SOLUTIONS MANUAL

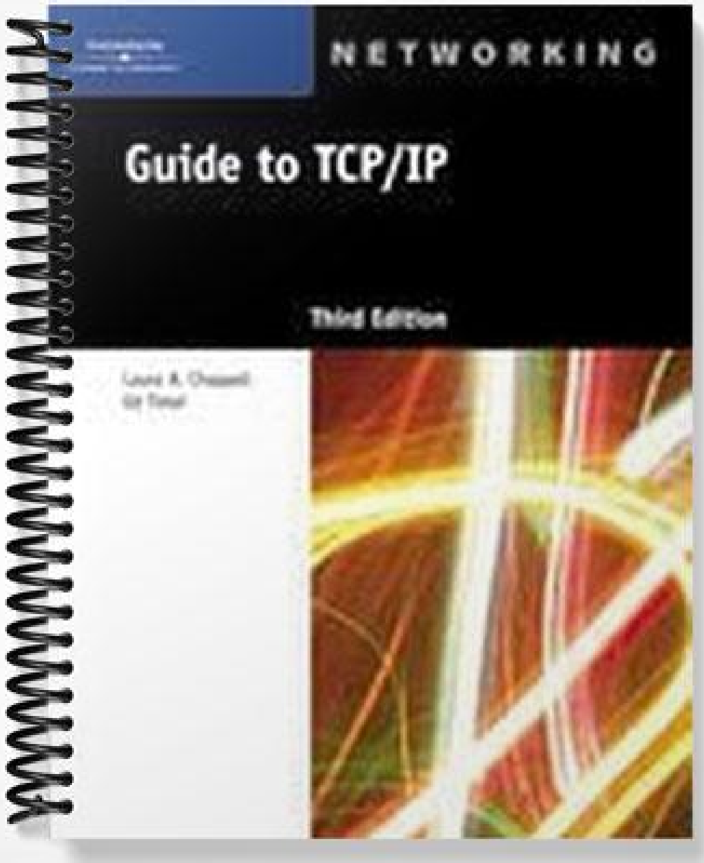

Chapter 2

IP Addressing and Related Topics

At a Glance

Instructor’s Manual Table of Contents

Lecture Notes

Overview

This chapter covers the structure and function of IP (Internet Protocol) addresses - those arcane four-number sequences that look like 24.29.72.3, but which uniquely identify all public network interfaces that use TCP/IP on the entire Internet. As you come to understand and appreciate IP addresses, you will learn how they are constructed, the classes into which they may (or may not) be relegated, and what roles these addresses play as traffic finds its way around a network.

Chapter Objectives

- Understand IP addressing, anatomy and structures, and addresses from a computer's point of view
- Recognize and describe the various IP address classes from A to E, and explain how they're composed and used
- Understand the nature of IP address limitations, and how techniques like Classless Inter-Domain Routing and Network Address Translation ease those limitations
- Define the terms subnet and supernet, and apply your knowledge of how subnets and supernets work to solve specific network design problems
- Understand how public and private Internet addresses are assigned, how to obtain them, and how to use them properly
- Recognize the importance and value of an IP addressing scheme

Teaching Tips

IP Addressing Basics

1. Often we think of the "language" of computers as binary. While this is true, the computer does not "see" binary numbers. Instead, it responds to the presence or absence of electrical current. Binary ones and zeros are actually the presence and absence of that current.

2. Referring back to Chapter 1, students can think of domain names in the same way they think of a model. Network communication does not require domain names in order to work. That is why they are *symbolic*. The naming system is there to make it easier for people to use and understand networks, just like modeling systems.

3. The method of expressing IP addresses as octets is for human convenience. It is easier for a person to remember a group of numbers that are "broken up" in some manner, like a social security number or phone number, than to remember a long string of digits. The octet system represents a long string of binary numbers.

