APPENDIX C

Answers to the "Do I Know This Already?" Quizzes

Chapter 1

1. D and F. Of the remaining answers, Ethernet defines both physical and data link protocols, PPP is a data link protocol, IP is a network layer protocol, and SMTP and HTTP are application layer protocols.

2. A and G. Of the remaining answers, IP is a network layer protocol, TCP and UDP are transport layer protocols, and SMTP and HTTP are application layer protocols.

3. B. Adjacent-layer interaction occurs on one computer, with two adjacent layers in the model. The higher layer requests services from the next lower layer, and the lower layer provides the services to the next higher layer.

4. B. Same-layer interaction occurs on multiple computers. The functions defined by that layer typically need to be accomplished by multiple computers—for example, the sender setting a sequence number for a segment, and the receiver acknowledging receipt of that segment. A single layer defines that process, but the implementation of that layer on multiple devices is required to accomplish the function.

5. A. Encapsulation is defined as the process of adding a header in front of data supplied by a higher layer (and possibly adding a trailer as well).

6. D. By convention, the term *frame* refers to the part of a network message that includes the data link header and trailer, with encapsulated data. The term *packet* omits the data link header and trailer, leaving the network layer header with its encapsulated data. The term *segment* omits the network layer header, leaving the transport layer header and its encapsulated data.

7. C. The network layer concerns itself with delivery of data over the complete end-to-end path. That requires a way to identify each device, using addresses, and the addresses must be logical addresses that are therefore not tied to the physical details of the network.

8. A. The OSI physical layer includes all standards that specify the shape of connectors, wiring in cabling, electrical details, and encoding that the electrical signals use to encode bits over a cable.

9. C and E. The layer names, from top to bottom, are application, presentation, session, transport, network, data link, and physical.

Chapter 2

1. A. The IEEE defines Ethernet LAN standards, with standard names that begin with 802.3, all of which happen to use cabling. The IEEE also defines wireless LAN standards, with standard names that begin with 802.11, which are separate standards from Ethernet.

2. C. The number before the word "BASE" defines the speed, in megabits per second (Mbps). 1000 Mbps equals 1 gigabit per second (1 Gbps). The "T" in the suffix implies twisted-pair or UTP cabling, so 1000BASE-T is the UTP-based Gigabit Ethernet standard name.

3. B. Crossover cables cross the wire at one node's transmit pin pair to the different pins used as the receive pins on the other device. For 10- and 100-Mbps Ethernet, the specific crossover cable wiring connects the pair at pins 1 and 2 on each end of the cable to pins 3 and 6 on the other end of the cable, respectively.

4. B, D, and E. Routers, wireless access point Ethernet ports, and PC NICs all send using pins 1 and 2, whereas hubs and LAN switches transmit on pins 3 and 6. Straight-through cables connect devices that use opposite pin pairs for sending, because the cable does not need to cross the pairs.

5. B. NICs (and switch ports) use the carrier sense multiple access with collision detection (CSMA/CD) algorithm to implement half-duplex logic. While CSMA/CD attempts to avoid collisions, it also notices when collisions do occur, with rules about how the Ethernet nodes should stop sending, wait, and try again later.

6. C. The 4-byte Ethernet FCS field, found in the Ethernet trailer, allows the receiving node to see what the sending node computed with a math formula that is a key part of the error-detection process. Note that Ethernet defines the process of detecting errors (error detection), but not error recovery.

7. B, C, and E. The preassigned universal MAC address, given to each Ethernet port when manufactured, breaks the address into two 3-byte halves. The first half is called the organizationally unique identifier (OUI), which the IEEE assigns to the company that builds the product as a unique hex number to be used only by that company.

8. C and D. Ethernet supports unicast addresses, which identify a single Ethernet node, and group addresses, which can be used to send one frame to multiple Ethernet nodes. The two types of group addresses are the *broadcast address* and *multicast addresses*.

Chapter 3

1. B. OSI Layer 1, the physical layer, defines how to physically transmit bits using some kind of physical medium, like a cable or radio waves. OSI Layer 2, the data link layer, defines protocols that control and manage the transmission, including the definition of data link frames, which include addresses, and fields to help perform error detection.

2. B. The four-wire circuit cable supplied by the telco connects to the device acting as the CSU/DSU. That can be an external CSU/DSU or a CSU/DSU integrated into a router serial interface card. LAN switches do not have serial interfaces, and router serial interfaces do not have transceivers.

3. C. Leased lines can run at various preset speeds. These include multiples of 64 kbps, up through 24 times 64 kbps. The speeds can also be multiples of T1 speed (1.544 Mbps), up to 28 times that speed.

4. B. The standard HDLC header does not include a Type field, which identifies the type of packet encapsulated inside the HDLC frame.

5. B and D. The physical installation uses a model in which each router uses a physical Ethernet link to connect to some SP device in an SP facility called a point of presence (POP). The Ethernet link does not span from each customer device to the other. From a data link perspective, both routers use the same Ethernet standard header and trailer used on LANs; HDLC does not matter on these Ethernet WAN links.

6. B and C. Leased lines transmit data at the same speed in both directions, making it a symmetric service. DSL and cable Internet offer asymmetric speeds, with a faster downstream speed. BGP is a routing protocol and is not an Internet access technology.

7. C. With DSL, the requirements on the phone's wiring are unchanged. The phone can connect to any working telephone jack, as if the DSL modem and router did not exist.

Chapter 4

1. A and C. The network layer defines logical addressing, in contrast to physical addressing. The logical address structure allows easy grouping of addresses, which makes routing more efficient. Path selection refers to the process of choosing the best routes to use in the network. Physical addressing and arbitration typically are data link layer functions, and error recovery typically is a transport layer function.

2. C and E. Because PC1 and PC2 are separated by routers, the PCs must rely on the routers to use routing logic. With routing, the router discards the data link header of an incoming frame, making a decision of where to forward the packet deencapsulated from the frame. As a result, the largest entity that passes from PC1 to PC2 is the IPv4 packet (or IPv6 packet), which goes by the more general name of L3PDU.

3. B. 224.1.1.1 is a Class D address.

4. D. The first octet of Class A addresses ranges from 1 to 126, inclusive; Class B, 128 to 191, inclusive; and Class C, 192 to 223 inclusive. 127 is technically in the Class A range, but it is a reserved address used as a loopback.

5. D and F. Without any subnetting in use, all addresses in the same network as 10.1.1.1—all addresses in Class A network 10.0.0.0—must be on the same LAN. Addresses separated from that network by some router cannot be in network 10.0.0.0. So, the two correct answers are the only two answers that list a valid unicast IP address that is not in network 10.0.0.0.

6. A. PC1 will send an Ethernet frame to Router1, with PC1's MAC address as the

6. A. PC1 will send an Ethernet frame to Router1, with PC1's MAC address as the source address and Router1's MAC address as the destination address. Router1 will remove the encapsulated IP packet from that Ethernet frame, discarding the frame header and trailer. Router1 will forward the IP packet by first encapsulating it inside an HDLC frame, but Router1 will not encapsulate the Ethernet frame in the HDLC frame, but rather the IP packet. Router2 will deencapsulate the IP packet from the HDLC frame and forward it onto the Ethernet LAN, adding a new Ethernet header and trailer, but this header will differ. It will list Router2's MAC address as the source address and PC2's MAC address as the destination address.

7. C. Routers compare the packet's destination IP address to the router's IP routing table, making a match and using the forwarding instructions in the matched route to forward the IP packet.

