CCNA R&S
(200-125)
Ultimate Practice Exam

by

Kevin Wallace
CCIEx2 #7945 (R&S and Collab.)

KEVIN WALLACE TRAINING, LLC

CCNA R&S (200-125)
Ultimate Practice Exam

by Kevin Wallace, CCIEx2 (R&S and Collaboration) #7945

Introduction

Thanks so much for your purchase of the **Ultimate CCNA R&S (200-125) Practice Exam**. Before diving into the questions, let's talk about the philosophy behind these questions. First of all, I believe that brain dumps of Cisco® exams, which you might find on the Internet, hurt our industry. They help people without an understanding of the technology become certified, which devalues the certifications that we've worked so hard to earn.

Here's how I'm combatting that. I've dramatically undercut the prices of the big brain dump vendors. Also, I've created a set of questions that addresses the topics found on the exam (according to Cisco's exam blueprint), and I've thoroughly explained why the correct answer is correct. So, this practice exam is not an attempt to have you memorize answers, but rather, the goal is to help you understand the technology at a deeper level and confirm your understanding of the various exam topic areas.

I salute you for your integrity in choosing a brain-dump-free practice exam.

Kevin Wallace, CCIEx2 (R&S and Collaboration) #7945

Video Resources:

To supplement your studies, I've put together a series of Cisco routing and switching training videos for you.

Get Your First Free Video Here:

http://kwtrain.com/mcast

Also, I personally created a video training series for Cisco Press that covers all of the exam blueprint topics, includes 300+ videos, 25+ hours of instruction, interactive exercises, Glossary Quizzes, Module Quizzes, and more full practice exams.

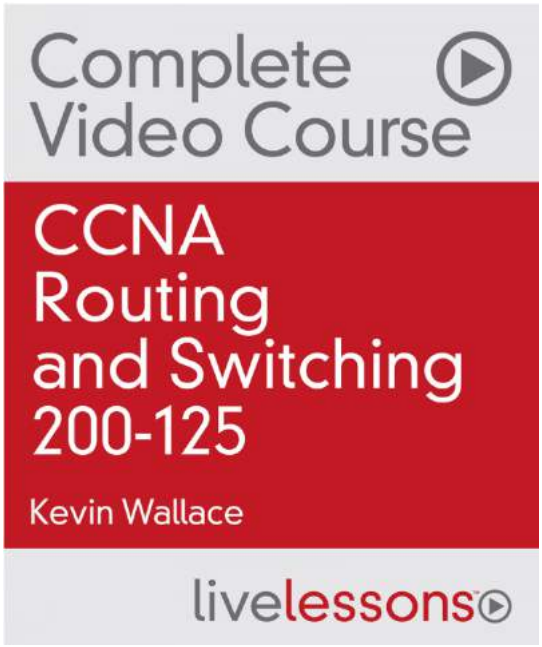

Learn More About the Complete Course Here:

http://kwtrain.com/ccnacourse

CCNA R&S (200-125) Ultimate Practice Exam Questions

Question #1

Which of the following IPv6 address ranges is used for *Global Unicast* IPv6 addresses?

a. 2000::/3
b. FF00::/8
c. FEC0::/10
d. FC00::/7
e. ::1
f. ::
g. FF02::1::FF00:0/104

Question #2

You are working with an enterprise router connecting out to two Internet Service Providers (ISPs). The router has a single link to each ISP. What type of topology is described by this scenario?
a. Single Homed
b. Dual Homed
c. Single Multihomed
d. Dual Multihomed

Question #3

What is the range of assignable IP addresses for a subnet containing an IP address of 172.16.1.10 /19?

a. 172.16.0.1 – 172.16.31.254
b. 172.16.0.1 – 172.16.63.254
c. 172.16.0.0 – 172.16.31.255
d. 172.16.0.1 – 172.16.31.255
e. 172.16.0.0 – 172.16.63.254

Question #4

Consider the following topology and configuration for router R1.

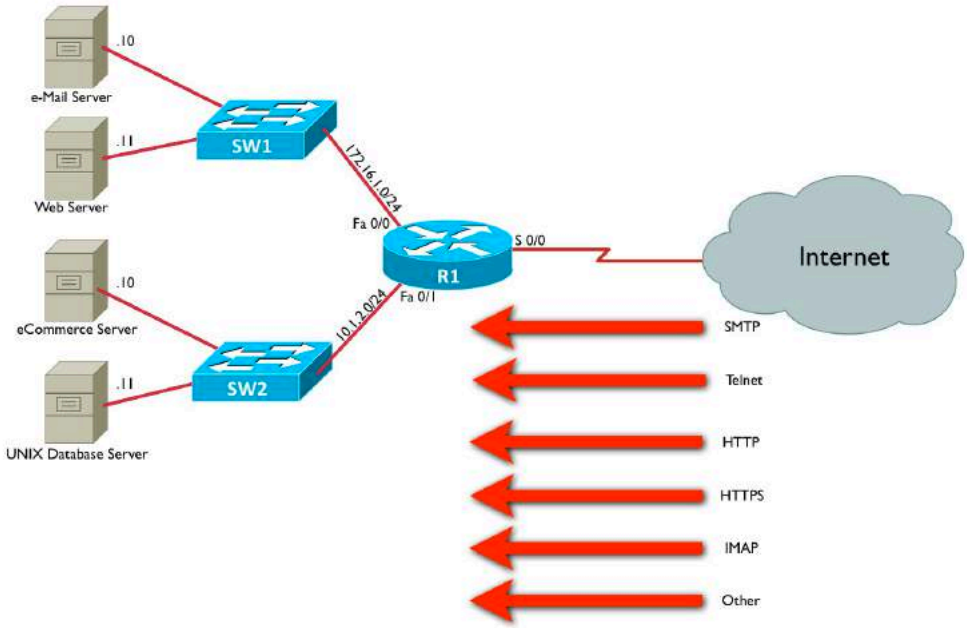


```
ip access-list extended TEST
 permit tcp any host 172.16.1.10 eq smtp
 deny ip any 172.16.1.0 0.0.0.255
 permit ip any 10.1.2.0 0.0.0.255
!
interface Serial 0/0
 ip access-group TEST in
```