4. IP or OSI model layer 3 addresses are routable and changeable, unlike the MAC or layer 2 addresses. This information will be explained in detail over the next two chapters but it is important for students to understand the distinction as soon as they can since how a computer communicates on a network differs greatly depending on which addressing system is being used. Also, both addressing systems work together to allow network communication as will be seen as the class progresses.

5. On an individual computer, the network interface card (NIC) has the MAC address permanently “burned in” or assigned to the NIC. As long as that particular NIC is part of the computer, the computer’s MAC address never changes.

6. When a computer sends a datagram out onto the wire, the layer 2 MAC address field contains the MAC address of the sending computer. The layer address field does permanently contain the sending computer’s address. In order for the datagram to be forwarded throughout the network, whenever the datagram is received by a router or switch and is forwarded, the layer 2 address encapsulation is stripped off and replaced with the layer 2 MAC address of the router or switch’s sending interface. When the datagram is finally received by the destination computer, it reads the source MAC address of the datagram as that of the interface of the switch that most recently forwarded the datagram to the receiver, not the MAC address of the source computer.

7. Remind the class that, in a sense, a computer’s MAC address is part of the NIC installed on the computer, the MAC address is considered permanent and unchanging. In the sense that a datagram traveling across a network must change the source MAC address field every time it passes through a router or switch, the MAC address as a field in a datagram is considered changed.

Anatomy of an IP Address

1. One of the things you can point out to your students is that, if an IP address is manually assigned to a computer and another computer is already using that address, a message will be generated stating this and requesting that the user sets a different IP address.

2. Most networks of any size at all use DHCP (Dynamic Host Configuration Protocol) services to automatically assign IP addresses to all the hosts on a network. If the DHCP server is configured correctly, there should not be an issue with IP address conflict.

IP Address Classes

1. The idea of a network versus a host address is perhaps like the difference between a person’s city address and street address. Someone may live in a “network” like Chicago, New York, or Denver, but the “host” address within that network is 321 Bannock Street or 5678 Glenwood Avenue. The first address gets the “mail” to the general area and the second address delivers it to the specific location.

More About Class A Addresses

1. Although it seems that a Class A address should include network addresses from 1 through 127, the 127 network address is reserved for loopback testing, as described in the text. You can have your class do a “mini-lab” right here by asking them to open a command window and type: “ping 127.0.0.1”. This is a basic test for NIC functioning. Also have them type: “ping localhost” and report on the results.

2. The First Octet Value of a Class A network is a range from 1 to 126. Your students will need to memorize these ranges for each class of network. Network technicians usually can recognize IP addresses by class easily. This often saves time when trying to diagnose a connection problem in order to rule out a possible incompatible address.

3. The private IP address range for a Class A network is 10.0.0.0 to 10.255.255.255. The concept of a private address range is discussed in the next section of the text.

More About Class B Addresses

1. The First Octet Value of a Class B network is a range from 128 to 151.

2. This section mentions public IP addresses. Each class of address has a range of addresses set aside for private use. These addresses are not routable on the Internet and, if a router encounters a packet with a source address from a private range, that packet will be dropped. Often these private addresses are used for small home networks, computer lab networks and even business networks in concert with NAT (Network Address Translation). NAT is discussed later on in this chapter.

3. The private IP address range for a Class B Network is 172.16.0.0 to 172.31.255.255.

More About Class C Addresses

1. The First Octet Value of a Class C network is a range from 192 to 223.

2. The private IP address range for a Class C Network is 192.168.0.0 to 192.168.255.255.

3. A Class C addressing scheme is the one most commonly encountered. Of the three public classes used, both Class A and B have almost no available addresses left. We will discuss the matter of the “disappearing” IP addresses and what strategies have been developed to combat this later on in the chapter.

More About Class D and E Addresses

1. The First Octet Value of a Class D network is a range from 224 to 239.

2. Class D addresses are known as “multicast” addresses and are often used by routers to transmit changes in routing tables to other routers using a single message.

3. The First Octet Value of a Class E network is a range from 240 to 254.

4. Class E addresses are exclusively reserved for Internet experimentation. As of the writing of this manual, there are no current Class E projects active.