8. B and C. IPv4 hosts generally use basic two-branch logic. To send an IP packet to another host on the same IP network or subnet that is on the same LAN, the sender sends the IP packet directly to that host. Otherwise, the sender sends the packet to its default router (also called the default gateway).

9. A and C. Routers do all the actions listed in all four answers. However, the routing protocol does the functions in the two listed answers. Independent of the routing protocol, a router learns routes for IP subnets and IP networks directly connected to its interfaces. Routers also forward (route) IP packets, but that process is called IP routing, or IP forwarding, and is an independent process compared to the work of a routing protocol.

10. C. Address Resolution Protocol (ARP) does allow PC1 to learn information, but the information is not stored on a server. The **ping** command does let the user at PC1 learn whether packets can flow in the network, but it again does not use a server. With the Domain Name System (DNS), PC1 acts as a DNS client, relying on a DNS server to respond with information about the IP addresses that match a given host name.

Chapter 5

1. D. Three of the answers define functions of TCP: error recovery, flow control, and segmentation. The correct answer defines a function not done by TCP nor UDP, but instead done by application layer protocols.

2. D and E. Many headers include a field that identifies the next header that follows inside a message. Ethernet uses the Ethernet Type field, and the IP header uses the Protocol field. The TCP and UDP headers identify the application that should receive the data that follows the TCP or UDP header by using the port number field in the TCP and UDP headers, respectively.

3. A, B, C, and F. IP, not TCP, defines routing. Many other protocols define encryption, but TCP does not. The correct answers simply list various TCP features.

4. C. TCP, not UDP, performs windowing, error recovery, and ordered data transfer. Neither performs routing or encryption.

5. C and F. The terms *packet* and *L3PDU* refer to the header plus data encapsulated by Layer 3. *Frame* and *L2PDU* refer to the header (and trailer), plus the data encapsulated by Layer 2. *Segment* and *L4PDU* refer to header and data encapsulated by the transport layer protocol.

6. B. Note that the host name is all the text between the // and the /. The text before the // identifies the application layer protocol, and the text after the / represents the name of the web page.

7. A and D. VoIP flows need better delay, jitter, and loss, with better meaning less delay, jitter, and loss, as compared with all data applications. VoIP typically requires less bandwidth than data applications.

Chapter 6

1. A. A switch compares the destination address to the MAC address table. If a matching entry is found, the switch knows out which interface to forward the frame. If no matching entry is found, the switch floods the frame.

2. C. A switch floods broadcast frames, multicast frames (if no multicast optimizations are enabled), and unknown unicast destination frames (frames whose destination MAC address is not in the MAC address table).

3. A. A switch floods broadcast frames, multicast frames (if no multicast optimizations are enabled), and unknown unicast destination frames (frames whose destination MAC address is not in the MAC address table).

4. B. Switches need to learn the location of each MAC address used in the LAN relative to that local switch. When a switch sends a frame, the source MAC identifies the sender. The interface in which the frame arrives identifies the local switch interface closest to that node in the LAN topology.

5. A and B. When the frame sent by PC3 arrives at the switch, the switch has learned a MAC address table entry for only 1111.1111.1111, PC1's MAC address. PC3's frame, addressed to 2222.2222.2222, is flooded, which means it is forwarded out all interfaces except for the interface on which the frame arrived.

6. A. A collision domain contains all devices whose frames could collide with frames sent by all the other devices in the domain. Bridges, switches, and routers separate or segment a LAN into multiple collision domains, whereas hubs and repeaters do not.

7. A, B, and C. A broadcast domain contains all devices whose sent broadcast frames should be delivered to all the other devices in the domain. Hubs, repeaters, bridges, and switches do not separate or segment a LAN into multiple broadcast domains, whereas routers do.

8. B and D. The IEEE Ethernet standards support 100-meter links when using UTP cabling. Most standards that use fiber-optic cabling, like the standards in the two correct answers, use lengths longer than 100 meters.

9. D. Full-duplex can be used on any Ethernet link between a switch and a PC. The IEEE autonegotiation process states that the two endpoints on a link should choose the best speed and the best duplex that both support, with full-duplex being better than half-duplex.

Chapter 7

1. A and B. The command in the question is an EXEC command that happens to require only user mode access. As an EXEC command, you cannot use the command in configuration mode, unless you preface the command with "do" (for example, **do show mac address-table**). As such, you can use this command in both user mode and enable mode.

2. B. The command referenced in the question, the **reload** command, is an EXEC command that happens to require privileged mode, also known as enable mode. As an EXEC command, you cannot use the command in configuration mode, unless you preface the command with **do** (for example, **do reload**). This command is not available in user mode.

3. B. SSH provides a secure remote login option, encrypting all data flows, including password exchanges. Telnet sends all data (including passwords) as clear text.

4. A. Switches (and routers) keep the currently used configuration in RAM, using NVRAM to store the configuration file that is loaded when the switch (or router) next loads the IOS.

5. F. The startup config file is in NVRAM, and the running config file is in RAM.

6. B and C. The **exit** command moves the user one config mode backward, toward global configuration mode, or if already in global configuration mode, it moves the user back to enable mode. From console mode, it moves the user back to global configuration mode. The **end** command and the Ctrl-Z key sequence both move the user back to enable mode regardless of the current configuration submode.

Chapter 8

1. B. If both commands are configured, IOS accepts only the password as configured in the **enable secret** command.

2. B and C. SSH requires the use of usernames in addition to a password. Using the **username** global command would be one way to define usernames (and matching passwords) to support SSH. The vty lines would also need to be configured to require the use of usernames, with the **login local** vty subcommand being one such option. The **transport input ssh** command could be part of a meaningful configuration, but it is not a global configuration command (as claimed in one wrong answer). Likewise, one answer refers to the **username** command as a command in vty config mode, which is also the wrong mode.

3. B. The first nonblank character after the **banner login** phrase is interpreted as the beginning delimiter character. In this case, it's the letter *t*. So, the second letter *t*—the first letter in "the"—is interpreted as the ending delimiter. The resulting login banner is the text between these two *t*s—namely, "his is."

4. B. The setting for the maximum number of MAC addresses has a default of 1, so the **switchport port-security maximum** command does not have to be configured. With sticky learning, you do not need to predefine the specific MAC addresses either. However, you must enable port security, which requires the **switchport port-security** interface subcommand.

5. A, D, and F. To allow access through Telnet, the switch must have password security enabled, at a minimum using the **password** vty line configuration subcommand. Additionally, the switch needs an IP address (configured under one VLAN interface) and a default gateway when the switch needs to communicate with hosts in a different subnet.

6. F. Cisco switches do not have a command to disable autonegotiation of speed and duplex. Instead, a switch port that has both **speed** and **duplex** configured disables autonegotiation.

7. E. Cisco switches can be configured for speed (with the **speed** command) and duplex (with the **duplex** command) in interface configuration mode.

Chapter 9

1. B. A VLAN is a set of devices in the same Layer 2 broadcast domain. A subnet often includes the exact same set of devices, but it is a Layer 3 concept. A collision domain refers to a set of Ethernet devices, but with different rules than VLAN rules for determining which devices are in the same collision domain.

2. D. Although a subnet and a VLAN are not equivalent concepts, the devices in one VLAN are typically in the same IP subnet and vice versa.

