protocol (ICMP)
2	Internet Group Management Protocol (IGMP)
6	Transmission Control Protocol (TCP)
17	User Datagram Protocol (UDP)
41	IPv6 encapsulation
50	Encapsulating Security Payload (ESP)
51	Authentication Header (AH)
58	ICMPv6
88	Enhanced Interior Gateway Routing Protocol (EIGRP)
89	Open Shortest Path First (OSPF)
103	Protocol-Independent Multicast (PIM)
112	Virtual Router Redundancy Protocol (VRRP)

- **Header Checksum:** This field is 16 bits in length. The checksum does not include the data portion of the packet in the calculation. The checksum is verified and recomputed at each point the IP header is processed.
- **Source Address:** This field is 32 bits in length. It is the sender's IP

address.

- **Destination Address:** This field is 32 bits in length. It is the receiver’s IP address.
- **IP Options:** This field is variable in length. The options provide for control functions that are useful in some situations but unnecessary for the most common communications. Specific options are security, loose source routing, strict source routing, record route, and timestamp.
- **Padding:** This field is variable in length. It ensures that the IP header ends on a 32-bit boundary.

Table 1-3 summarizes the fields of the IP header.

Key Topic

Table 1-3 IPv4 Header Fields

Field	Length	Description
Version	4 bits	Indicates the IP header’s format, based on the version number. Set to 0100 for IPv4.
IHL	4 bits	Length of the header, in 32-bit words.
ToS	8 bits	QoS parameters.
Total Length	16 bits	Length of the packet, in bytes, including header and data.
Identification	16 bits	Identifies a fragment.
Flags	3 bits	Indicates whether a packet is fragmented and whether more fragments follow.
Fragment Offset	13 bits	Location of the fragment in the total packet.
Time to Live	8 bits	Decremented by 1 by each router. When this is 0,

		the router discards the packet.
Protocol	8 bits	Indicates the upper-layer protocol.
Header Checksum	16 bits	Checksum of the IP header; does not include the data portion.
Source Address	32 bits	IP address of the sending host.
Destination Address	32 bits	IP address of the destination host.
IP Options	Variable	Options for security, loose source routing, record route, and timestamp.
Padding	Variable	Added to ensure that the header ends in a 32-bit boundary.

ToS

Key Topic

The ToS field of the IP header is used to specify QoS parameters. Routers and Layer 3 switches look at the ToS field to apply policies, such as priority, to IP packets based on the markings. An example is a router prioritizing time-sensitive IP packets over regular data traffic such as web or email, which is not time sensitive.

The ToS field has undergone several definitions since RFC 791. Figure 1-2 shows the several formats of the ToS service field, based on the evolution of RFCs 791 (1981), 1349 (1992), 2474 (1998), and 3168 (2001). The following paragraphs describe this evolution.

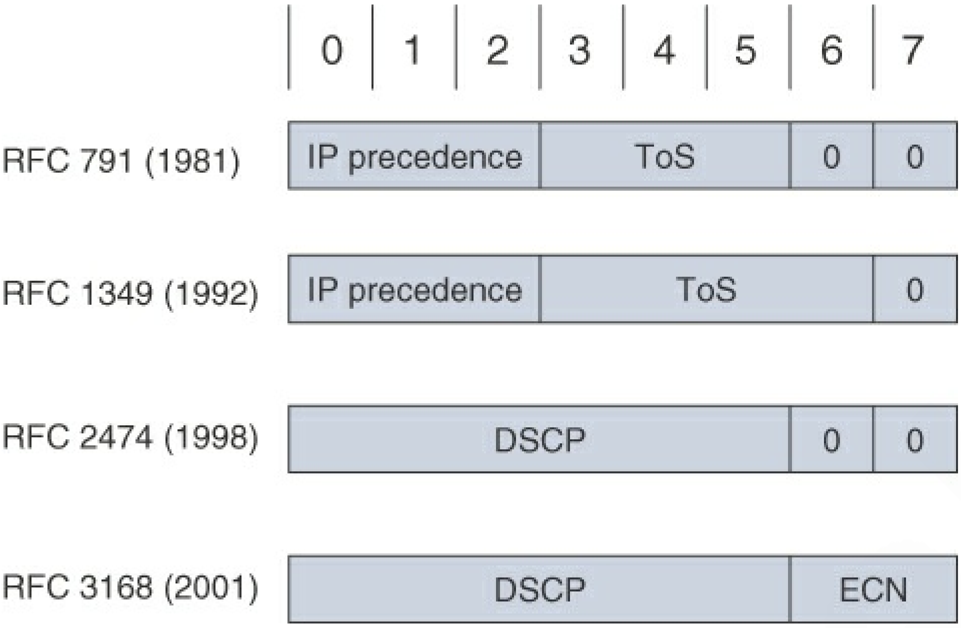

Figure 1-2 *Evolution of the IPv4 ToS Field*

The first 3 (leftmost) bits are the IP precedence bits. These bits define values that are used by QoS methods. The precedence bits especially help in marking packets to give them differentiated treatment with different priorities. For example, Voice over IP (VoIP) packets can get preferential treatment over regular data packets. RFC 791 describes the precedence bits as shown in Table 1-4.