What traffic from the Internet will be allowed to pass through router R1? (Choose all that apply.)

a. SMTP traffic destined for the e-mail server
b. HTTP traffic destined for the web server
c. HTTPS traffic destined for the eCommerce server
d. Telnet traffic destined for the UNIX database server

Question #5

Which of the following is a set of wiring standards that describe the color coding of wires in RJ-45 connections?
a. CSMA/CD
b. Multimode
c. EIA/TIA 568
d. RS-232

Question #6

You are configuring Network Address Translation (NAT). A PC on the inside of your network has an IP address of 192.168.1.10. However, since that IP address

is a private IP address, it needs to be translated into an address that is routable on the public Internet. What NAT terminology is used to describe the PC's 192.168.1.10 IP address?

a. inside global address
b. outside local address
c. inside local address
d. outside global address

Question #7

What is the default Holdtime for HSRP?
a. 1 second
b. 3 seconds
c. 5 seconds
d. 10 seconds

Question #8

In the following topology, how many collision domains are represented?

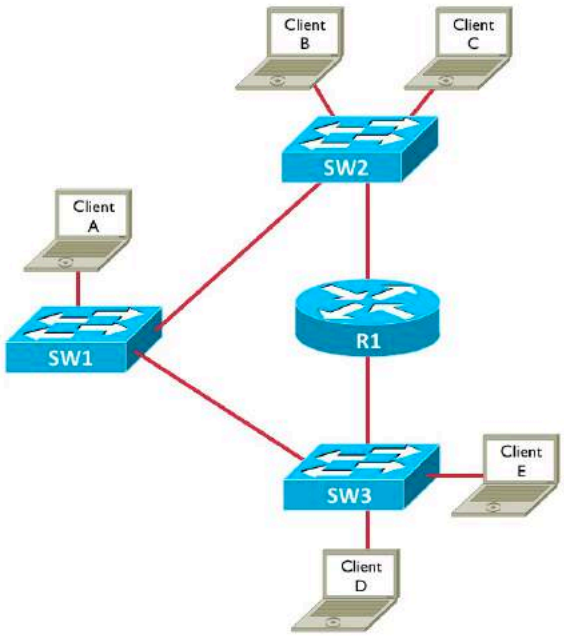

a. 3
b. 4
c. 5
d. 9

Question #9

Consider the following port security configuration:

```
Switch(config)# int gig 1/0/1
Switch(config-if)# switchport mode access
Switch(config-if)# switchport port-security
Switch(config-if)# switchport port-security maximum 3
Switch(config-if)# switchport port-security mac-address sticky
Switch(config-if)# switchport port-security violation shutdown
```

What does the **shutdown** option in the bottom command do?
a. The **shutdown** option causes the port to go into an err-disable state if a port security violation occurs.
b. The **shutdown** option causes the port to be administratively shutdown if a port security violation occurs.
c. The **shutdown** option disables port security on this port.
d. The **shutdown** option drops traffic from the fourth MAC address seen on this port, while permitting traffic from the first three MAC addresses learned on this port. However, the security violation counter does not get incremented.
e. The **shutdown** option drops traffic from the fourth MAC address seen on this port, while permitting traffic from the first three MAC addresses learned on this port. Also, the security violation counter is incremented.

Question #10

You connect a laptop to an available port on a Cisco Catalyst switch. By default, how long does it take that port to start forwarding traffic?
a. 0 seconds
b. 15 seconds
c. 30 seconds
d. 50 seconds

Question #11

Consider the output of the **show controllers** commands issued on routers R1 and R2. Of these two routers, which one should have a clock rate set for the specified serial interface, and how do you know from the output?

```
R1#show controllers serial 0/2/0
Interface Serial0/2/0
Hardware is GT96K
DTE V.35idb at 0x47FFCA4C, driver data structure at 0x48004290
wic_info 0x480048BC
Physical Port 5, SCC Num 5
MPSC Registers:
MMCR_L=0x000304C0, MMCR_H=0x00000000, MPCR=0x00000100
CHR1=0x00FE007E, CHR2=0x00000000, CHR3=0x0000064A, CHR4=0x00000000
CHR5=0x00000000, CHR6=0x00000000, CHR7=0x00000000, CHR8=0x00000000
CHR9=0x00000000, CHR10=0x00003008
SDMA Registers:

... OUTPUT OMITTED ...
```

```
R2#show controllers serial 1/0
CD2430 Slot 1, Port 0, Controller 0, Channel 0, Revision 19
Channel mode is synchronous serial
idb 0x466767F0, buffer size 1524, V.35 DCE cable, clockrate 128000

Global registers
  rpilr 0x2, rir 0x2, risr 0x0, rfoc 0x0, rdr 0x0
  tpilr 0x1, tir 0x0, tisr 0x68, tftc 0x0, tdr 0x90
  mpilr 0x3, mir 0x1, misr 0x60
  bercnt 0xFF, stk 0x0
Per-channel registers for channel 0

... OUTPUT OMITTED ...
```

a. R1 should have the clock rate set, because the DTE end of the serial cable connects to R1.
b. R1 should have the clock rate set, because the DCE end of the serial cable connects to R1.
c. R2 should have the clock rate set, because the DTE end of the serial cable connects to R2.
d. R2 should have the clock rate set, because the DCE end of the serial cable connects to R2.