3. B. 802.1Q defines a 4-byte header, inserted after the original frame's destination and source MAC address fields. The insertion of this header does not change the original frame's source or destination address. The header itself holds a 12-bit VLAN ID field, which identifies the VLAN associated with the frame.

4. A. 802.1Q defines the native VLAN as one designated VLAN on a trunk for which the devices choose to not add an 802.1Q header for frames in that VLAN. The switches can set the native VLAN to any VLAN ID, but the switches should agree. The default native VLAN is VLAN 1. Note that only one such native VLAN is allowed on any one trunk; otherwise, that VLAN associated with untagged frames could not be discerned by the receiving switch.

5. A and C. The **dynamic auto** setting means that the switch can negotiate trunking, but it can only respond to negotiation messages, and it cannot initiate the negotiation process. So, the other switch must be configured to trunk or to initiate the negotiation process (based on being configured with the **dynamic desirable** option).

6. A and B. The default VTP setting of VTP transparent mode means that the switch can configure VLANs, so the VLAN is configured. Additionally, the VLAN configuration details, including the VLAN name, show up as part of the running-config file.

7. B and C. The **show interfaces switchport** command lists both the administrative and operational status of each port. When a switch considers a port to be trunking, this command lists an operational trunking state of "trunk." The **show interfaces trunk** command lists a set of interfaces: the interfaces that are currently operating as trunks. So, both these commands identify interfaces that are operational trunks.

Chapter 10

1. E and F. CDP discovers information about neighbors. **show cdp** gives you several options that display more or less information, depending on the parameters used.

2. E and F. The **show cdp neighbors** command lists one line of output per neighbor. However, it does list the platform information of the neighbor, which typically includes the hardware model number. The **show cdp entry Hannah** command lists a group of messages about the neighboring router, including more detail about the hardware model and the IOS version.

3. A, B, and D. The disabled state in the **show interfaces status** command is the same as an "administratively down and down" state shown in the **show interfaces** command. The interface must be in a connected state (per the **show interfaces status** command) before the switch can send frames out the interface.

4. A and D. SW2 has effectively disabled IEEE standard autonegotiation by configuring both speed and duplex. However, Cisco switches can detect the speed used by the other device, even with autonegotiation turned off. Also, at 1 Gbps, the IEEE autonegotiation standard says to use full-duplex if the duplex setting cannot be negotiated, so both ends use 1 Gbps, full-duplex.

5. B and D. The **show interfaces** command lists the actual speed and duplex setting, but it does not imply anything about how the settings were configured or negotiated. The **show interfaces status** command lists a prefix of **a-** in front of the speed and duplex setting to imply that the setting was autonegotiated, leaving off this prefix if the setting was configured.

6. B and C. IOS adds MAC addresses configured by the port security feature as static MAC addresses, so they do not show up in the output of the **show mac address-table dynamic** command. **show mac address-table port-security** is not a valid command.

7. A and C. The **show mac address-table** command lists all entries in the switch's MAC address table, including dynamically learned and statically defined addresses. From left to right, the output lists the VLAN ID, the MAC address, the type (static or dynamic), and the associated ports (out which frames sent to the listed address will be forwarded by this switch).

8. D. The **show interface switchport** command lists a switch's configured value of the **switchport mode** command as the administrative mode, telling us that SW1 uses the **switchport mode trunk** configuration command. The operation state shows the current working state (trunk). Of the four answers, one incorrect answer claims that SW2's operational state must also be "trunk"; however, one switch can use trunking, while the other does not. Another answer claims that the administrative mode must

be "trunk," which basically means that the command **switchport mode trunk** must be used. However, SW2 has a few other configuration options that would work, like the **switchport mode dynamic desirable** and **switchport mode dynamic auto** commands. The correct answer refers to one of these commands that would work correctly.

Chapter 11

1. B and D. The general rule to determine whether two devices' interfaces should be in the same subnet is whether the two interfaces are separated from each other by a router. To provide a way for hosts in one VLAN to send data to hosts outside that VLAN, a local router must connect its LAN interface to the same VLAN as the hosts, and have an address in the same subnet as the hosts. All the hosts in that same VLAN on the same switch would not be separated from each other by a router, so these hosts would also be in the same subnet. However, another PC, connected to the same switch but in a different VLAN, will require its packets to flow through a router to reach Host A, so Host A's IP address would need to be in a different subnet compared to this new host.

2. D. By definition, two address values in every IPv4 subnet cannot be used as host IPv4 addresses: the first (lowest) numeric value in the subnet for the subnet ID, and the last (highest) numeric value in the subnet for the subnet broadcast address.

3. B and C. At least 7 subnet bits are needed, because $2^6 = 64$, so 6 subnet bits could not number 100 different subnets. Seven subnet bits could, because $2^7 = 128 => 100$. Similarly, 6 host bits is not enough, because $2^6 - 2 = 62$, but 7 host bits is enough, because $2^7 - 2 = 126 => 100$.

   The number of network, subnet, and host bits must total 32 bits, making one of the answers incorrect. The answer with 8 network bits cannot be correct because the question states that a Class B network is used, so the number of network bits must always be 16. The two correct answers have 16 network bits (required because the question states the use of a Class B network), and at least 7 subnet and host bits each.

4. A and C. The private IPv4 networks, defined by RFC 1918, are Class A network 10.0.0.0, the 16 Class B networks from 172.16.0.0 to 172.31.0.0, and the 256 Class C networks that begin with 192.168.

5. A, D, and E. The private IPv4 networks, defined by RFC 1918, are Class A network 10.0.0.0, the 16 Class B networks from 172.16.0.0 to 172.31.0.0, and the 256 Class C networks that begin with 192.168. The three correct answers are from the public IP network range, and none are reserved values.

6. A and C. An unsubnetted Class A, B, or C network has two parts: the network and host parts.

7. B. An unsubnetted Class A, B, or C network has two parts: the network and host parts. To perform subnetting, the engineer creates a new subnet part by borrowing host bits, shrinking the number of host bits. The subnet part of the address structure exists only after the engineer chooses a nondefault mask. The network part remains a constant size.

8. C and D. Subnet ID (short for subnet identifier), subnet address, and subnet number are all synonyms and refer to the number that identifies the subnet. The actual value is a dotted-decimal number, so the term *subnet name* does not apply. The term *subnet broadcast*, a synonym for the subnet broadcast address, refers to the last (highest) numeric value in a subnet.

Chapter 12

1. B and C. Class A networks have a first octet in the range of 1–126, inclusive, and their network IDs have a 0 in the last three octets. 130.0.0.0 is actually a Class B network (first octet range 128–191, inclusive). All addresses that begin with 127 are reserved, so 127.0.0.0 is not a Class A network.

2. E. Class B networks all begin with values between 128 and 191, inclusive, in their first octets. The network ID has any value in the 128–191 range in the first octet, and any value from 0–255 inclusive in the second octet, with decimal 0s in the final two octets. Two of the answers show a 255 in the second octet, which is acceptable. Two of the answers show a 0 in the second octet, which is also acceptable.

3. B and D. The first octet (172) is in the range of values for Class B addresses (128–191). As a result, the network ID can be formed by copying the first two octets (172.16) and writing 0s for the last two octets (172.16.0.0). The default mask for all Class B networks is 255.255.0.0, and the number of host bits in all unsubnetted Class B networks is 16.

4. A and C. The first octet (192) is in the range of values for Class C addresses (192–223). As a result, the network ID can be formed by copying the first three octets (192.168.6) and writing 0 for the last octet (192.168.6.0). The default mask for all Class C networks is 255.255.255.0, and the number of host bits in all unsubnetted Class C networks is 8.