Quick Quiz 1

1. The physical numeric address functions at a sub-layer of the Data Link layer in the OSI network reference model, called the ____________________ layer.
   Answer: Media Access Control (MAC)

2. True or False: Multicast addresses come in handy when a class of devices, such as routers, must be updated with the same information on a regular basis.
   Answer: True

3. A(n) ____________________ is the router or other device that will forward traffic to the host's physical network.
   Answer: IP gateway

4. True or False: Duplication of numeric IP addresses is not allowed because that would lead to confusion.
   Answer: True

Network, Broadcast, Multicast, and Other Special IP Addresses

1. As mentioned before, your students can think of addresses as being broken down from larger areas (like cities) to smaller areas (like street addresses). This is how IP addresses work. They are hierarchical and each portion describes a different aspect of the address. This is also how a zip code works on a letter a person mails.

2. Part of the hierarchy of addresses is the host address and the broadcast address. These are the two addresses that the text mentions are held out or deducted from the total pool of addresses in any class.

3. The network portion is the address for the overall network. If you mail a letter to a friend in Austin, it first has to make it to Austin before finding the particular house. "Austin" is the network address.

4. A broadcast address is a bit more difficult to explain in the above type of analogy. If someone wants to send a message to all the devices on a single network, it would be like sending copies of the same letter to everyone in Austin. Perhaps this is like a mass mailing of some advertisement. Everybody gets the same information at pretty much the same time. Other than a broadcast, a network device will ignore any traffic on the network that is not addressed specifically to them.

Broadcast Packet Structures

1. The text mentions that IP packets have two address fields: one at the Network Layer and one at the Data Link Layer. Broadcasts use the Data Link Layer (Layer 2 on the OSI Model) and these addresses are not forwarded by routers. The relationship between Network and Data Link Layer addresses and how computers communicate using them will be developed as class progresses.

Multicast Packet and Address Structures

1. It was mentioned before that generally, network devices will ignore any traffic on the wire except for broadcasts and messages specifically addressed to them. This includes multicast (Class D) transmissions. Routers using multicast transmissions must be configured to "listen" for them or they will be ignored. This will result in a router that is "out of touch" with changes to the routing tables. It would be as if a driver on the modern U.S. freeway system were using a map from 1964.

The Vanishing IP Address Space

1. As corporate network use and the Internet grew in popularity, vast numbers of IP address were purchased including large sections of Class A addresses. As this subject will be covered later in this chapter, a single Class A network address includes a huge number of hosts per network. Owners of these Class A networks, even now, possess a large storehouse of unused host addresses.

2. In addition to the program the text mentions about a "brisk trade" in IP addresses, there is a voluntary "buy back" program run by ICANN to reclaim portions of the above-mentioned addresses that have never been used.

3. The text mentions how many companies rent rather than buy their addresses from ISPs. You might mention to your class that they are also part of the group that rents IP addresses. Every time we go on the Internet, our ISP temporarily assigns us an IP address from a pool they own, allowing us to have an address that is routable on the Internet and saving us the expense of buying one of our own.

4. One thing to mention about private networks is that, if your network will never go on the internet, you can use any address at all. However, if you try to use those addresses to surf the web, you will find they are already owned. Highly recommend to your students that they use a private IP addressing scheme, even for their home or lab networks.