Question #12

A router resides at what layer of the TCP/IP Model?

a. Application
b. Transport
c. Internet
d. Network Interface
e. Session

Question #13

Which of the following is not a Rapid Spanning Tree Protocol (RSTP) port state?
a. Listening
b. Discarding
c. Learning
d. Forwarding

Question #14

In the diagram below, Client A is sending a packet to Host 1. As the packet is coming into the Fa 0/0 interface on router R2, what is the source IP address in the packet's header?

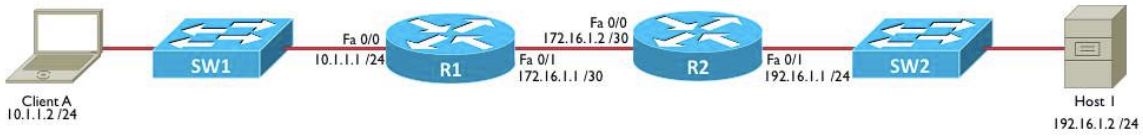

a. 10.1.1.1
b. 172.16.1.2
c. 192.16.1.1
d. 10.1.1.2
e. 172.16.1.1
f. 192.16.1.2

Question #15

In the diagram below, Client A is sending a packet to Host 1. Which devices will make a packet forwarding decision based on a destination IP address? (Choose 2.)

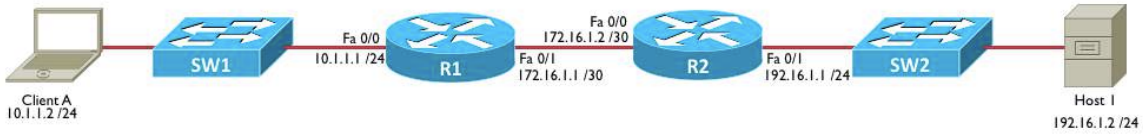

a. switch SW1
b. router R1
c. router R2
d. switch SW2

Question #16

Which of the following is considered to be a reliable Transport Layer protocol?

a. IP
b. UDP
c. TCP
d. ICMP
e. PPP

Question #17

You are configuring eBGP on router R1 seen in the figure. You want to form a BGP neighborship with router R2. Which command would you issue from within BGP router configuration mode on router R1?

a. `neighbor 198.51.100.2 remote-as 64495`
b. `neighbor 198.51.100.2 remote-as 64500`
c. `neighbor 198.51.100.1 remote-as 64495`
d. `neighbor 198.51.100.1 remote-as 64500`

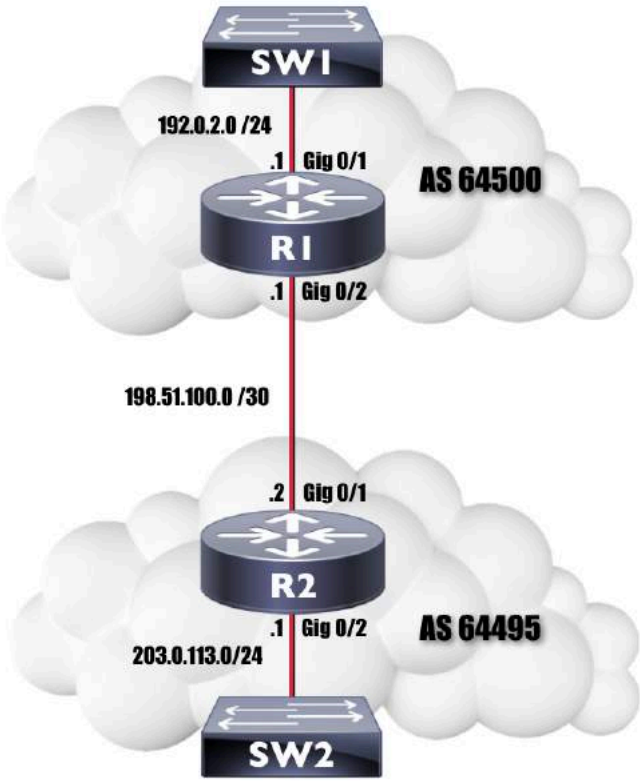

Question #18

You are assigning IP addresses to hosts in the 192.168.4.0 /26 subnet. Which two of the following IP addresses are assignable IP addresses that reside in that subnet?

a. 192.168.4.0
b. 192.168.4.63
c. 192.168.4.62
d. 192.168.4.32
e. 192.168.4.64

Question #19

Which of the following are security best practices when setting up a Cisco router? (Choose all the apply.)

a. Enable the password encryption service.
b. Enable Telnet for remote access.
c. Create a welcome.
d. Use the **enable password** command to restrict access to the router's privileged mode.
e. Issue the **exec-timeout 0 0** command to log out any idle sessions.

Question #20

A host in your network has been assigned an IP address of 192.168.181.182 /25. What is the subnet to which the host belongs?

a. 192.168.181.128 /25
b. 192.168.181.0 /25
c. 192.168.181.176 /25
d. 192.168.181.192 /25
e. 192.168.181.160 /25

Question #21

The PING utility uses which of the following protocols to check for network connectivity to an IP address?

a. UDP
b. TCP
c. IGMP
d. ICMP
e. OSPF

Question #22

On an IEEE 802.1Q trunk, what do we call the VLAN that is not tagged by 802.1Q?
a. Access VLAN
b. Voice VLAN
c. Trunk VLAN
d. Native VLAN
e. Passive VLAN

Question #23

You wish to Telnet into a Layer 2 Cisco Catalyst switch. Which of the following must be configured on the switch before you can Telnet into the switch?

a. You must issue the **no switchport** command on one of the interfaces.
b. You must assign an IP address to the switch.
c. You must enable IP routing on the switch.
d. You must issue the **service telnet** global configuration command.
e. You must configure a *router-on-a-stick* topology to provide a Layer 3 path to the switch.