5. D. To find the network broadcast address, first determine the class, and then determine the number of host octets. At that point, convert the host octets to 255 to create the network broadcast address. In this case, 10.1.255.255 is in a Class A network, with the last three octets as host octets, for a network broadcast address of 10.255.255.255. For 192.168.255.1, it is a Class C address, with the last octet as the host part, for a network broadcast address of 192.168.255.255. Address 224.1.1.255 is a class D address, so it is not in any unicast IP network, so the question does not apply. For 172.30.255.255, it is a Class B address, with the last two octets as host octets, so the network broadcast address is 172.30.255.255.

6. B. To find the network ID, first determine the class, and then determine the number of host octets. At that point, convert the host octets to 0 to create the network ID. In this case, 10.1.0.0 is in a Class A network, with the last three octets as host octets, for a network ID of 10.0.0.0. For 192.168.1.0, it is a Class C address, with the last octet as the host part, for a network ID of 192.168.1.0. Address 127.0.0.0 looks like a Class A network ID, but it begins with a reserved value (127), so it is not in any Class A, B, or C network. 172.20.0.1 is a Class B address, with the last two octets as host octets, so the network ID is 172.20.0.0.

Chapter 13

1. C. Thinking about the conversion one octet at a time, the first two octets each convert to 8 binary 1s. 254 converts to 8-bit binary 11111110, and decimal 0 converts to 8-bit binary 00000000. So, the total number of binary 1s (which defines the prefix length) is 8+8+7+0 = /23.

2. B. Thinking about the conversion one octet at a time, the first three octets each convert to 8 binary 1s. 240 converts to 8-bit binary 11110000, so the total number of binary 1s (which defines the prefix length) is 8+8+8+4 = /28.

3. C. /24 is the equivalent of the mask that in binary has 24 binary 1s. To convert that to DDN format, write down all the binary 1s (24 in this case), followed by binary 0s for the remainder of the 32-bit mask. Then take 8 bits at a time, and convert from binary to decimal (or memorize the nine possible DDN mask octet values and their binary equivalents). Using the /24 mask in this question, the binary mask is 11111111 11111111 11111111 00000000. Each of the first three octets is all binary 1, so each converts to 255. The last octet, all binary 0s, converts to decimal 0, for a DDN mask of 255.255.255.0. See Appendix A for a decimal/binary conversion table.

4. B. /30 is the equivalent of the mask that in binary has 30 binary 1s. To convert that to DDN format, write down all the binary 1s (30 in this case), followed by binary 0s for the remainder of the 32-bit mask. Then take 8 bits at a time, and convert from binary to decimal (or memorize the nine possible DDN mask octet values and their binary equivalents). Using the /30 mask in this question, the binary mask is 11111111 11111111 11111111 11111100. Each of the first three octets is all binary 1, so each converts to 255. The last octet, 11111100, converts to 252, for a DDN mask of 255.255.255.252. See Appendix A for a decimal/binary conversion table.

5. C. The size of the network part is always either 8, 16, or 24 bits, based on whether it is Class A, B, or C, respectively. As a Class A address, N=8. The mask 255.255.255.0, converted to prefix format, is /24. The number of subnet bits is the difference between the prefix length (24) and N, so S=16 in this case. The size of the host part is a number that, when added to the prefix length (24), gives you 32, so H=8 in this case.

6. A. The size of the network part is always either 8, 16, or 24 bits, based on whether it is Class A, B, or C, respectively. As a Class C address, N=24. The number of subnet bits is the difference between the prefix length (27) and N, so S=3 in this case. The size of the host part is a number that, when added to the prefix length (27), gives you 32, so H=5 in this case.

7. B and D. Classless addressing rules define a two-part IP address structure: the prefix and the host part. The host part is defined the same way as with classful IP addressing rules. The classless address rules' prefix length is the length of the combined network and subnet parts when using classful IP addressing concepts. Mathematically, the prefix length is equal to the number of binary 1s in the mask. In this case, with a mask of 255.255.255.0, the prefix length is 24 bits. The host length is the number of bits added to 24 to total 32, for 8 host bits.

8. D. Classless addressing rules define a two-part IP address structure: the prefix and the host part. This logic ignores Class A, B, and C rules, and can be applied to the 32-bit IPv4 addresses from any address class. By ignoring Class A, B, and C rules, classless addressing ignores any distinction as to the network part of an IPv4 address.

9. A and B. The masks in binary define a number of binary 1s, and the number of binary 1s defines the length of the prefix (network + subnet) part. With a Class B network, the network part is 16 bits. To support 100 subnets, the subnet part must be at least 7 bits long. Six subnet bits would supply only 2^6 = 64 subnets, while 7 subnet bits supply 2^7 = 128 subnets. The /24 answer supplies 8 subnet bits, and the 255.255.255.252 answer supplies 14 subnet bits.

Chapter 14

1. D. When using classful IP addressing concepts as described in Chapter 13, "Analyzing Subnet Masks," addresses have three parts: network, subnet, and host. For addresses in a single classful network, the network parts must be identical for the numbers to be in the same network. For addresses in the same subnet, both the network and subnet parts must have identical values. The host part differs when comparing different addresses in the same subnet.

2. B and D. In any subnet, the subnet ID is the smallest number in the range, the subnet broadcast address is the largest number, and the usable IP addresses sit between them. All numbers in a subnet have identical binary values in the prefix part (classless view) and network + subnet part (classful view). To be the lowest number, the subnet ID must have the lowest possible binary value (all 0s) in the host part. To be the largest number, the broadcast address must have the highest possible binary value (all binary 1s) in the host part. The usable addresses do not include the subnet ID and subnet broadcast address, so the addresses in the range of usable IP address never have a value of all 0s or 1s in their host parts.

3. C. The mask converts to 255.255.255.0. To find the subnet ID, for each octet of the mask that is 255, you can copy the IP address's corresponding values. For mask octets of decimal 0, you can record a 0 in that octet of the subnet ID. As such, copy the 10.7.99 and write a 0 for the fourth octet, for a subnet ID of 10.7.99.0.

4. C. First, the resident subnet (the subnet ID of the subnet in which the address resides) must be numerically smaller than the IP address, which rules out one of the answers. The mask converts to 255.255.255.252. As such, you can copy the first three octets of the IP address because of their value of 255. For the fourth octet, the subnet ID value must be a multiple of 4, because 256 – 252 (mask) = 4. Those multiples include 96 and 100, and the right choice is the multiple closest to the IP address value in that octet (97) without going over. So, the correct subnet ID is 192.168.44.96.

5. C. The resident subnet ID in this case is 172.31.77.192. You can find the subnet broadcast address based on the subnet ID and mask using several methods. Following the decimal process in the book, the mask converts to 255.255.255.224, making the interesting octet be octet 4, with magic number 256 – 224 = 32. For the three octets where the mask = 255, copy the subnet ID (172.31.77). For the interesting octet, take the subnet ID value (192), add magic (32), and subtract 1, for 223. That makes the subnet broadcast address 172.31.77.223.