Quick Quiz 2

1. A(n) ____________________ is a network address that all hosts on a network must read.
   Answer: broadcast address

2. ____________________ permits existing addresses to be combined into larger routing domains with more host addresses than simple addition of the normal number of host addresses for each domain would supply.
   Answer: Classless Inter-Domain Routing (CIDR)

3. True or False: RFC 1918 reserves three ranges of IP addresses for private use - a single Class A, 16 Class Bs, and 56 Class Cs.
   Answer: False

4. True or False: Originally, when IP addresses were assigned for public use, they were assigned on a per-network basis.
   Answer: True

Understanding Basic Binary Arithmetic

1. Point out that students must master four different kinds of binary calculations:
   - Converting binary to decimal
   - Converting decimal to binary
   - Understanding how setting increasing numbers of high-order bits to 1 in eight-bit binary numbers corresponds to specific decimal numbers
   - Understanding how setting increasing low-order bits to 1 in eight-bit binary numbers corresponds to specific decimal numbers

Converting Decimal to Binary

1. Your students will need to learn to manually calculate conversions between decimal and binary. Nevertheless, the scientific calculator found on their Windows XP machine under Programs -> Accessories can also do the conversion.

2. The text takes the students through a series of exercises to help them do these conversions. This works well for some students but others seem to get better results from just memorizing the values of each bit.

3. Writing the value of each of the eight bits on the board and keeping it there will give your students a visual reference while doing these conversions:
   a. 128 64 32 16 8 4 2 1 If all these were “on”, in binary it would be 11111111 and would total 255.

4. Here is a simple exercise to go over with the class. Let’s say we are trying to convert “125” into binary. We know by looking at the chart that the leading digit must be 0 and that subtracting 128 from 255 gives us 127. We need only subtract the binary digit with the “2” value to achieve 125. You end up with a binary value of 01111101.

Converting Binary to Decimal

1. This is so much easier if you know the value of each bit.

High-Order Bit Patterns

1. The numbers block in this section is a great way to remember binary to decimal relations.

Low-Order Bit Patterns

1. Either method is valid in terms of doing conversions. People seem to have different styles in dealing with these sorts of problems. Encourage your students to use the method that works best for them.

IP Networks, Subnets, and Masks

1. Introduce the concept of a subnet mask. Note that this topic can be particularly challenging for students.

IP Subnets and Supernets

1. The key to supernetting is to use subnets that are contiguous, that is, their ranges are numerically (in binary) “next to” each other. This allows two or more subnets to be combined. Typically, you will see Class C addresses most often supernetted.

Calculating Subnet Masks

1. Constant-length subnet masks (CLSM) are what most people think of when they think of subnet masks. In a production environment, you are more likely to hear the word “subnet” thrown around than “CLSM”.

2. A variable-length subnet mask (VLSM) is subnetting across a class boundary. Basically, it is subnetting a subnet. The protocol used by routers in these network environments must support extended network prefix information.

3. Generally when you create a particular subnet, you are trying to satisfy a set of requirements for a number of networks in your environment and a number of valid host addresses within each network. Do not forget to figure in potential growth. If a network designer calculates subnets to only satisfy the current requirements of the organization, these networks will not have the ability to expand when more users are added.

Calculating Supernets

1. Supernetting is actually a form of Classless Inter-Domain Routing (CIDR) as will be seen in the next major section.

2. The one restriction in supernetting is one of boundaries. The value of the 3rd octet of the lower address must be divisible by 2 in order to combine two subnets, divisible by 4 to combine four subnets, and so on.

Classless Inter-Domain Routing (CIDR)

1. As mentioned earlier, all addresses in a CIDR address must be contiguous.

2. To clarify this point, address aggregation is where a single address will represent multiple networks in a routing table. These "multiple addresses" are combined by CIDR to appear as a single network.

3. Unlike CLSM where you lose a significant number of host addresses by subdividing a network, in CIDR you can use the entire range of addresses available. For example, with the network address of 224.127.97.8/20, the "/20" (called the "CIDR notation") is interpreted to mean that the network portion of the address is the first 20 bits leaving the remaining 12 bits for host addresses. That results in 4094 host addresses available in this network.

4. To relate this to "classful" networks, a standard Class A network uses 8 bits for the network portion of its address so it is a /8 address. A Class B network uses 16 bits for the network address so it is a /16 and a Class C network uses 24 bits for the network address so it is a /24. In CIDR, you can use any number of bits for the network address ignoring class limitations (leaving two bits available to support hosts, of course).

