

Appendix B: IP Addressing Basics

Contents

Overview	1
What Is an IP Address?	2
Network ID and Host ID	3
Converting IP Addresses from Binary to Decimal	4
Address Classes	6
Address Class Summary	8
Addressing Guidelines	9
Assigning Network IDs	10
Assigning Host IDs	11
What Is a Subnet Mask?	12
Default Subnet Masks	13
Determining the Destination of a Packet	14
Understanding Subnetting	16



Information in this document, including URL and other Internet Web site references, is subject to change without notice. Unless otherwise noted, the example companies, organizations, products, domain names, e-mail addresses, logos, people, places, and events depicted herein are fictitious, and no association with any real company, organization, product, domain name, e-mail address, logo, person, places or events is intended or should be inferred. Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

© 2000 Microsoft Corporation. All rights reserved.

Microsoft, MS-DOS, Windows, Windows NT, Active Directory, BackOffice, FrontPage, IntelliMirror, NetShow, Outlook, PowerPoint, Visual Studio, and Windows Media are either registered trademarks or trademarks of Microsoft Corporation in the U.S.A. and/or other countries.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

Simulations and interactive exercises were built with Macromedia Authorware

Overview

- What Is an IP Address?
- Network ID and Host ID
- Converting IP Addresses from Binary to Decimal
- Address Classes
- Address Class Summary
- Addressing Guidelines
- Assigning Network IDs
- Assigning Host IDs
- What Is a Subnet Mask?

*****ILLEGAL FOR NON-TRAINER USE*****


Using Transmission Control Protocol/Internet Protocol (TCP/IP) requires that an IP address be provided for each computer, either automatically by Microsoft® Windows® 2000, dynamically through the Dynamic Host Configuration Protocol (DHCP) service, or statically by using an IP address that you have obtained from your Internet service provider (ISP).

Objectives

At the end of this module, you will be able to:

- Identify the network identifier (ID) and host ID in a class A, B, or C IP address.
- Distinguish between a valid and invalid class A, B, or C IP address.
- Identify the network components that require an IP address.
- Describe subnetting.

What Is an IP Address?



*****ILLEGAL FOR NON-TRAINER USE*****

Each TCP/IP host is identified by a logical IP address. A unique IP address is required for each host and network component that communicates by using TCP/IP.

The IP address identifies a system's location on the network in the same way that a street address identifies a house on a city block. Just as a street address must identify a unique residence, an IP address must be globally unique and have a uniform format.


Network ID

Each IP address defines the network ID and host ID. The network ID identifies the systems that are located on the same physical segment. All systems on the same physical segment must have the same network ID. The network ID must be unique to the internetwork.

Host ID

The host ID identifies a workstation, server, router, or other TCP/IP host within a segment. The address for each host must be unique to the network ID.

Network ID and Host ID




*****ILLEGAL FOR NON-TRAINER USE*****

Each IP address is 32 bits long and is composed of four 8-bit fields, called *octets*. Octets are separated by periods. The octet represents a decimal number in the range 0–255. This format is called *dotted decimal notation*. The following is an example of an IP address in binary and dotted decimal formats.

Binary format	Dotted decimal notation
10000011 01101011 00000011 00011000	131.107.3.24

Converting IP Addresses from Binary to Decimal



*****ILLEGAL FOR NON-TRAINER USE*****

Each bit position in an octet has an assigned decimal value. A bit that is set to 0 always has a zero value.

A bit that is set to 1 can be converted to a decimal value. The low-order bit represents a decimal value of one. The high-order bit represents a decimal value of 128. The highest decimal value of an octet is 255—that is, when all bits are set to 1.

Converting Bits From Binary to Decimal

The following table shows how the bits in one octet are converted from binary code to a decimal value.

Binary code	Bit values	Decimal value
00000000	0	0
00000001	1	1
00000011	1+2	3
00000111	1+2+4	7
00001111	1+2+4+8	15
00011111	1+2+4+8+16	31
00111111	1+2+4+8+16+32	63
01111111	1+2+4+8+16+32+64	127
11111111	1+2+4+8+16+32+64+128	255

Practice

1. Convert the following binary numbers to decimal format.


Binary value	Decimal value
10001011	139
10101010	170
10111111 11100000 00000111 10000001	191.224.7.129
01111111 00000000 00000000 00000001	127.0.0.1

2. Convert the following decimal values to binary format.

Decimal value	Binary value
250	11111010
19	00010011
109.128.255.254	01101101 10000000 11111111 11111110
131.107.2.89	10000011 01101011 00000010 01011001

Tip Use the calculator (scientific view) in the **Accessories** group to convert decimal format to binary format, and vice versa.