6. C. To answer this question, you need to find the range of addresses in the subnet, which typically then means you need to calculate the subnet ID and subnet broadcast address. With subnet ID/mask of 10.1.4.0/23, the mask converts to 255.255.254.0. To find the subnet broadcast address, following the decimal process described in this chapter, you can copy the subnet ID's first two octets because the mask's value is 255 in each octet. You write a 255 in the fourth octet because the mask has a 0 on the fourth octet. In octet 3, the interesting octet, add the magic number (2) to the subnet ID's value (4), minus 1, for a value of 2 + 4 – 1 = 5. (The magic number in this case is calculated as 256 – 254 = 2.) That makes the broadcast address 10.1.5.255. The last usable address is 1 less: 10.1.5.254. The range that includes the last 100 addresses is 10.1.5.155–10.1.5.254.

7. B. To answer this question, you do not actually need to calculate the subnet broadcast address, because you only need to know the low end of the range of addresses in the subnet. The first IP address in the subnet is 1 more than the subnet ID, or 192.168.9.97. The first 20 addresses then include 192.168.9.97–192.168.9.116.

Chapter 15

1. B and E. Cisco routers have an on/off switch, but Cisco switches generally do not.

2. A. Both switches and routers configure IP addresses, so the **ip address** *address mask* and **ip address dhcp** commands could be used on both routers and switches. The **interface vlan 1** command applies only to switches.

3. B and D. To route packets, a router interface must have an IP address assigned and be in an "up and up" interface state. For a serial link created in a lab, without using CSU/DSUs, one router must be configured with a **clock rate** command to the speed of the link. The **bandwidth** and **description** commands are not required to make a link operational.

4. C. If the first of the two status codes is "down," it typically means that a Layer 1 problem exists (for example, the physical cable is not connected to the interface).

5. C and E. The **show ip interface brief** command lists all the interface IPv4 addresses but none of the masks. The **show version** command lists none of the IP addresses and none of the masks. The other three commands list both the address and mask.

6. B. A router has one IPv4 address for each interface in use, whereas a LAN switch has a single IPv4 address that is just used for accessing the switch. The rest of the answers list configuration settings that use the same conventions on both routers and switches.

Chapter 16

1. B. PCs use two-option logic: send local packets (destined for hosts in the same subnet) directly, and send remote packets (destined for hosts in other subnets) to the default gateway/router. In this case, the PC's own IP address is 192.168.5.77, with mask 255.255.224.0, so it is in subnet 192.168.4.64/27. This subnet has a range of addresses from 192.168.4.64 to 192.168.4.95, including the subnet and broadcast address. As a result, the PC sends the packet to its default gateway. As for the other incorrect answers, if the **ping** command had used a host name, it would have first asked a DNS server to resolve the name. Also, the PC already has an IP address, so DHCP would not be needed.

2. A and C. The route defines the group of addresses represented by the route using the subnet ID and mask. The router can use those numbers to find the range of addresses that should be matched by this route. The other two answers list facts useful when forwarding packets that happen to match the route.

3. A and F. Of all the commands listed, only the two correct answers are syntactically correct router configuration commands. The command to enable 802.1Q trunking is **encapsulation dot1q** *vlan_id*.

4. C. With the **no ip subnet-zero** command configured, the router will not allow any interfaces to be configured with an IP address in the zero subnet. Of the listed answers, subnet 10.0.0.0 255.254.0.0 is a zero subnet, with a range of addresses from 10.0.0.1 to 10.1.255.254. The **ip address 10.1.2.2 255.254.0.0** command would be rejected.

5. C. The configuration of the Layer 3 switch's routing feature uses VLAN interfaces, with the interface number matching the VLAN ID. The matching connected routes, like all connected IP routes, will list the interfaces but not a next-hop IP address. The three connected routes will list the VLAN interfaces 1, 2, and 3, respectively.

6. C. The **ip route** command can refer to the IP address of the next-hop router or to the local router's interface. It also refers to the subnet ID and matching subnet mask, defining the range of addresses matched by the route.

7. A. The correct syntax lists a subnet number, then a subnet mask in dotted-decimal form, and then either an outgoing interface or a next-hop IP address.

Chapter 17

1. D and E. Both versions of RIP use distance vector logic, and EIGRP uses a different kind of logic, characterized either as advanced distance vector or a balanced hybrid.

2. C and D. Both versions of RIP use the same hop-count metric.

3. B, C, D, and E. Only the old RIP Version 1 (RIP-1) protocol does not support variable-length subnet masks (VLSM).

4. B. Link-state protocols do not exchange data that lists routes. They do list metric information, but it is per-interface information, and it is not tied to a subnet. Link-state protocols do require the SPF algorithm to take the varied pieces of information and create routes based on that information.

5. A and D. OSPF uses the SPF algorithm, conceived by a mathematician named Dijkstra.

6. B. The **network 10.0.0.0 0.255.255.255 area 0** command works, because it matches all interfaces whose first octet is 10. The rest of the commands match as follows: all addresses that end with 0.0.0 (wildcard mask 255.0.0.0); all addresses that begin with 10.0.0 (wildcard mask 0.0.0.255); and all addresses that begin 10.0 (wildcard mask 0.0.255.255).

7. A. The **network 0.0.0.0 255.255.255.255 area 0** command matches all IP addresses as a result of the 255.255.255.255 wildcard mask, so this command enables OSPF in Area 0 on all interfaces. The answer with wildcard mask 0.255.255.0 is illegal, because it represents more than one string of binary 0s separated by binary 1s. The answer with x's is syntactically incorrect. The answer with wildcard mask 255.0.0.0 means "Match all addresses whose last three octets are 0.0.0," so none of the three interfaces are matched.

8. A and E. Of the three wrong answers, two are real commands that simply do not happen to list the OSPF neighbors. **show ip ospf interface brief** lists interfaces on which OSPF is enabled, but does not list neighbors. **show ip interface** lists IPv4 details about interfaces, but none related to OSPF. One incorrect answer, **show ip neighbor**, is not a valid IOS command.

Chapter 18

1. B and D. The client sends a Discover message, with the server returning an Offer message. The client then sends a Request, with the server sending back the IP address in the Acknowledgment message.

2. E. Of the commands in the answers, three list syntactically correct commands, but one does not: **ip dhcp-server 10.1.10.1**. The answer with this command is incorrect. The **ip helper-address 10.1.10.1** command is indeed needed as an interface subcommand, but only on remote routers like the Boston router. This command is not needed for the DNS server, so while the **ip helper-address 10.1.10.2** command would be accepted, it does not help DNS or DHCP work. Finally, the **ip name-server 10.1.10.2** command on the Atlanta router works, but it simply lets users of the Atlanta router CLI make use of the DNS server; it has no effect on user DNS traffic.

3. B. The configuration puts the per-client settings in a DHCP pool: client IP address, mask, default router, and DNS server IP addresses. The pool also lists the length of time for which the client can keep the lease for the address. Only the exclude list (**ip dhcp exclude-address**) sits outside the DHCP pool.

4. B and D. A PC needs to know its own IP address, the associated mask, the IP address of the default router, and the IP addresses of the DNS servers.

5. A. A ping of a host's own IP address does not test whether the LAN is working, because the packet does not have to traverse the LAN. A ping that requires the packet to go from PC1 to the default gateway (R1) proves the LAN works, at least between PC1 and R1. The only answer that lists a command that causes a packet to need to cross the LAN from PC1 to R1 (although that process fails) is the **ping 10.1.1.1** command.

6. A, C, and E. The **tracert** (Microsoft operating systems) and **traceroute** (Cisco IOS Software) commands list the IP address of the intermediate routers and end host. The commands list the router's IP address closest to the host that issued the command.