Public Versus Private IP Addresses

1. Most private networks including home labs and small office/home office (SOHO) networks use private IP addresses. It is very typical to have one computer act as the interface between the internal network and the internet, even on home networks. This computer will have a NIC configured to a private IP address on the same subnet as the other internal computers and a dial-up, xDSL (Digital Subscriber Service), or cable modem creating the link to the Internet. The modem interface will be assigned a public IP address from the ISP and this computer must be configured to share its Internet connection with the other computers on the private network.

2. The text mentions that IP Security (IPSec), a form of secure, encrypted information transfer, cannot be used in conjunction with NAT since the private address cannot be translated and thus routed to the Internet. This is true but there is a way around this. Instead of establishing an IPSec tunnel directly from computer to computer, establish it from perimeter device that does the NAT translation to the other computer outside the network. This is typically done firewall to firewall and would only apply to this particular link. All other standard traffic to and from the Internet would go through NAT translation.

3. The text mentions the issue of lag time in changing name to address resolution in this section. A practical example would be eBay or Amazon. Imagine how much revenue each one of them would lose if they had to re-establish name to address resolution, waiting up to 72 hours to be able to send and receive traffic on the Internet.

4. This would be a good time to review RFCs 2709 and 3104 with your students.

Quick Quiz 3

1. A(n) ____________________ is a special bit pattern that “blocks off” the network portion of an IP address with an all-ones pattern.
   Answer: subnet mask

2. The simplest form of subnet masking uses a technique called ____________________, in which each subnet includes the same number of stations and represents a simple division of the address space made available by subnetting into multiple equal segments.
   Answer: constant-length subnet masking (CLSM)

3. True or False: CIDR addresses are commonly applied to Class A addresses.
   Answer: False

4. A(n) ____________________ is a device that interconnects multiple IP networks or subnets.
   Answer: IP gateway

Managing Access to IP Address Information

1. Although it is valid to use NAT as part of your network security strategy, it is generally recommended that multiple methods of security be employed. Reverse proxying would be part of a layered security approach.

Obtaining Public IP Addresses

1. To emphasize what the text already mentions, it is extremely common for organizations and individuals alike to lease their public addresses from an ISP rather than purchase them.

IP Addressing Schemes

1. In this section, you will discuss the need for IP addressing schemes, and how to create and document one.

The Network Space

1. There are a number of critical factors that typically constrain IP addressing schemes, and we look at these in two groups. The first group of constraints determines the number and size of networks. These are:
   - Number of physical locations
   - Number of network devices at each location
   - Amount of broadcast traffic at each location
   - Availability of IP addresses

- Delay caused by routing from one network to another

2. The second group that helps us determine how to choose which IP addresses go where are these design objectives:
   - Minimize the size of the routing tables.
   - Minimize the time required for the network to "converge."
   - Maximize flexibility and facilitate management and troubleshooting.

The Host Space

1. The point of planning growth for networks was previously mentioned. You can re-emphasize it here. One of the important things that go along with an organized host space and network addressing scheme is accurate documentation of the network. While this is not a requirement of the class, it would be helpful for the students to see the relationship between having a logical and organized network and being able to document it.

Teaching Tip	For more information on documentation tools, see: http://www.more.net/technical/netserv/diagrams/index.html

Quick Quiz 4

1. ____________________ permits the proxy server to front for servers inside the boundary by advertising only the proxy server's address to the outside world, and then forwarding only legitimate requests for service to internal servers for further processing.
   Answer: Reverse proxying

2. Because all devices accessible to the Internet must have public IP addresses, changing providers often means going through a tedious exercise called ____________________.
   Answer: IP renumbering

3. Switches make their decisions with specialized hardware known as ____________________.
   Answer: Application Specific Integrated Circuits (ASICs)

4. True or False: The time it takes to route from one network to another is affected by the size of the routing table.
   Answer: True

Class Discussion Topics

1. Have the class discuss why a large corporation that bought a Class A network 20 years ago would be reluctant to sell back even a portion of their unused host addresses to help conserve overall resources.