Key Topic

Table 1-4 IP Precedence Bit Values

Decimal	Binary	IP Precedence Description
0	000	Routine
1	001	Priority
2	010	Immediate
3	011	Flash

4	100	Flash override
5	101	Critical
6	110	Internetwork control
7	111	Network control

All default traffic is set with 000 in the precedence bits. Voice traffic is usually set to 101 (critical) to give it priority over normal traffic. An application such as FTP is assigned a normal priority because it tolerates network latency and packet loss. Packet retransmissions are typically acceptable for normal traffic.

> **Note**
>
> It is common to see voice traffic classified as IP precedence 5, video traffic classified as IP precedence 4, and voice and video signaling classified as IP precedence 3. Default traffic remains as IP precedence 0.

RFC 1349 redefined bits 3 and 6 (expanding for ToS bits) to reflect a desired type of service optimization. Table 1-5 shows the ToS field values that indicate service parameters to use for IP packets.

Table 1-5 ToS Field Values

ToS Bits 3 to 6	Description
0000	Normal service
1000	Minimize delay
0100	Maximize throughput
0010	Maximize reliability
0001	Minimize monetary cost

In 1998, RFC 2474 redefined the ToS octet as the Differentiated Services (DS) field and further specified bits 0 through 5 as the Differentiated Services Codepoint (DSCP) bits to support differentiated services. RFC 3168 (2001) updates RFC 2474, with the specification of an Explicit Congestion Notification (ECN) field.

The DS field takes the form shown in Figure 1-2. The DS field provides more granular levels of packet classification by using 6 bits for packet marking. DS has 2^6 = 64 levels of classification, which is significantly higher than the eight levels of the IP precedence bits. These 64 levels are called codepoints, and they have been defined to be backward compatible with IP precedence values. RFCs 2474, 3246, and 3260 define three sets of PHBs: Class Selector (CS), Assured Forwarding (AF), and Expedited Forwarding (EF). The CS PHB set is for DSCP values that are compatible with IP precedence bits. The AF PHB set is used for queuing and congestion avoidance. The EF PHB set is used for premium service. The CS per-hop behaviors (PHB), in the form of xxx000, make it backward compatible with IP precedence.

A network designer uses DSCP to give priority to IP packets using Cisco routers. Routers should be configured to map these codepoints to PHBs with queuing or other bandwidth-management techniques. Table 1-6 compares DSCP and IP precedence values used to assign priority and apply policies to IP packets.

Table 1-6 DSCP and IP Precedence Values

IP Precedence	**Limitation**		**DSCP**		
Service Type	**Decimal**	**Binary**	**Class**	**Decimal**	**Codepoint**
Routine	0	000	Best effort	0	000000
Priority	1	001	Assured Forwarding (AF) Class 1	8 to 14	001xxx
Immediate	2	010	AF Class 2	16 to 22	010xxx

Flash	3	011	AF Class 3	24 to 30	011xxx
Flash override	4	100	AF Class 4	32 to 38	100xxx
Critical	5	101	Expedited Forwarding (EF)	40 to 46	101xxx
Internetwork control	6	110	Control	48	110xxx
Network control	7	111	Control	56	111xxx

RFC 2597 defines recommended values for AF codepoints with low, medium, and high packet-drop precedence. Table 1-7 shows the recommended AF codepoint values.

Table 1-7 DSCP AF Packet-Drop Precedence Values

Precedence	AF Class 1	AF Class 2	AF Class 3	AF Class 4
Low drop precedence	001010	010010	011010	100010
Medium drop precedence	001100	010100	011100	100100
High drop precedence	001110	010110	011110	100110

RFC 2598 defines the EF PHB for low loss, loss latency, and assured bandwidth types of traffic. This is considered a premium service. Traffic such as VoIP is classified as EF. The codepoint for EF is 101110, which corresponds to a DSCP value of 46.

When you are configuring Cisco routers, some options are preconfigured and summarize the defined values for DSCP (see Table 1-8).