7. C. The asterisk beside connection 2 identifies the connection number to which the **resume** command will connect the user if the **resume** command does not have any parameters.

Chapter 19

1. A. With 50 percent growth, the mask needs to define enough subnet bits to create 150 subnets. As a result, the mask needs at least 8 subnet bits (7 subnet bits supply 2^7, or 128, subnets, and 8 subnet bits supply 2^8, or 256, subnets). Similarly, the need for 50 percent growth in the size for the largest subnet means that the host part needs enough bits to number 750 hosts/subnet. Nine host bits are not enough ($2^9 - 2 = 510$), but 10 host bits supply 1022 hosts/subnet ($2^{10} - 2 = 1022$). With 16 network bits existing because of the choice to use a Class B network, the design needs a total of 34 bits (at least) in the mask (16 network, 8 subnet, 10 host), but only 32 bits exist—so no single mask meets the requirements.

2. B. With a growth of 20 percent, the design needs to support 240 subnets. To meet that need, 7 subnet bits do not meet the need ($2^7 = 128$), but 8 subnet bits do meet the need ($2^8 = 256$). Similarly, the minimum number of host bits is also 8, because the need, after the 20 percent growth, would be 144 hosts/subnet. That number requires 8 host bits ($2^8 - 2 = 254$). These numbers are minimum numbers of subnet and host bits.

   The right answer, 10.0.0.0/22, has 8 network bits because the network class is Class A, 14 subnet bits (/22 – 8 = 14), and 10 host bits (32 – 22 = 10). This mask supplies at least 8 subnet bits and at least 8 host bits. The masks in the other answers either do not supply at least 8 host bits or do not supply at least 8 subnet bits.

3. B. To support 1000 subnets, 10 subnet bits ($2^{10} = 1024$) are needed. The design uses a Class B network, which means that 16 network bits exist as well. So, the shortest mask that meets the requirements is 255.255.255.192, or /26, comprised of 16 network plus 10 subnet bits. The /28 answer also supplies enough subnets to meet the need, but compared to /26, /28 supplies fewer host bits and so fewer hosts/subnet.

4. D. To choose the mask that maximizes the number of subnets, choose the mask that uses the minimum number of host bits, which in turn maximizes the number of subnet bits. In this case, the design requires 200 hosts/subnet. Seven host bits do not supply enough hosts per subnet ($2^7 - 2 = 126$), but 8 host bits do supply enough hosts per

subnet ($2^8 - 2 = 254$). The design uses a Class A network, with 8 network bits. A prefix mask of /24 (a mask with 8 host bits, the minimum to support the right number of hosts/subnet) maximizes the size of the subnet part, in this case supplying 16 subnet bits.

5. C and D. The mask converts to 255.255.252.0, so the difference from subnet ID to subnet ID (called the magic number in this chapter) is 256 – 252 = 4. So, the subnet IDs start with 172.30.0.0, then 172.30.4.0, then 172.30.8.0, and so on, adding 4 to the third octet. The mask, used with a Class B network, implies 6 subnet bits, for 64 total subnet IDs. The last of these, 172.30.252.0, can be recognized in part because the third octet, where the subnet bits sit, has the same value as the mask in that third octet.

6. A. The first (numerically lowest) subnet ID is the same number as the classful network number, or 192.168.9.0. The remaining subnet IDs are each 8 larger than the previous subnet ID, in sequence, or 192.168.9.8, 192.168.9.16, 192.168.9.24, 192.168.9.32, and so on, through 192.168.9.248.

7. D. Using mask /24 (255.255.255.0), the subnet IDs increment by 1 in the third octet. The reasoning is that with a Class B network, 16 network bits exist, and with mask /24, the next 8 bits are subnet bits, so the entire third octet contains subnet bits. All the subnet IDs will have a 0 as the last octet, because the entire fourth octet consists of host bits. Note that 172.19.0.0 (the zero subnet) and 172.19.255.0 (the broadcast subnet) might look odd, but are valid subnet IDs.

8. D. Following the processes in this chapter, this Class A network and /25 mask defines 8 network bits (per the class) and 17 subnet bits (all of octets 2 and 3, plus 1 subnet bit in octet 4). As such, the magic number can be used in the fourth octet to find the value to add to the previous subnet ID to find the next subnet ID. The magic number in this case is 256 – 128 = 128.

   As a result, the fourth octet will be either a 0 or a 128. Octets 2 and 3 could be any value from 0 through 255 inclusive. Only one answer (10.1.1.192) breaks these range values.

Chapter 20

1. B, C, and D. Classless routing protocols, by definition, support VLSM because classless routing protocols transmit the subnet mask in their routing updates. Of the listed answers, only RIP version 1 (RIP-1) is not a classless routing protocol.

2. A. Note that sometimes VLSM stands for variable-length subnet masking, which refers to the process of using different masks in the same classful network, whereas variable-length subnet mask refers to the subnet masks themselves.

3. A. Subnet 10.5.0.0 255.255.240.0 implies range 10.5.0.0–10.5.15.255, which does not overlap. 10.4.0.0 255.254.0.0 implies range 10.4.0.0–10.5.255.255, which does overlap. 10.5.32.0 255.255.224.0 implies range 10.5.32.0–10.5.63.255, which does overlap. 10.5.0.0 255.255.128.0 implies range 10.5.0.0–10.5.127.255, which does overlap.

4. D. The four answers imply the following ranges: 172.16.0.0/21: 172.16.0.0–172.16.7.255. 172.16.6.0/23: 172.16.6.0–172.16.7.255. 172.16.16.0/20: 172.16.16.0–172.16.31.255. 172.16.11.0/25: 172.16.11.0–172.16.11.255. The subnet in the question, 172.16.8.0/22, implies a range of 172.16.8.0–172.16.11.127, which includes the range of numbers in subnet 172.16.11.0/25.

5. C. The question lists three existing subnets, which together consume part of Class C network 192.168.1.0. Just listing the last octet values, these subnets consume 0–63, 128–131, and 160–167. The new subnet, with a /28 mask, needs 16 consecutive numbers, and the subnet numbers will all be a multiple of 16 in the last octet (0, 16, 32, and so on). Looking at the consumed numbers again, the first opening starts at 64 and runs up through 127, so it has more than enough space for 16 addresses. So the numerically lowest subnet number is 192.168.1.64/28, with range 192.168.1.64–192.168.1.79.

Chapter 21

1. C. All the listed answers include the range of all three subnets, except for 10.3.64.0 255.255.224.0, which implies an address range of 10.3.64.0–10.3.95.255. Of the other three answers, 10.3.64.0 255.255.192.0 is the smallest range (10.3.64.0–10.3.127.255). It also happens to be the smallest single summary route that includes all three subnets listed in the question.

2. C and D. 10.0.0.0 255.0.0.0 implies a range of all addresses that begin with 10, and 10.1.0.0 255.255.0.0 implies a range of all addresses that begin with 10.1, so both of these answers include all the address ranges listed in the question. 10.1.32.0 255.255.224.0 implies a range of 10.1.32.0–10.1.63.255, which includes all addresses listed in the question. 10.1.55.0 255.255.255.0 implies a range of only 10.1.55.0–10.1.55.255, which does not include all addresses. 10.1.48.0 255.255.248.0 implies a range of 10.1.48.0–10.1.55.255, which omits two of the subnets listed in the question.