2. Have the class discuss the pros and cons of using a subnet mask calculator versus manually calculating subnet masks.

3. Have the class discuss the benefits of using reverse proxying from a security standpoint. Are additional methods required to protect the network from external attack? Why or why not? If more security methods are required, which ones would they choose?

Additional Projects

1. Assign the class the task of briefly documenting the classroom LAN including PCs, switches, routers, and printers, making sure they assign an IP addressing scheme. It does not have to be detailed or complicated. The class can break up into small groups for this project.

2. Have the class open a command prompt on their computers and type in: “ipconfig/all”. Have them identify the class of address used, the subnet mask, the default gateway address and the DNS server address(es). If your class LAN uses a customized subnet mask, how many bits were borrowed? Have them calculate the number of networks and hosts supported by this network address configuration.

Additional Resources

1. For a good essay on IP multicast as well as the OSI model in general, go to http://ntrg.cs.tcd.ie/undergrad/4ba2/multicast/.

2. This one is a little basic but sometimes it helps to have a different way to define core concepts. The article: “What is an IP Address?” can be found at http://www.howstuffworks.com. Just do a search for the title of the article once on the site.

3. For the resource: “Connected: An Internet Encyclopedia”, the students can go to http://www.freesoft.org. Then they can click on the title given above and use the various links to find information on the various subjects covered in this chapter.

4. A very good article on NAT can be found at: http://www.cisco.com/en/US/tech/tk648/tk361/technologies_tech_note09186a0080094831.shtml

5. A review of the relevant Request for Comment documents will be helpful. The students should be able to look up this information with a search engine or by entering the RFC numbers in the browser address window:

NAT and IPSec	RFC 2709 and 3104
Subnets	RFC 950
Private Addresses	RFC 1918
VLSN	RFC 1878
CIDR	RFC 1517, 1518, and 1519
Multicast extensions to OSPF (open shortest path first)	RFC 1584

Key Terms

- **baseline performance—**The average performance of a computer under normal working conditions.
- **anycast address** —A type of address in IPv6, an anycast address is an ordinary address that can be assigned to more than one host or interface. Packets pointed to an anycast address are delivered to the holder of that address nearest to the sender in terms of routing distance. An anycast address does not apply to IPv4.
- **address masquerading** — A method of mapping many internal (i.e., private), nonroutable addresses to a single external (i.e., public) IP address for the purpose of sharing a single Internet connection (also referred to as “address hiding”).
- **Application Specific Integrated Circuit (ASIC)** — A special-purpose form of integrated circuit. An ASIC provides a way to implement specific programming logic directly into chip form, thereby also providing the fastest possible execution of such programming logic when processing data. ASICs are what make it possible for high-speed, high-volume routers to perform complex address recognition and management functions that can keep up with data volumes and time-sensitive processing needs.
- **broadcast address** —The all-ones address for a network or subnet, this address provides a way to send the same information to all interfaces on a network.
- **Classless Inter-Domain Routing (CIDR)** —A form of subnet masking that does away with placing network and host address portions precisely on octet boundaries, but instead uses the /*n* prefix notation, in which *n* indicates the number of bits in the network portion of whatever address is presented.
- **constant-length subnet masking (CLSM)** — An IP subnetting scheme in which all subnets use the same size subnet mask, which therefore divides the subnetted address space into a fixed number of equal-size subnets.
- **domain name** —A symbolic name for a TCP/IP network resource; the Domain Name System (DNS) translates such names into numeric IP addresses so outbound traffic may be addressed properly.
- **Domain Name System (DNS)** —The TCP/IP Application layer protocol and service that manages an Internet-wide distributed database of symbolic domain names and numeric IP addresses so users can ask for resources by name, and get those names translated into the correct numeric IP addresses.
- **dot quad** — *See* dotted decimal notation.