Address Classes



*****ILLEGAL FOR NON-TRAINER USE*****

The Internet community has defined five IP address classes to accommodate networks of varying sizes. TCP/IP in Windows 2000 supports class A, B, and C addresses assigned to hosts. The class of address defines which bits are used for the network ID and which bits are used for the host ID. It also defines the possible number of networks and the number of hosts per network.

The following chart shows the network and host ID fields for class A, B, and C IP addressing:

Class	IP address	Network ID	Host ID
A	w.x.y.z	w	x.y.z
B	w.x.y.z	w.x	y.z
C	w.x.y.z	w.x.y	z

Class A

Class A addresses are assigned to networks with a very large number of hosts. The high-order bit in a class A address is always set to zero. The next seven bits (completing the first octet) complete the network ID. The remaining 24 bits (the last three octets) represent the host ID. This allows for 126 networks and approximately 17 million hosts per network.

Class B

Class B addresses are assigned to medium-sized to large-sized networks. The two high-order bits in a class B address are always set to binary 1 0. The next 14 bits (completing the first two octets) complete the network ID. The remaining 16 bits (last two octets) represent the host ID. This allows for 16,384 networks and approximately 65,000 hosts per network.

Class C

Class C addresses are used for small local area networks (LANs). The three high-order bits in a class C address are always set to binary 1 1 0. The next 21 bits (completing the first three octets) complete the network ID. The remaining 8 bits (last octet) represent the host ID. This allows for approximately 2 million networks and 254 hosts per network.

Class D

Class D addresses are used for multicast group usage. A multicast group may contain one or more hosts, or none at all. The four high-order bits in a class D address are always set to binary 1 1 1 0. The remaining bits designate the specific group in which the client participates. There are no network or host bits in the multicast operations. Packets are passed to a selected subset of hosts on a network. Only those hosts registered for the multicast address accept the packet. Microsoft uses class D addresses for applications to multicast data to hosts on an internetwork, including Windows Internet Name Service (WINS) and Microsoft NetShow®.

Class E

Class E is an experimental address that is not available for general use; it is reserved for future use. The high-order bits in a class E address are set to 1 1 1 1.

Note For more information about multicasting, see the white paper *Multicasting* under **Additional Reading** on the Web page on the Student Materials compact disc.

Address Class Summary

	Number of Networks	Number of Hosts per Network	Range of Network IDs (First Octet)
Class A	126	16,777,214	1–126
Class B	16,384	65,534	128–191
Class C	2,097,152	254	192–223

*****ILLEGAL FOR NON-TRAINER USE*****

The graphic in the previous slide summarizes the number of networks and number of hosts per network, and the range of network IDs in class A, B, and C IP addresses. The 32-bit IP addressing scheme presented in the graphic supports a total of 3,720,314,628 hosts.

Addressing Guidelines


- **Network ID Cannot Be 127**
 - 127 is reserved for loopback functions
- **Network ID and Host ID Cannot Be 255 (All Bits Set to 1)**
 - 255 is a broadcast address
- **Network ID and Host ID Cannot Be 0 (All Bits Set to 0)**
 - 0 means “this network only”
- **Host ID Must Be Unique to the Network**

*****ILLEGAL FOR NON-TRAINER USE*****

Follow these guidelines when assigning network IDs and host IDs:

- The network ID cannot be 127. This ID is reserved for loopback functions.
- The network ID and host ID cannot be 255 (all bits set to 1). If all bits are set to 1, the address is interpreted as a broadcast rather than a host ID.
- The network ID and host ID bits cannot all be set to 0. If all bits are set to 0, the address is interpreted to mean “this network only.”
- The host ID must be unique to the local network ID.

Assigning Network IDs



*****ILLEGAL FOR NON-TRAINER USE*****

The network ID identifies the TCP/IP hosts that are located on the same physical network. All hosts on the same physical network must be assigned the same network ID to communicate with each other.