Question #24

You are working with a Class B network with the private IP address of 172.16.0.0 /16. You need to maximize the number of broadcast domains, where each

broadcast domain can accommodate 1000 hosts. What subnet mask should you use?

a. /22
b. /23
c. /24
d. /25
e. /26

Question #25

CDP (Cisco Discovery Protocol) resides at which layer of the OSI Model?

a. Layer 1
b. Layer 2
c. Layer 3
d. Layer 4
e. Layer 5

Question #26

Which of the following technologies allows VPN tunnels to be setup and torn down on an as-needed basis?
a. IPsec
b. DMVPN
c. GRE
d. PPP

Question #27

Typically, HTTP uses what Transport Layer protocol and what port number?

a. UDP port 80
b. TCP port 443
c. UDP port 443
d. TCP port 80
e. TCP port 143

Question #28

BPDUGuard should be only be configured on ports that have which feature enabled?
a. RootFast
b. PortFast
c. RootGuard
d. UplinkFast

Question #29

A PC has just booted up and wants to communicate with a host on a remote subnet. The PC knows the IP addresses of its default gateway and the remote host. However, in order to properly construct a frame, the PC needs an appropriate destination MAC address. What MAC address does the PC need to learn, and what protocol will the PC use to learn that MAC address?

a. the MAC address of the PC's default gateway, learned via ARP
b. the MAC address of the remote host, learned via DNS
c. the MAC address of the PC's default gateway, learned via DNS
d. the MAC address of the remote host, learned via ARP

Question #30

What is the directed broadcast address of a subnet containing an IP address of 172.16.1.10 /19?

a. 172.16.15.255
b. 172.16.31.255
c. 172.16.255.255
d. 172.16.95.255
e. 172.16.0.255

Question #31

Of the following protocols, select the ones that are UDP-based. (Choose 2)

a. SSH
b. NTP
c. IMAP4
d. DHCP
e. POP3

Question #32

Consider the following topology. Client F has just booted up, and its MAC address has not been learned by switch SW1 (as seen in the output of **the show mac address-table** command). If Client A sends a frame destined for Client F's MAC address, what will switch SW1 do with the frame? (Assume all switch ports are assigned to VLAN 100.)

a. SW1 will flood the frame out of all of its ports.
b. SW1 will flood the frame out of all of its ports, other than the port on which the frame was received.
c. SW1 will drop the frame, because there is no entry for the destination MAC address in switch SW1's MAC address table.
d. SW1 will forward the frame out of port Gi 1/0/2 only.

e. SW1 will forward the frame out of port Gi 1/0/1 only.

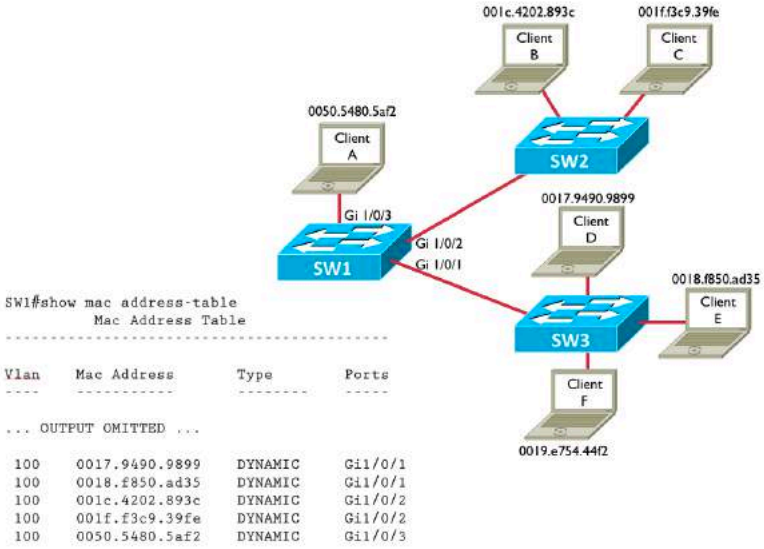

Question #33

A customer is using a Class C network of 192.168.10.0 subnetted with a 28-bit subnet mask. How many subnets can be created by using this subnet mask?

a. 32
b. 16
c. 30
d. 8
e. 14

Question #34

A PC requires which of the following configuration parameters in order to Telnet to a host on a remote subnet? (Choose 3.)

a. IP address
b. subnet mask
c. routing protocol
d. VLAN assignment
e. default gateway

Question #35

Which of the follow are dynamic routing protocols? (Choose 2.)
a. ICMP
b. OSPF
c. SNMP
d. ARP
e. EIGRP

Question #36
Identify the three layers of Cisco's traditional hierarchical network design model. (Choose 3.)

a. Access
b. Cross-connect
c. Distribution
d. Uplink
e. Core

Question #37
Which of the following devices can be used to interconnect broadcast domains? (Choose 2.)

a. router
b. Ethernet hub
c. Layer 2 switch
d. Ethernet bridge
e. Layer 3 switch

Question #38
ARP can perform which of the following functions?

a. ARP allows a network host to learn the MAC address corresponding to a known IP address.
b. ARP allows a network host to dynamically obtain an IP address.
c. ARP allows a network host to learn an IP address corresponding to a known domain name.
d. ARP allows a network host to learn the IP address corresponding to a known MAC address.

Question #39
You're viewing the running configuration on a Cisco router and notice the **enable secret** command. The command begins with **enable secret**, followed by the number **9**, followed by a hash digest. What does the **9** indicate?
a. The password is encrypted using AES.
b. The password is hashed using Message Digest 5 (MD5).
c. The password is hashed using Password-Based Key Derivation Function 2 (PBKDF2).
d. The password is hashed using SCRYPT.