- **dotted decimal notation** — The name for the format used to denote numeric IP addresses, such as 172.16.1.7, wherein four numbers are separated by periods (dots).
- **end-to-end connection** — A network connection in which the original sending and receiving IP addresses may not be altered, and where a communications connection extends all the way from sender to receiver while that connection remains active.
- **extended network prefix** —The portion of an IP address that represents the sum of the network portion of the address, plus the number of bits used for subnetting that network address. A Class B address with a three-bit subnetting scheme would have an extended network prefix of /19, 16 bits for the default network portion, plus three bits for the subnetting portion of that address, with a corresponding subnet mask of 255.255.224.0.
- **firewall** —A network boundary device that sits between the public and private sides of a network, and provides a variety of screening and inspection services to ensure that only safe, authorized traffic flows from outside to inside (used in the sense of a barrier designed specifically to block the spread of fire in houses or cars).
- **hop** — A single transfer of data from one network to another, through some kind of networking device. Router-to-router transfers are often called hops. The number of hops often provides a rough metric of the distance between a sender's network and a receiver's network. The number of routers that a packet must cross, or the number of routers that a packet crosses, represents the hop count from the source network to the target network.
- **host portion** — The rightmost bits in an IP address, allocated to identify hosts on a supernetwork, network, or subnetwork. **Internet Assigned Numbers Authority (IANA)** — The arm of the ISOC originally responsible for registering domain names and allocating public IP addresses. This job is now the responsibility of ICANN.
- **Internet Service Provider (ISP)** — An organization that provides Internet access to individuals or organizations as a primary line of business. Currently, ISPs are the source for public IP addresses for most organizations seeking Internet access.
- **IP gateway** —TCP/IP terminology for a router that provides access to resources outside the local subnet network address. (A default gateway is the name given to the TCP/IP configuration entry for clients that identifies the router they must use to send data outside their local subnet areas.)
- **IP renumbering**—The process of replacing one set of numeric IP addresses with another set of numeric IP addresses because of a change in ISPs, or an address reassignment.
- **layer-3 switch** — A type of networking device that combines hub, router, and network management functions within a single box. Layer-3 switches make it possible to create and manage multiple virtual subnets in a single device, while offering extremely high bandwidth to individual connections between pairs of devices attached to that device.
- **loopback**—An address that points directly back to the sender. In IPv4, the ClassA domain 127.0.0.0 (or 127.0.0.1 for a specific machine address) is reserved for loopback. In IPv6, there is a single loopback address, written "::1" (all 0s, except for that last bit, which is 1). By passing traffic down through the TCP/IP stack, then back up again, the loopback address can be used to test a computer's TCP/IP software.
- **Media Access Control (MAC) layer** —A sub-layer of the Data Link layer. This layer is part of the Media Access Control definition, in which network access methods, such as Ethernet and token ring, apply.
- **multicast address** — One of a block of addresses reserved for use in sending the same message to multiple interfaces or nodes. Members of a community of interest subscribe to a multicast address in order to receive router updates, streaming data (video, audio,

teleconferencing), and so on. In IPv4, the Class D block of addresses is reserved for multicast. In IPv6, all multicast addresses begin with 0xFF. ICANN, with the help of IANA, manages all such address adjustments.