3. D. Following the process in the book, the range of addresses in the subnets are 10.1.12.0–10.1.12.255, 10.1.14.0–10.1.14.255, and 10.1.15.0–10.1.15.255. So, the entire range includes 10.1.12.0–10.1.15.255. Then, the lowest subnet ID from the list is 10.1.12.0, and the initial value of /P to try should be /23.

   Calculating a subnet ID and broadcast from these two numbers gives us 10.1.12.0/23, broadcast 10.1.13.255. This range does not include all addresses in the original subnets, so repeat the process with 10.1.12.0 and /22. This combination gives us subnet 10.1.12.0, broadcast 10.1.15.255, which does include all addresses, identifying the best summary as using subnet/mask 10.1.12.0/22.

4. B. Following the process in the book, the range of addresses in the subnets are 192.168.1.64–192.168.1.79, 192.168.1.80–192.168.1.95, and 192.168.1.96–192.168.1.111. So, the entire range includes 192.168.1.64–192.168.1.111. Then, the lowest subnet ID from the list is 192.168.1.64, and the initial value of /P to try should be /27.

Calculating a subnet ID and broadcast from these two numbers gives us 192.168.1.64/27, broadcast 192.168.1.95. This range does not include all addresses in the original subnets, so repeat the process with 192.168.1.64 and /26. This combination gives us subnet 192.168.1.64, broadcast 192.168.1.127, which does include all addresses, identifying the best summary as using subnet/mask 192.168.1.64/26.

Chapter 22

1. A and C. Standard ACLs check the source IP address. The address range 10.1.1.1–10.1.1.4 can be matched by an ACL, but it requires multiple **access-list** commands. Matching all hosts in Barney's subnet can be accomplished with the **access-list 1 permit 10.1.1.0 0.0.0.255** command.

2. A and D. The range of valid ACL numbers for standard numbered IP ACLs is 1–99 and 1300–1999, inclusive.

3. D. 0.0.0.255 matches all packets that have the same first three octets. This is useful when you want to match a subnet in which the subnet part comprises the first three octets, as in this case.

4. E. 0.0.15.255 matches all packets with the same first 20 bits. This is useful when you want to match a subnet in which the subnet part comprises the first 20 bits, as in this case.

5. A. The router always searches the ACL statements in order, and stops trying to match ACL statements after a statement is matched. In other words, it uses first-match logic. A packet with source IP address 1.1.1.1 would match any of the three explicitly configured commands described in the question. As a result, the first statement will be used.

6. B. One wrong answer, with wildcard mask 0.0.255.0, matches all packets that begin with 172.16, with a 5 in the last octet. One wrong answer matches only specific IP address 172.16.5.0. One wrong answer uses a wildcard mask of 0.0.0.128, which has only one wildcard bit (in binary), and happens to only match addresses 172.16.5.0 and 172.16.5.128. The correct answer matches the range of addresses 172.16.4.0–172.16.5.255, which includes all addresses in the subnet listed in the question (172.16.5.0–172.16.5.127).

Chapter 23

1. E and F. Extended ACLs can look at the Layer 3 (IP) and Layer 4 (TCP, UDP) headers and a few others, but not any application layer information. Named extended ACLs can look for the same fields as numbered extended ACLs.

2. A and E. The correct range of ACL numbers for extended IP access lists is 100 to 199 and 2000 to 2699. The answers that list the **eq www** parameter after 10.1.1.1 match the source port number, and the packets are going toward the web server, not away from it.

3. E. Because the packet is going toward any web client, you need to check for the web server's port number as a source port. The client IP address range is not specified in the question, but the servers are, so the source address beginning with 172.16.5 is the correct answer.

4. E. Named extended IP ACLs can match the exact same set of fields as can numbered extended IP ACLs.

5. A and C. Before IOS 12.3, numbered ACLs must be removed and then reconfigured to remove a line from the ACL. As of IOS 12.3, you can also use ACL configuration mode and sequence numbers to delete one ACL line at a time.

6. C. The authorized Cisco curriculum makes the suggestion in answer C for extended IP ACLs, suggesting that standard ACLs be placed as close to the destination as possible.

7. A. NTP uses protocol messages between clients and servers so that the clients can adjust their time-of-day clock to match the server. NTP is totally unrelated to serial line clocking. It also does not count CPU cycles, instead relying on messages from the NTP server. Also, the client defines the IP address of the server, and does not have to be in the same subnet.

Chapter 24

1. F. The word *classless* in this term refers to the fact that the routes advertise a block of addresses with a mask, and can essentially ignore the Class A, B, and C boundaries with regard to what address block is used. Also, because the motivation was to describe routes between different companies (different routing domains), the term uses the word *interdomain*.

2. D. CIDR's original intent was to allow the summarization of multiple Class A, B, and C networks to reduce the size of Internet routing tables. Of the answers, only 200.1.0.0 255.255.0.0 summarizes multiple networks.

3. B and E. RFC 1918 identifies private network numbers. It includes Class A network 10.0.0.0, Class B networks 172.16.0.0 through 172.31.0.0, and Class C networks 192.168.0.0 through 192.168.255.0.

4. C. With static NAT, the entries are statically configured. Because the question mentions translation for inside addresses, the **inside** keyword is needed in the command.

5. A. With dynamic NAT, the entries are created as a result of the first packet flow from the inside network.

6. A. The **list 1** parameter references an IP ACL, which matches packets, identifying the inside local addresses.

7. A and C. The configuration is missing the **overload** keyword in the **ip nat inside source** command and in the **ip nat outside** interface subcommand on the serial interface.

8. B. The last line mentions that the pool has seven addresses, with all seven allocated, with the misses counter close to 1000—meaning that close to 1000 new flows were rejected because of insufficient space in the NAT pool.

Chapter 25

1. C. NAT, specifically the PAT feature that allows many hosts to use private IPv4 addresses while being supported by a single public IPv4 address, was one short-term solution to the IPv4 address exhaustion problem. IP version 5 existed briefly as an experimental protocol and had nothing to do with IPv4 address exhaustion. IPv6 directly addresses the IPv4 address exhaustion problem, but it is a long-term solution. ARP has no impact on the number of IPv4 addresses used.

2. A. Routers use the same process steps when routing IPv6 packets as they do when routing IPv4 packets. Routers route IPv6 packets based on the IPv6 addresses, listed inside the IPv6 header in the IPv6 packets, by comparing the destination IPv6 address to the router's IPv6 routing table. As a result, the router discards the incoming frame's data link header and trailer, leaving an IPv6 packet. The router compares the destination (not source) IPv6 address in the header to the router's IPv6 (not IPv4) routing table and then forwards the packet based on the matched route.

3. D. If following the steps in the book, the first step removes up to three leading 0s in each quartet, leaving FE80:0:0:100:0:0:0:123. This leaves two strings of consecutive all-0 quartets; by changing the longest string of all 0s to ::, the address is FE80:0:0:100::123.

4. B. This question has many quartets that make it easy to make a common mistake: removing trailing 0s in a quartet of hex digits. To abbreviate IPv6 addresses, only leading 0s in a quartet should be removed. Many of the quartets have trailing 0s (0s on the right side of the quartet), so make sure to not remove those 0s.

5. A. The unabbreviated version of an IPv6 address must have 32 digits, and only one answer has 32 hex digits. In this case, the original number shows four quartets and a ::. So, the :: was replaced with four quartets of 0000, making the number have eight quartets. Then, for each quartet with fewer than four digits, leading 0s were added so that each quartet has four hex digits.