Question #40

An Ethernet hub has 12 Ethernet ports. How many collision domains and how many broadcast domains exist on the Ethernet hub?

a. 12 Collision Domains and 12 Broadcast Domains
b. 1 Collision Domain and 12 Broadcast Domains
c. 12 Collision Domains and 1 Broadcast Domain
d. 1 Collision Domain and 1 Broadcast Domain

Question #41

A router performs which of the following tasks? (Choose the 3 best answers.)

a. packet switching
b. rewriting the MAC addresses of Ethernet frames
c. determining the full path to a destination network
d. determining the next hop to a destination network
e. rewriting the IP addresses of IP packets

Question #42

Given the following topology, what would be the most efficient way of configuring routing on the BR1 router, such that devices at the BR1 site would have connectivity to the HQ site and the Internet?

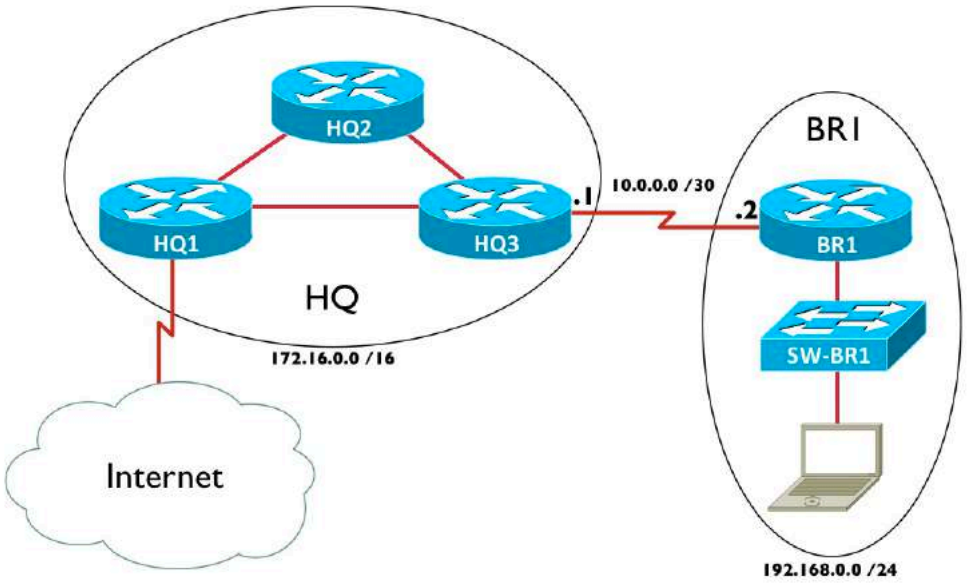

a. Create a static route to the 172.16.0.0 /16 network.
b. Create a default static route.
c. Configure OSPFv3.
d. Configure EIGRP.
e. Configure BGP.

Question #43
Which of the following are components of an Ethernet frame? (Choose 3.)
a. Preamble
b. ToS Byte
c. SOF
d. DF bit
e. FCS

Question #44
To what address does EIGRP for IPv6 send updates?
a. 224.0.0.10
b. FF02::A
c. ::1
d. FE80::10

Question #45
Your router has more than one routing information source that is telling the router how to reach the 10.10.10.0 /24 network. One source is a statically configured route. Another source is the OSPF routing protocol, and yet another source is the EIGRP routing protocol. What determines which routing information source the router will use?

a. Static routes are always trusted over dynamically learned routes.
b. The route with the lowest metric will be used.
c. The information source with the lowest administrative distance will be used.
d. The route will load-balance among the three available routes.

Question #46
You are sitting at the privileged mode command line prompt on RouterA, and you want to remotely connect to RouterB (which has an IP address of 172.16.1.1). What command would you issue on RouterA?

a. **ssh 172.16.1.1**
b. **traceroute 172.16.1.1**
c. **ping 172.16.1.1**
d. **telnet 172.16.1.1**

Question #47
You are configuring an EtherChannel connection between two Cisco Catalyst switches, using PAgP for dynamic EtherChannel formation. Ports on one side of the EtherChannel are configured to the Auto PAgP channel mode. To which PAgP channel mode must the ports at the other side of the EtherChannel be configured?

a. On
b. Auto
c. Desirable
d. Active
e. Passive

Question #48

Consider the following topology. You issue the **show cdp neighbors** command on router R1. Which of the following devices will appear in the output of the command? (Choose all that apply.)

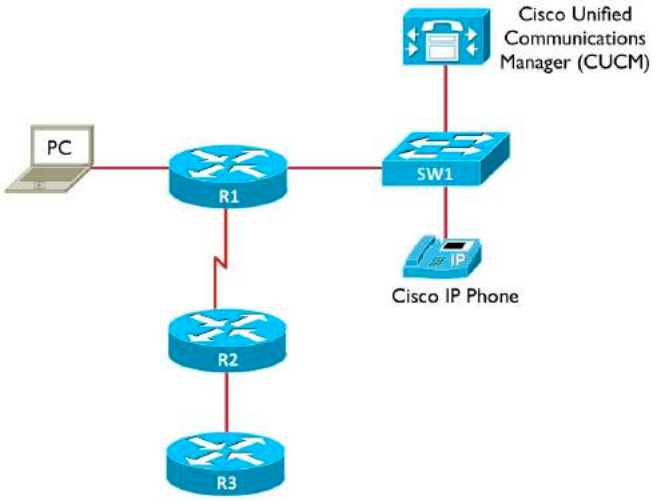

a. PC
b. CUCM server
c. Cisco IP Phone
d. R2
e. R3
f. SW1

Question #49

The following output was generated from what command?

```
Global values:
  Internet Protocol routing is enabled
FastEthernet0/0 is up, line protocol is up
  Internet address is 10.1.1.1/24
FastEthernet0/1 is up, line protocol is up
  Internet address is 4.4.4.4/24
FastEthernet1/0 is administratively down, line protocol is down
FastEthernet1/1 is administratively down, line protocol is down
```

a. **show cdp neighbor**
b. **show protocols**
c. **show version**
d. **show ip interface brief**

e. **show interfaces**
f. **show running-config**

Question #50

The default OSPF cost for a 1 Gbps link is 1. What is the default cost for a 100 Mbps link?
a. 1
b. 10
c. 100
d. 1000