If your networks are connected by routers, a unique network ID is required for each wide area connection. For example, in the graphic:

- Networks 1 and 3 represent two routed networks.
- Network 2 represents the wide area network (WAN) connection between the routers. Network 2 requires a network ID so that the interfaces between the two routers can be assigned unique host IDs.

Notes If you plan to connect your network to the Internet, you must obtain the network ID portion of the IP address to guarantee IP network ID uniqueness. For domain name registration and IP network number assignment, see your ISP.

For more information about IP address allocation for private networks, see RFC 1918 under **Additional Reading** on the Web page on the Student Materials compact disc.

Assigning Host IDs



*****ILLEGAL FOR NON-TRAINER USE*****

The host ID identifies a TCP/IP host within a network and must be unique to the network ID. All TCP/IP hosts, including interfaces to routers, require unique host IDs.

The host ID of the router interface is the IP address configured as a default gateway to the workstation when TCP/IP is installed. For example, for the host on subnet 1 with an IP address of 124.0.0.27, the IP address of the default gateway is 124.0.0.1.

Valid Host IDs

The following table lists the valid ranges of host IDs for a private internetwork.


Address class	Beginning range	Ending range
Class A	w.0.0.1	w.255.255.254
Class B	w.x.0.1	w.x.255.254
Class C	w.x.y.1	w.x.y.254

Suggestions for Assigning Host IDs

There are no rules for assigning valid IP addresses. You can number all TCP/IP hosts consecutively, or you can number them so they can easily be identified—for example:

- Assign host IDs in groups based on host or server type.
- Designate routers by their IP address.

What Is a Subnet Mask?



*****ILLEGAL FOR NON-TRAINER USE*****

A subnet mask is a 32-bit address used to:

- Block out a portion of the IP address to distinguish the network ID from the host ID.
- Specify whether the destination host's IP address is located on a local network or a remote network.

Each host on a TCP/IP network requires a subnet mask—either a default subnet mask, which is used when a network is not divided into subnets, or a custom subnet mask, which is used when a network is divided into subnets.

Default Subnet Masks

Address Class	Bits Used for Subnet Mask				Dotted Decimal Notation
Class A	11111111	00000000	00000000	00000000	255.0.0.0
Class B	11111111	11111111	00000000	00000000	255.255.0.0
Class C	11111111	11111111	11111111	00000000	255.255.255.0

Class B Example

IP Address	131.107.	16.200
Subnet Mask	255.255.	0.0
Network ID	131.107.	y.z
Host ID	w.x.	16.200

*****ILLEGAL FOR NON-TRAINER USE*****

A default subnet mask is used on TCP/IP networks that are not divided into subnets. All TCP/IP hosts require a subnet mask, even on a single-segment network. The default subnet mask that you will use depends on the address class.

All bits that correspond to the network ID are set to 1. The decimal value in each octet is 255.

All bits that correspond to the host ID are set to 0.

Determining the Destination of a Packet

■ Local and Destination Host's Subnet Masks Are ANDed

- 1 AND 1 = 1
- Other combinations = 0
- If ANDed results of source and destination hosts match, the destination is local.

IP Address	10011111	11100000	00000111	10000001
Subnet Mask	11111111	11111111	00000000	00000000
Result	10011111	11100000	00000000	00000000

*****ILLEGAL FOR NON-TRAINER USE*****

ANDing is the internal process that TCP/IP uses to determine whether a packet is destined for a host on a local network or a remote network.

When TCP/IP is initialized, the host's IP address is ANDed with its subnet mask. Before a packet is sent, the destination IP address is ANDed with the same subnet mask. If both results match, IP knows that the packet belongs to a host on the local network. If the results do not match, the packet is sent to the IP address of an IP router.

To AND the IP address to a subnet mask, TCP/IP compares each bit in the IP address to the corresponding bit in the subnet mask. If both bits are set to 1, the resulting bit is 1. If there is any other combination, the resulting bit is 0.

Bit combination	Result
1 AND 1	1
1 AND 0	0
0 AND 0	0
0 AND 1	0

Practice

AND the following IP addresses to determine whether the destination IP address belongs to a host on a local network or a remote network.