- **network address** —That portion of an IP address that consists of the network prefix for that address; an extended network prefix also includes any subnetting bits. All bits that belong to the extended network prefix show up as 1s in the corresponding subnet mask for that network.
- **Network Address Translation (NAT)** — A special type of networking software that manages network connections on behalf of multiple clients on an internal network and translates the source address for all outbound traffic from the original source to the address of the outbound network interface. NAT software also manages forwarding replies to all outgoing traffic back to its original sender.NAT software is often used to allow clients using private IP addresses to access the Internet.
- **network portion —** The leftmost octets or bits in a numeric IP address, the network portion of an IP address identifies the network and subnet portions of that address. The value assigned to the prefix number identifies the number of bits in the network portion of any IP address. (For example, 10.0.0.0/8 indicates that the first eight bits of the address are the network portion for the public Class A IP address.)
- **network prefix** —That portion of an IP address that corresponds to the network portion of the address; for example, the network prefix for a Class B address is /16 (meaning that the first 16 bits represent the network portion of the address, and 255.255.0.0 is the corresponding default subnet mask).
- **numeric address —** *See* numeric IP address.
- **numeric IP address** —An IP address expressed in dotted decimal or binary notation.
- **octet** —TCP/IP terminology for an eight-bit number; numeric IPv4 addresses consist of four octets.
- **organizationally unique identifier (OUI)** —A unique identifier assigned by IANA or ICANN that's used as the first three bytes of a NIC's MAC layer address to identify its maker or manufacturer. **physical numeric address** —Another term for MAC layer address (or MAC address).
- **private IP address** —Any of a series of Class A, B, and C IP addresses reserved by IANA for private use, documented in RFC 1918, and intended for uncontrolled private use in organizations. Private IP addresses may not be routed across the Internet because there is no guarantee that any such address is unique.
- **proxy server** —A special type of network boundary service that interposes itself between internal network addresses and external network addresses. For internal clients, a proxy server makes a connection to external resources on the client's behalf and provides address masquerading. For external clients, a proxy server presents internal resources to the public Internet as if they are present on the proxy server itself.
- **public IP address**—Any TCP/IP address allocated for the exclusive use of some particular organization, either by IANA or ICANN, or by an ISP to one of its clients.
- **Quality of Service (QoS)** — A specific level of service guarantee associated with Application layer protocols in which time-sensitivity requirements for data (such as voice or video) require that delays be controlled within definite guidelines to deliver viewable or audible data streams.
- **reverse proxying** —The technique whereby a proxy server presents an internal network resource (for example, a Web, e-mail, or FTP server) as if it were present on the proxy server itself so external clients can access internal network resources without seeing internal network IP address structures.

- **route aggregation**—A form of IP address analysis that permits routers to indicate general interest in a particular network prefix that represents the "common portion" of a series of IP network addresses, as a way of reducing the number of individual routing table entries that routers must manage.
- **subnet mask** —A special bit pattern that masks off the network portion of an IP address with all 1s.
- **subnetting—**The operation of using bits borrowed from the host portion of an IP address to extend and subdivide the address space that falls beneath the network portion of a range of IP addresses.
- **summary address** — A form of specialized IP network address that identifies the "common portion" of a series of IP network addresses used when route aggregation is in effect. This approach speeds routing behavior and decreases the number of entries necessary for routing tables.
- **supernetting** — The technique of borrowing bits from the network portion of an IP address and lending those bits to the host part, creating a larger address space for host addresses.
- **symbolic name**—A human-readable name for an Internet resource, such as *www.course.com* or *www. microsoft.com*. Also, a name used to represent a device instead of an address. For example, the name *serv1* could be a symbolic name for a device that uses the IP address 10.2.10.2.
- **variable-length subnet masking (VLSM)** —A subnetting scheme for IP addresses that permits containers of various sizes to be defined for a network prefix. The largest subnet defines the maximum container size, and any individual container in that address space may be further subdivided into multiple, smaller sub-containers (sometimes called sub-subnets).

Technical Notes for Hands-On Projects

The lab setup for Chapter 2 includes the following elements:

HANDS-ON PROJECT	NETWORK DEVICES REQUIRED	WORKSTATION OPERATING SYSTEM REQUIRED	OTHER RESOURCES REQUIRED
2 – 1		Windows XP Professional	The CD-ROM accompanying the text containing the IP Subnet Calculator Demo
2 – 2	Internet Connection	Windows XP Professional	
2 – 3	Internet Connection	Windows XP Professional	

2 – 4	Internet Connection	Windows XP Professional	
2 – 5		Windows XP Professional	IP Subnet Calculator Demo loaded onto Workstation
2 – 6		Windows XP Professional	IP Subnet Calculator Demo loaded onto Workstation