Question #51

Which of the follow protocols is used to determine the MAC address for a known IP address?
a. DNS
b. DHCP
c. WINS
d. ARP
e. ICMP

Question #52

A traditional Ethernet switch operates at which layer of the OSI Model?

a. Transport
b. Data Link
c. Network
d. Physical
e. Session

Question #53

You are in interface configuration mode of a Cisco router, and you want to assign an IP address of 172.16.1.1 /24 to the interface. Which of the following is the command you should enter?
a. Router1(config-if)# **ip address 172.16.1.1 /24**
b. Router1(config-if)# **ip address 172.16.1.1 0.0.0.255**
c. Router1(config-if)# **ip address 172.16.1.1 255.255.255.0**
d. Router1(config-if)# **ip address 172.16.1.1 classful**

Question #54

What file transfer protocol uses a connectionless Layer 4 transport protocol and does not required user authentication?

a. TFTP
b. SFTP

c. FTP
d. SSH
e. Telnet

Question #55

A Layer 2 Ethernet switch with 12 ports, where all ports belong to the same VLAN, has how many collision domains and how many broadcast domains?

a. 12 collision domains and 12 broadcast domains
b. 1 collision domain and 1 broadcast domain
c. 12 collision domains and 1 broadcast domain
d. 1 collision domain and 12 broadcast domains

Question #56

RIPv2 advertisements are sent to what multicast IP address?
a. 224.0.0.5
b. 224.0.0.6
c. 224.0.0.9
d. 224.0.0.10

Question #57

In the diagram below, Client A is sending a packet to Host 1. As the frame is coming into the Fa 0/0 interface on router R2, what is the destination MAC address in the frame's header?

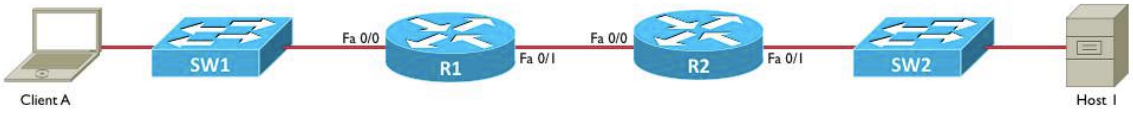

a. Host 1's MAC address
b. Client A's MAC address
c. Router R1's Fa 0/0 MAC address
d. Router R2's Fa 0/0 MAC address
e. Router R2's Fa 0/1 MAC address

Question #58

You are connected to the console line of RouterA. From there, you connect to RouterB via Telnet. Without terminating the Telnet session, what key sequence could you enter to return to the RouterA prompt?
a. **<CTRL-SHIFT-6> x**
b. **<CTRL-Break>**
c. **<CTRL-ALT-DELETE>**
d. **<CTRL-ALT-x>**

Question #59

Which three of the following are components of a network secured using IEEE 802.1x? (Choose 3.)

a. Encryption Server
b. Supplicant
c. Authorization Server
d. Key Manager
e. Authenticator
f. Authentication Server

Question #60

What protocol allows multiple hosts to dynamically obtain IP addresses from a server?

a. DNS
b. DHCP
c. WINS
d. ARP
e. ICMP

Question #61

Which of the following are true regarding CSMA/CD? (Choose 3.)

a. CSMA/CD is an Ethernet technology.
b. CSMA/CD is a wireless technology.
c. CSMA/CD should run on every port of an Ethernet switch.
d. CSMA/CD is used for half-duplex connections.
e. CSMA/CD is used by devices connecting to Ethernet hubs.

Question #62

The following output was generated from what command?

```
FastEthernet0/1 is up, line protocol is up
  Hardware is i82543 (Livengood), address is ca03.1af0.0006 (bia
ca03.1af0.0006)
  Internet address is 4.4.4.4/24
  MTU 1500 bytes, BW 100000 Kbit/sec, DLY 100 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 100Mb/s, 100BaseTX/FX
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:00, output 00:00:03, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
```

```
476 packets input, 41501 bytes
Received 333 broadcasts (0 IP multicasts)
0 runts, 0 giants, 0 throttles
0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
0 watchdog
0 input packets with dribble condition detected
481 packets output, 41819 bytes, 0 underruns
0 output errors, 0 collisions, 1 interface resets
0 unknown protocol drops
0 babbles, 0 late collision, 0 deferred
0 lost carrier, 0 no carrier
0 output buffer failures, 0 output buffers swapped out
```

a. **show cdp neighbor**
b. **show protocols**
c. **show version**
d. **show ip interface brief**
e. **show interfaces**
f. **show running-config**

Question #63

You connect a laptop to an available port on a Cisco Catalyst switch. By default, how long does it take that port to start forwarding traffic?
a. 0 seconds
b. 15 seconds
c. 30 seconds
d. 50 seconds

Question #64

What is another term for Port Address Translation?

a. static NAT
b. dynamic NAT
c. NAT overloading
d. NAT pooling

Question #65

By default, a DHCP Discover message cannot pass through a router, because it is a broadcast packet. What interface configuration-mode command can cause the router to forward the DHCP Discover message to a target IP address or subnet?

a. Router(config-if)# **dhcp-relay *ip-address***
b. Router(config-if)# **ip helper-address *ip-address***
c. Router(config-if)# **ip discover-forward *ip-address***
d. Router(config-if)# **forward-bootp *ip-address***

Question #66

What is the Administrative Distance (AD) for External EIGRP?
a. 90
b. 110
c. 120
d. 170

Question #67

You enter the following commands in a router:

```
Router(config)# enable secret Pa$$1
Router(config)# enable password Pa$$2
```

What password must you enter the next time you attempt to enter privileged mode on the router?

a. You must enter both passwords.
b. Pa$$1
c. Pa$$2
d. Since the passwords do not match, remote authentication is disabled.