Source (host) IP address	10011001 10101010 00100101 10100011
Subnet mask	<u>11111111 11111111 00000000 00000000</u>
Result	

Destination IP address	11011001 10101010 10101100 11101001
Subnet mask	<u>11111111 11111111 00000000 00000000</u>
Result	

1. Do the results match?

No.

2. Is the destination IP address located on a local or remote network?

Remote.

Note ANDing is a process that IP uses internally and is not a process that a user would typically do.


◆ Understanding Subnetting

- What Is a Subnet?
- Implementing Subnetting
- What Are Subnet Mask Bits?
- Defining a Subnet Mask
- Subnetting More than One Octet
- Defining Subnet IDs
- Shortcut to Defining Subnet IDs
- Defining Host IDs for a Subnet
- Supernetting

*****ILLEGAL FOR NON-TRAINER USE*****

Network IDs and host IDs within an IP address are distinguished by using a subnet mask. Each subnet mask is a 32-bit number that uses consecutive bit groups for identification. The network ID is identified by bit groups that are all set to 1, and the host ID portions of an IP address are identified by bit groups that are all set to 0.

What Is a Subnet?



*****ILLEGAL FOR NON-TRAINER USE*****

A subnet is a physical segment in a TCP/IP environment that uses IP addresses derived from a single network ID. Typically, an organization acquires one network ID from its ISP.

Dividing the network into subnets requires that each segment use a different network ID or subnet ID. A unique subnet ID is created for each segment by partitioning the bits in the host ID into two parts. One part is used to identify the segment as a unique network, and the other part is used to identify the hosts. This is referred to as *subnetting* or *subnetworking*.

Subnetting Benefits

Organizations use subnetting to apply one network across multiple physical segments. Therefore, you can:

- Mix different technologies, such as Ethernet and token ring.
- Overcome limitations of current technologies, such as exceeding the maximum number of hosts per segment.
- Reduce network congestion by redirecting traffic and reducing broadcasts.

Note For more information about subnetting, see RFC 950 under **Additional Reading** on the Web page on the Student Materials compact disc.

Implementing Subnetting


- **Determine the Number of Required Network IDs**
 - One for each subnet
 - One for each wide-area network connection
- **Determine the Number of Required Host IDs per Subnet**
 - One for each TCP/IP host
 - One for each router interface
- **Based on Your Requirements, Define**
 - One subnet mask based on requirements
 - A unique subnet ID for each physical segment based on the subnet mask
 - Valid host IDs for each subnet based on the subnet ID

*****ILLEGAL FOR NON-TRAINER USE*****

Before you implement subnetting, you need to determine your current requirements and plan for future requirements. Follow these guidelines:

- Determine the number of required network IDs. You require:
 - One network ID for each subnet.
 - One network ID for each wide-area connection.
- Determine the number of required host IDs per subnet.
 - Each TCP/IP host requires at least one IP address.
 - Each router interface requires at least one IP address.
- Based on your requirements, define:
 - One subnet mask for your entire network based on your requirements.
 - A unique subnet ID for each physical segment based on the subnet mask.
 - A range of valid host IDs for each subnet based on the subnet ID.

What Are Subnet Mask Bits?




*****ILLEGAL FOR NON-TRAINER USE*****

Before you define a subnet mask, you should determine the number of segments and hosts per segment that you will require in the future.

As the graphic in the previous slide illustrates, when more bits are used for the subnet mask, more subnets are available, but fewer hosts are available per subnet. Using more bits than needed will allow for growth in the number of subnets, but will limit the growth in the number of hosts. Using fewer bits than needed will allow for growth in the number of hosts, but will limit the growth in the number of subnets.

Defining a Subnet Mask



*****ILLEGAL FOR NON-TRAINER USE*****

Defining a subnet mask is required if you are dividing your network into subnets. Follow these steps to define a subnet mask:

1. Once you have determined the number of physical segments in your network environment, convert this number to binary format.
2. Count the number of bits required to represent the number of physical segments in binary. For example, if you need six subnets, the binary value is 110. Representing six in binary requires three bits.
3. Convert the required number of bits to decimal format in high order (from left to right). For example, if three bits are required, configure the first three bits of the host ID as the subnet ID. The decimal value for binary 11100000 is 224. The subnet mask is 255.255.224.0 (for a class B address).