Key Topic

Table 1-8 IP DSCP Values

DSCP Class	DSCP Codepoint Value	DSCP Decimal
Default	000000	0
CS1	001000	8
AF11	001010	10
AF12	001100	12
AF13	001110	14
CS2	010000	16
AF21	010010	18
AF22	010100	20
AF23	010110	22
CS3	011000	24
AF31	011010	26
AF32	011100	28
AF33	011110	30
CS4	100000	32
AF41	100010	34
AF42	100100	36
AF43	100110	38
CS5	101000	40

EF	101110	46
CS6	110000	48
CS7	111000	56

IPv4 Fragmentation

One key characteristic of IPv4 is fragmentation and reassembly. Although the maximum length of an IP packet is 65,535 bytes, most of the common lower-layer protocols do not support such large MTUs. For example, the MTU for Ethernet is approximately 1518 bytes. When the IP layer receives a packet to send, it first queries the outgoing interface to get its MTU. If the packet’s size is greater than the interface’s MTU, the layer fragments the packet.

When a packet is fragmented, it is not reassembled until it reaches the destination IP layer. The destination IP layer performs the reassembly. Any router in the path can fragment a packet, and any router in the path can fragment a fragmented packet again. Each fragmented packet receives its own IP header and identifier, and it is routed independently from other packets. Routers and Layer 3 switches in the path do not reassemble the fragments. The destination host performs the reassembly and places the fragments in the correct order by looking at the Identification and Fragment Offset fields.

If one or more fragments are lost, the entire packet must be retransmitted. Retransmission is the responsibility of a higher-layer protocol (such as TCP). Also, you can set the Flags field in the IP header to Do Not Fragment; in this case, the packet is discarded if the outgoing MTU is smaller than the packet.

IPv4 Addressing

This section covers the IPv4 address classes, private addressing, and NAT. The IPv4 address space was initially divided into five classes. Each IP address class is identified by the initial bits of the address. Classes A, B, and C are unicast IP addresses, meaning that the destination is a single host. IP Class D addresses are multicast addresses, which are sent to multiple hosts. IP Class E addresses are reserved. This section introduces IPv4 private addresses, which are selected address ranges that are reserved for use by companies in their private networks. These private addresses are not routed on the Internet. NAT translates between private and public addresses.

An IP address is a unique logical number to a network device or interface. An IP address is 32 bits in length. To make the number easier to read, the dotted-decimal format is used. The bits are combined into four 8-bit groups, each converted into decimal numbers (for example, 10.1.1.1). If you are not familiar with binary numbers, see Appendix C, which provides a review of binary and hexadecimal number manipulation.

Consider an example involving the binary IP address 01101110 00110010 11110010 00001010. Convert each byte into decimal.

Convert the first octet as follows:

0	1	1	0	1	1	1	0
0	+64	+32	+0	+8	+4	+2	+0 = 110

01101110 = 110

Convert the second octet as follows:

0	0	1	1	0	0	1	0
0	+0	+32	+16	+0	+0	+2	+0 = 50

00110010 = 50

Convert the third octet as follows:

1	1	1	1	0	0	1	0
128	+64	+32	+16	+0	+0	+2	+0 = 242

11110010 = 242

Convert the fourth octet as follows:

0	0	0	0	1	0	1	0

0	+0	+0	+0	+8	+0	+2	+0 = 10

00001010 = 10

The IP address in decimal is 110.50.242.10.

IPv4 Address Classes

IPv4 addresses have five classes: A, B, C, D, and E. In classful addressing, the most significant bits of the first byte determine the address class of the IP address. Table 1-9 shows the high-order bits of each IP address class.

Table 1-9 High-Order Bits of IPv4 Address Classes

Address Class	High-Order Bits
A	0xxxxxxx
B	10xxxxxx
C	110xxxxx
D	1110xxxx
E	1111xxxx

Again, the IPv4 Class A, B, and C addresses are unicast addresses. Such an address represents a single destination. Class D is for multicast addresses. Packets sent to a multicast address are sent to a group of hosts. Class E addresses are reserved for experimental use. IANA allocates the IPv4 address space. IANA delegates regional assignments to the five Regional Internet Registries (RIR):

- **ARIN (American Registry for Internet Numbers):** Covers USA, Canada, and some Caribbean Islands
- **RIPE NCC (Reseaux IP Europeens Network Control Center):** Covers Europe, the Middle East, and Central Asia

- **APNIC (Asia Pacific Network Information Center):** Covers the Asia Pacific Region
- **LACNIC (Latin America and Caribbean Network Information Center):** Covers Latin America and some Caribbean Islands
- **AfriNIC (African Network Information Centre):** Covers Africa

Updates to the IPv4 address space can be found at https://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.xhtml.

The following sections discuss each of these classes in detail.

Class A Addresses

Class A addresses range from 0 (00000000) to 127 (01111111) in the first byte. Network numbers available for assignment to organizations are from 1.0.0.0 to 126.0.0.0. Networks 0 and 127 are reserved. For example, 127.0.0.1 is reserved for the local host or host loopback. A packet sent to a local host address is sent to the local machine.