Question #68

Your new employee tells you that they are unable to log into router R1 via Telnet. You examine the router's configuration and find the following configuration:

```
... OUTPUT OMITTED ...
line vty 0 4
 exec-timeout 0 0
 logging synchronous
 login
... OUTPUT OMITTED ...
```

Why is your new employee unable to log into router R1 via Telnet?

a. The logging synchronous command requires that the login be authenticated by a AAA server. Therefore, the issue must be with the AAA server.
b. The exec-timeout 0 0 command causes an instantaneous timeout whenever someone attempts to log into the router.
c. Due to SSH being more secure, Telnet access is disabled by default.
d. No password is configured for the VTY lines.

Question #69

You are examining a router's running configuration and notice that the password for the VTY lines is in clear text:

```
...OUTPUT OMITTED...
```

```
line vty 0 4
 exec-timeout 0 0
 password cisco
 logging synchronous
 login
...OUTPUT OMITTED...
```

You want the VTY line password to be encrypted in the running configuration, as follows:

```
line vty 0 4
 exec-timeout 0 0
 password 7 02050D480809
 logging synchronous
 login
```

What command should you enter to encrypt the VTY line password?

a. Router(config-line)# **enable password-encryption**
b. Router(config)# **enable password-encryption**
c. Router(config-line)# **service encryption**
d. Router(config)# **service password-encryption**

Question #70

Which type of DNS record is used to map a hostname to an IPv6 address?
a. A
b. AAAA
c. SOA
d. MX

Question #71

Consider the following port security configuration:

```
Switch(config)# int gig 1/0/1
Switch(config-if)# switchport mode access
Switch(config-if)# switchport port-security
Switch(config-if)# switchport port-security maximum 3
Switch(config-if)# switchport port-security mac-address sticky
Switch(config-if)# switchport port-security violation restrict
```

What does the **restrict** option in the bottom command do?

a. The **restrict** option causes the port to go into an err-disable state if a port security violation occurs.
b. The **restrict** option causes the port to be administratively shutdown if a port security violation occurs.
c. The **restrict** option disables port security on this port.

d. The **restrict** option drops traffic from the fourth MAC address seen on this port, while permitting traffic from the first three MAC addresses learned on this port. However, the security violation counter does not get incremented.
e. The **restrict** option drops traffic from the fourth MAC address seen on this port, while permitting traffic from the first three MAC addresses learned on this port. Also, the security violation counter is incremented.

Question #72

If Spanning Tree Protocol (STP) operation fails in a network with redundant links, what symptoms could result? (Choose 3.)

a. broadcast storm
b. MAC address table corruption
c. duplex mismatch
d. duplicate frames received by the intended receiver

Question #73

You wish to remotely connect to a Cisco Catalyst 2960 switch. Which of the following parameters must be configured on the switch? (Choose 2.)
a. a default static route
b. an IP address
c. a default gateway
d. a routing protocol

Question #74

Which of the following features allows a Cisco Catalyst switch to create a copy of frames appearing on a switch port or in a VLAN, and send those copied frames out of a designated port?
a. SPAN
b. CEF
c. HSRP
d. VRRP

Question #75

Given a subnet of 172.16.56.0 /21, identify which of the following IP addresses belong to this subnet. (Select 2.)

a. 172.16.54.129
b. 172.16.62.255
c. 172.16.61.0
d. 172.16.65.255
e. 172.16.64.1

Question #76

What is the subnet address of the IP address 192.168.5.55 with a subnet mask of 255.255.255.224?
a. 192.168.5.0 /27
b. 192.168.5.16 /27
c. 192.168.5.32 /27
d. 192.168.5.48 /27
e. 192.168.5.64 /27

Question #77

Which of the following protocols allow you to view Layer 2 adjacent network devices from a Cisco router or Cisco Catalyst switch command prompt? (Choose 2.)
a. MLP
b. CDP
c. RTP
d. LLDP

Question #78

You are working for a company that will be using the 192.168.1.0 /24 private IP address space for IP addressing inside their organization.

They have multiple geographical locations and want to carve up the 192.168.1.0 /24 address space into subnets. Their largest subnet will need 13 hosts.

What subnet mask should you use to accommodate at least 13 hosts per subnet, while maximizing the number of subnets that can be created?

a. 255.255.255.248
b. 255.255.255.224
c. 255.255.255.252
d. 255.255.255.192
e. 255.255.255.240

Question #79

A customer is using a Class C network of 192.168.10.0 subnetted with a 28-bit subnet mask. How many assignable addresses are available in each of the subnets?

a. 32
b. 16
c. 30
d. 8
e. 14

Question #80
Which of the following are configurable versions of Simple Network Management Protocol (SNMP) within Cisco IOS? (Choose 3.)
a. v1
b. v1c
c. v2
d. v2c
e. v3
f. v3c

Question #81
An IP address of 192.168.0.100 /27 belongs to which of the following subnets?

a. 192.168.0.92
b. 192.168.0.128
c. 192.168.0.64
d. 192.168.0.96
e. 192.168.0.32

Question #82
Which of the following is considered to be an unreliable Transport Layer protocol?

a. IP
b. UDP
c. TCP
d. ICMP
e. PPP

Question #83
What subnet mask should be used to subnet the 192.168.10.0 network to support the number of subnets and IP addresses per subnet shown in the following topology?