Contiguous Mask Bits

Because subnets are defined by the subnet mask, there is nothing to prevent an administrator from using low-order or unordered bits to determine the subnet ID. When subnetting was initially defined in RFC 950, it was recommended that subnet IDs be derived from high-order bits. Today, however, few router vendors support the use of low-order or non-order bits in subnet IDs. Furthermore, it is now a requirement that the subnet ID make use of contiguous, high-order bits of the local address portion of the subnet mask.

Conversion Tables

The following table lists the subnet masks already converted using one octet for class A networks.

Number of subnets	Required number of bits	Subnet mask	Number of hosts per subnet
0	1	Invalid	Invalid
2	2	255.192.0.0	4,194,302
6	3	255.224.0.0	2,097,150
14	4	255.240.0.0	1,048,574
30	5	255.248.0.0	524,286
62	6	255.252.0.0	262,142
126	7	255.254.0.0	131,070
254	8	255.255.0.0	65,534


The following table lists the subnet masks already converted using one octet for class B networks.

Number of subnets	Required number of bits	Subnet mask	Number of hosts per subnet
0	1	Invalid	Invalid
2	2	255.255.192.0	16,382
6	3	255.255.224.0	8,190
14	4	255.255.240.0	4,094
30	5	255.255.248.0	2,046
62	6	255.255.252.0	1,022
126	7	255.255.254.0	510
254	8	255.255.255.0	254

The following table lists the subnet masks already converted using one octet for class C networks.

Required Number of subnets	Required number of bits	Subnet mask	Number of hosts per subnet
Invalid	1	Invalid	Invalid
1–2	2	255.255.255.192	62
3–6	3	255.255.255.224	30
7–14	4	255.255.255.240	14
15–30	5	255.255.255.248	6
31–62	6	255.255.255.252	2
Invalid	7	Invalid	Invalid
Invalid	8	Invalid	Invalid

Subnetting More than One Octet



*****ILLEGAL FOR NON-TRAINER USE*****

Until this point, we have worked within one octet to define a subnet mask. At times, it may be advantageous to subnet using more than one octet, or more than eight bits.

For example, suppose you are on a team responsible for configuring an intranet for a large corporation. The corporation plans to internally connect its sites that are distributed across Europe, North America, and Asia. This totals approximately 30 geographical locations with almost 1,000 subnets and an average of 750 hosts per subnet.


It is possible to use several class B network IDs and further subnet them. To meet our host requirements per subnet with a class B network address, we will need to use a subnet mask of 255.255.252.0. Further adding our requirement of subnets, we will need at least 16 class B addresses.

However, there is an easier way. Because we are on an intranet, we can use a private network. If we choose to allocate a class A network ID of 10.0.0.0, we can plan for growth and meet our requirements at the same time. Obviously, subnetting only the second octet will not meet our requirements of 1,000 subnets. However, if we subnet both the second octet and a portion of the third octet, we can meet all of our requirements with one network ID.

Network ID	Subnet mask	Subnet mask (binary)
10.0.0.0	255.255.248.0	11111111 11111111 11111000 00000000

Using 13 bits for the subnet ID in a class A address, we have allocated 8,190 subnets, each with up to 2,046 hosts. We have met our requirements with flexibility for growth.

Defining Subnet IDs



*****ILLEGAL FOR NON-TRAINER USE*****

The subnet ID for a physical segment is defined using the same number of host bits as used for the subnet mask. The possible bit combinations are evaluated and then converted to a decimal format. Follow these steps to define a range of subnet IDs for an internetwork:

1. Using the same number of bits as used for the subnet mask, list all possible bit combinations.
2. Cross out values that use all 0s or 1s. All 0s and 1s are invalid IP addresses and network IDs, because all 0s indicate “this network only” and all 1s match the subnet mask.
3. Convert to decimal the subnet ID bits for each subnet. Each decimal value represents a single subnet. This value is used to define the range of host IDs for a subnet.

Special Case Subnet Addresses

Subnet IDs comprising all 0s or all 1s are called *special-case subnet addresses*. A subnet ID of all 1s indicates a subnet broadcast, and a subnet ID of all 0s indicates “this subnet.” When subnetting, it is recommended that you do not use these subnet IDs. However, it is possible to use these special-case subnet addresses if they are supported by all routers and hardware on your network. RFC 950 discusses the limitations imposed when using special-case addresses.