6. C. The /64 prefix length means that the last 64 bits, or last 16 digits, of the address should be changed to all 0s. That process leaves the unabbreviated prefix as 2000:0000:0000:0005:0000:0000:0000:0000. The last four quartets are all 0s, making that string of all 0s be the longest and best string of 0s to replace with ::. After removing the leading 0s in other quartets, the answer is 2000:0:0:5::/64.

Chapter 26

1. C. Unique local addresses begin with FD in the first two digits.

2. A. Global unicast addresses can begin with many different initial values, but most commonly, they begin with either a hex 2 or 3.

3. D. The global routing prefix is the address block, represented as a prefix value and prefix length, given to an organization by some numbering authority. All IPv6 addresses inside the company have the same value in these initial bits of their IPv6 addresses. Similarly, when a company uses a public IPv4 address block, all the addresses have the same value in the network part.

4. B. Subnetting a global unicast address block, using a single prefix length for all subnets, breaks the addresses into three parts. The parts are the global routing prefix, subnet, and interface ID.

5. D. Unique local addresses begin with a 2-hex-digit prefix of FD, followed by the 10-hex-digit global ID.

Chapter 27

1. A. The one correct answer lists the exact same IPv6 address listed in the question, with a /64 prefix length, and no spaces in the syntax of the answer. Another (incorrect) answer is identical, except it leaves a space between the address and prefix length, which is incorrect syntax. The two answers that list the **eui-64** parameter both list an address and not a prefix; they should list a prefix to be correct, although neither would have resulted in the IPv6 address listed in the question.

2. B. With the **eui-64** parameter, the router will calculate the interface ID portion of the IPv6 address based on its MAC address. Beginning with 5055.4444.3333, the router injects FF FE in the middle (5055.44FF.FE44.3333). Then, the router inverts the seventh bit in the first byte. Mentally, this converts hex 50 to binary 01010000, changing bit 7 so that the string is 0101 0010, and converting back to hex 52. The final interface ID value is 5255:44FF:FE44:3333. The wrong answers simply list a different value.

3. A and C. Of the four answers, the two correct answers show the minimal required configuration to support IPv6 on a Cisco router: enabling IPv6 routing (**ipv6 unicast-routing**) and enabling IPv6 on each interface, typically by adding a unicast address to each interface (**ipv6 address...**). The two incorrect answers list nonexistent commands.

4. A. With an **ipv6 address** command configured for a global unicast address, but without a link-local address configured with an **ipv6 address** command, the router calculates its link-local address on the interface based on its MAC address and EUI-64 rules. The first half of the link-local address begins FE80:0000:0000:0000. The router then calculates the second half of the link-local address value by taking the MAC address (0200.0001.000A), injecting FF FE in the middle (0200.00FF.FE01.000A), and flipping the seventh bit (0000.00FF.FE01.000A).

5. B. FF02::1 is used by all IPv6 hosts on the link, FF02::5 is used by all OSPFv3 routers, and FF02::A is used by all EIGRPv6 routers. FF02::2 is used to send packets to all IPv6 routers on a link.

6. C. R1 will form the solicited node multicast address by using a format of FF02::1:FFxx:xxxx, replacing xx:xxxx with the last six hex digits of the unicast address. If the link-local address and unicast address have different values in their last six hex digits, the router creates a solicited node multicast for both. To form the solicited node multicast address FF02::1:FF00:A, the last six hex digits of an address must be hex 00:000A. Of three answers that define an IPv6 unicast addresses, only one lists a value whose last six hex digits are 00:000A. Also, in this case, the link-local address will be FE80::FF:FE01:A, resulting in a solicited node multicast address of FF02::1:FF01:A, which is not the address listed in the question.

Chapter 28

1. B. PC1 needs to discover PC2's MAC address. Unlike IPv4, IPv6 does not use ARP, instead using NDP. Specifically, PC1 uses the NDP Neighbor Solicitation (NS) message to request that PC2 send back an NDP Neighbor Advertisement (NA). SLAAC relates to address assignment, and not to discovering a neighbor's MAC address.

2. D. Hosts can ask for (solicit) all routers to identify themselves by sending an NDP Router Solicitation (RS) message, with the routers sending back an NDP Router Advertisement (RA) message. NDP Neighbor Solicitation (NS) can also be used by PC1, but not for the purpose of learning its default router IPv6 address. DAD is a function that happens to use NDP NS and NA messages, but its function does not include the discovery of the default router address. Finally, EUI-64 does not define a protocol or message.

3. A and C. The NDP RA lists the router IPv6 address, the IPv6 prefixes known on the link, and the matching prefix lengths. When using DHCPv6, the host learns the IPv6 address of the DNS server through DHCPv6 messages. For MAC addresses of on-link neighbors, hosts use NDP NS and NA messages.

4. C. Hosts use NDP messages, namely Router Solicitation (RS) and Router Advertisement (RA), to learn the IPv6 address of potential default routers. The other settings typically come from the stateful DHCPv6 server.

5. D. SLAAC gives the host a means to choose its unicast address. The host also uses NDP to learn its prefix length, plus the address(es) of any default routers. It then uses stateless DHCP to learn the addresses of the DNS server(s).

6. B and D. With SLAAC, the host learns the prefix from a router using NDP RS/RA messages, and then the host builds the rest of the address (the interface ID). The host can use EUI-64 rules or use a defined process to randomly generate the interface ID value. The host does not learn the interface ID from any other device, which helps make the process stateless, because no other device needs to assign the host its full address.

7. A. The **show ipv6 neighbors** command lists all IPv6 addresses of neighbors (both routers and hosts), plus their matching MAC addresses. It does not note which are routers, leaving that information for the **show ipv6 routers** command.

Chapter 29

1. A and C. With an IPv6 address on a working interface, the router adds a connected route for the prefix (subnet) implied by the **ipv6 address** command. It also adds a local host route (with a /128 prefix length) based on the unicast address. The router does not add a route based on the link-local address.

2. A. All four answers show examples of commands that use an outgoing interface. The two commands that begin with **ip route** define only IPv4 routes; the commands would be rejected because of the IPv6 prefixes listed in the commands. The two commands that begin with **ipv6 route** are syntactically correct, but the command should list the local router's interface (an interface on the router on which the command is being configured). R5 needs to use its local S0/1/1 interface as the outgoing interface.

3. B. All four answers show examples of commands that use a next-hop router IPv6 address. Two of the answers list R5's own IPv6 address (unicast or link-local), which is incorrect; the address should be an address on the neighboring router, R6 in this case. For the two answers that list addresses on Router R6, the one that lists R6's global unicast address is correct. The one that lists R6's link-local address would also require R5's outgoing interface, so the answer that lists FE80::FF:FE00:6 would be rejected as well.

4. A and D. Because the router already has IPv4 configured, the new IPv6 OSPFv3 process will be able to choose its own router ID automatically. So, the **router-id** command for OSPFv3 (one of the answers), while valid, is not required. OSPFv3 does not use a **network** *prefix/length* command, so that answer is not correct. The other two answers list commands required for the configuration.

5. B and D. The **show ipv6 ospf interface brief** command lists one line of output for each interface on which OSPV3 is enabled. The **show ipv6 protocols** command lists information about all sources of IPv6 routing information; the messages about OSPFv3 list the interfaces on which it is enabled.

6. B and C. OSPFv2 support IPv4 routes, while OSPFv3 supports IPv6. EIGRP for IPv6 is typically called EIGRPv6.