By default, for Class A addresses, the first byte is the network number, and the three remaining bytes are the host number. The format is *N.H.H.H*, where *N* is the network part and *H* is the host part. With 24 bits available, there are $2^{24} - 2 = 16,777,214$ IP addresses for host assignment per Class A network. We subtract 2 for the network number (all 0s) and broadcast address (all 1s). A network with this many hosts will surely not work with so many hosts attempting to broadcast on the network. As discussed later in this chapter, subnetting can be used to define smaller networks within a larger network address.

Class B Addresses

Class B addresses range from 128 (10000000) to 191 (10111111) in the first byte. Network numbers assigned to companies or other organizations are from 128.0.0.0 to 191.255.0.0. This section discusses the 16 networks reserved for private use later.

By default, for Class B addresses, the first 2 bytes are the network number, and the remaining 2 bytes are the host number. The format is *N.N.H.H*. With 16 bits available, there are $2^{16} - 2 = 65{,}534$ IP addresses for host assignment per Class B network. As with Class A addresses, having a segment with more than 65,000 hosts broadcasting will surely not work; you resolve this issue with subnetting.

Class C Addresses

Class C addresses range from 192 (11000000) to 223 (11011111) in the first byte. Network numbers assigned to companies are from 192.0.0.0 to 223.255.255.0. The format is *N.N.N.H*. With 8 bits available, there are $2^8 - 2 = 254$ IP addresses for host assignment per Class C network. *H* = 0 is the network number; *H* = 255 is the broadcast address.

Class D Addresses

Class D addresses range from 224 (11100000) to 239 (11101111) in the first byte. Network numbers assigned to multicast groups range from 224.0.0.1 to 239.255.255.255. These addresses do not have a host or network part. Some multicast addresses are already assigned; for example, routers running EIGRP use 224.0.0.10. You can find a full list of assigned multicast addresses at www.iana.org/assignments/multicast-addresses.

Class E Addresses

Class E addresses range from 240 (11110000) to 254 (11111110) in the first byte. These addresses are reserved for experimental networks. Network 255 is reserved for the broadcast address, such as 255.255.255.255. Table 1-10 summarizes the IPv4 address classes. Again, each address class can be uniquely identified in binary by the high-order bits.

Table 1-10 IPv4 Address Classes

Address Class	High-Order Bits	Network Numbers
A	0xxxxxxx	1.0.0.0 to 126.0.0.0*
B	10xxxxxx	128.0.0.0 to 191.255.0.0

C	110xxxxx	192.0.0.0 to 223.255.255.0
D	1110xxxx	224.0.0.1 to 239.255.255.255
E	1111xxxx	240.0.0.0 to 254.255.255.255

*Networks 0.0.0.0 and 127.0.0.0 are reserved as special-use addresses.

IPv4 Address Types

Key Topic

IPv4 addresses can be classified as one of three types:

- Unicast
- Broadcast
- Multicast

A unicast address represents a single interface of a host (PC, router, or server). It can be a source or destination IP address. A broadcast address is a destination IP address that is set to all other devices in a given address range; normally it is sent to all devices in the IP subnet. A multicast address is a destination IP address sent to a specific set of hosts. Table 1-11 summarizes IPv4 address types.

Key Topic

Table 1-11 IPv4 Address Type

IPv4 Address Type	Description
Unicast	The IP address of an interface on a single host. It can be a source address or a destination address.

Broadcast	An IP address that reaches all hosts in an address range. It is only a destination address.
Multicast	An IP address that reaches a group of hosts. It is only a destination address.

IPv4 Private Addresses

Key Topic

Some network numbers in the IPv4 address space are reserved for private use. These numbers are not routed on the Internet, so there is no way to reach them over an Internet connection. Many organizations today use private addresses in their internal networks with NAT to access the Internet. (NAT is covered later in this chapter.) Private addresses are explained in RFC 1918: *Address Allocation for Private Internets*, published in 1996. Creating private addresses was one of the first steps in dealing with the concern that the globally unique IPv4 address space would become exhausted. The availability of private addresses combined with NAT reduces the need for organizations to carefully define subnets to minimize the waste of assigned public global IP addresses.

The IP network address space reserved for private internetworks is 10/8, 172.16/12, and 192.168/16. It includes one Class A network, 16 Class B networks, and 256 Class C networks. Table 1-12 summarizes private address space. Large organizations can use network 10.0.0.0/8 to assign address space throughout the enterprise. Midsize organizations can use one of the Class B private networks 172.16.0.0/16 through 172.31.0.0/16 for IP addresses. The smaller Class C addresses, which begin with 192.168, can be used by corporations and are commonly used in home routers.

Key Topic

Table 1-12 IPv4 Private Address Space

Class Type	Start Address	End Address