Shortcut to Defining Subnet IDs

- 1 List the Number of Bits (High Order) Used for Subnet Mask
- 2 Convert the Bit with the Lowest Value to Decimal
- 3 Increment the Value for Each Bit Combination

11000000

64

0		
+ 64		
= 64	w.x.64.1	→ w.x.127.254
+ 64		
= 128	w.x.128.1	→ w.x.191.254
+ 64		
192		

*****ILLEGAL FOR NON-TRAINER USE*****

Using the previous method is impractical when you are using more than four bits for your subnet mask because it requires listing and converting many bit combinations. Follow these steps to define a range of subnet IDs:

1. List the number of bits (in high order) used for the subnet ID. For example, if two bits are used for the subnet mask, the binary octet is 11000000.
2. Convert the bit with the lowest value to decimal format. This is the increment value to determine each subnet. For example, if you use two bits, the lowest value is 64.
3. Starting with zero, increment the value for each bit combination until the next increment is 256.


Tip If you know the number of bits you need, you can raise two to the power of the bit, and then subtract two to determine the possible bit combinations.

Determining the Number of Valid Subnets

To determine the number of valid subnets:

1. Convert the number of bits used for the subnet ID to low order.
2. Convert the low order binary number to decimal format.
3. Subtract one.

Defining Host IDs for a Subnet



*****ILLEGAL FOR NON-TRAINER USE*****

The result of each incremented value indicates the beginning of a range of host IDs for a subnet. If you increment the value one additional time, you can determine the end of the range (one less than the subnet mask).

The following table shows the valid range of host IDs on a class B subnet using three bits for the subnet mask.

Bit values	Decimal value	Beginning range value	Ending range value
00000000	0	Invalid	Invalid
00100000	32	x.y.32.1	x.y.63.254
01000000	64	x.y.64.1	x.y.95.254
01100000	96	x.y.96.1	x.y.127.254
10000000	128	x.y.128.1	x.y.159.254
10100000	160	x.y.160.1	x.y.191.254
11000000	192	x.y.192.1	x.y.223.254
11100000	224	Invalid	Invalid


Determining the Number of Host per Subnet

To determine the number of hosts per subnet:

1. Calculate the number of bits available for the host ID. For example, if you are given a class B address that uses 16 bits for the network ID and two bits for the subnet ID, you have 14 bits remaining for the host ID.
2. Convert the binary host ID bits to decimal. For example, 11111111111111 in binary is converted to 16,383 in decimal format.
3. Subtract one.

Tip If you know the number of host ID bits that you need, you can raise two to the power of the number of host ID bits, and then subtract two.

Supernetting



*****ILLEGAL FOR NON-TRAINER USE*****

To prevent the depletion of network IDs, Internet authorities devised a scheme called *supernetting*. In opposition to subnetting, supernetting borrows bits from the network ID and masks them as the host ID for more efficient routing. For example, rather than allocating a Class B network ID to an organization that has 2,000 hosts, the American Registry for Internet Numbers (ARIN) allocates a range of eight Class C network IDs. Each class C network ID accommodates 254 hosts for a total of 2,032 host IDs.

Classless Inter-Domain Routing

While this technique helps conserve Class B network IDs, it creates a new problem. Using conventional routing techniques, the routers on the Internet now must have an additional seven entries in their routing tables to route IP packets to the organization. To prevent overwhelming the Internet routers, a technique called *Classless Inter-Domain Routing* (CIDR) is used to collapse the eight entries used in the above graphic to a single entry corresponding to all of the class C network IDs used by that organization.

Allocating Network IDs

To express the situation in which eight class C network IDs are allocated starting with the network ID 220.78.168.0 and ending with network ID 220.78.175.0, the entry in the routing table becomes:

Network ID	Subnet mask	Subnet mask (binary)
220.78.168.0	255.255.248.0	1111111111 11111111 11111000 00000000

In supernetting, the destination of a packet is determined by ANDing the destination IP address and the subnet mask of the routing entry. If a match is found to the Network ID, the route is used.

Note For more information about CIDR, see RFC 1518 and RFC 1519 under **Additional Reading** on the Web page on the Student Materials compact disc.