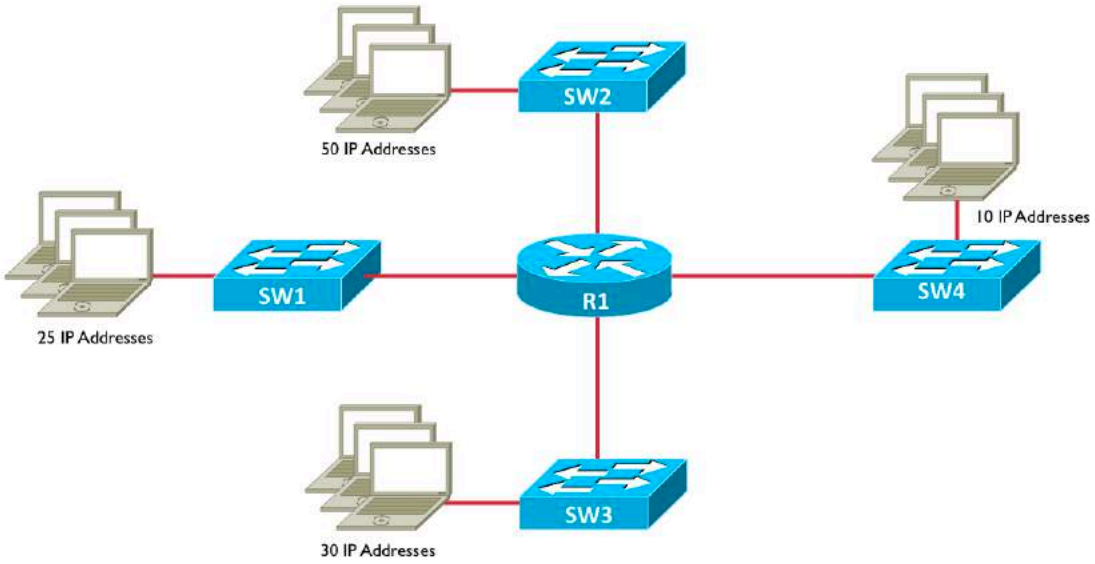

a. 255.255.255.0
b. 255.255.255.128
c. 255.255.255.192
d. 255.255.255.224
e. 255.255.255.240

Question #84

How can an IPv6 address of **2200:5678:0001:0000:0000:000A:0000:0001** can be abbreviated?

a. 22:5678:1::A:0:1
b. 2200:5678:1::A:0:1
c. 2200:5678:1::A::1
d. 22:5678:1::a:0:1
e. 2200:5678:1::A:0::1

Question #85

IPv6 unique local addresses are similar to IPv4 private IP addresses, because they cannot be routed over the public Internet. A unique local IPv6 address begins with which of the following patterns?

a. 2000::/3
b. FE80::/10
c. FF02::1:FF
d. FC00::/7
e. FF

Question #86

In an attempt to recover a lost password on a Cisco router, you issue a Break during the router's boot sequence. This takes you to the ROM Monitor prompt. From there, you want to set the configuration register such that the router's startup configuration will be ignored the next time it boots. To which of the following values should you set the configuration register?
a. 2102
b. 2142
c. 0x2142
d. 0x2102

Question #87

You issue the **ping 192.168.1.2** command from a router, and the response displayed on the router is:

```
M.M.M
```

What does this response indicate?

a. The router had to ARP for the MAC address of the next hop IP address.
b. The router is attempting to load balance across two links, and one of the links is not working.
c. The Ping packets needed to be fragmented, but the packets have their DF bit set (which says they cannot be fragmented).
d. The Ping is successful, and the alternating M and dot characters indicate the two directions of the bidirectional communication.

Question #88

What network architecture layers are combined in a **collapsed core** design?
a. Access and Distribution
b. Distribution and Core
c. Access and Core
d. Access, Distribution, and Core

Question #89

What type of Cisco Catalyst Switch port configuration allows a port to be an access port that supports two VLANs, if and only if one of the two VLANs is designated as a voice VLAN?
a. Voice VLAN Access Port
b. Voice Trunk Port
c. Multi-VLAN Access Port
d. 802.1p Access Port

Question #90

What type of queuing adds a priority queue to CB-WFQ?
a. ECN

b. LLQ
c. WRED
d. WFQ

Question #91

Which of the following are features of Point-to-Point Protocol (PPP)? [Choose 4.]
a. Authentication
b. Encryption
c. Error Detection and Correction
d. Logical Bundling of Multiple Links
e. Compression

Question #92

A ***Cisco Application Policy Infrastructure Controller - Enterprise Module*** (APIC-EM) uses Northbound APIs to connect to what?
a. Network Devices
b. Peer Controllers
c. Autonomous Controllers
d. Applications

Question #93

What Cisco technology allows you to interconnect multiple physical switches into a single logical switch?
a. SmartNet
b. Optimum Switching
c. Stackwise
d. Collapsed Core

Question #94

What technology allows an enterprise to more easily change their cloud provider (e.g. change from Microsoft to AWS)?
a. CloudFront
b. Intercloud Exchange
c. MP-BGP
d. APIC-EM

Question #95

A **Unique Local IPv6 Address**, which cannot be routed over the public Internet, begins with what hexadecimal prefix?
a. FE80::/10
b. FC00::/7
c. FF02::1:FF
d. 2000::/3

CCNA R&S (200-125) Ultimate Practice Exam Answers

Question #1

Which of the following IPv6 address ranges is used for *Global Unicast* IPv6 addresses?

a. 2000::/3
b. FF00::/8
c. FEC0::/10
d. FC00::/7
e. ::1
f. ::
g. FF02::1::FF00:0/104

Answer: a

IP version 6 (IPv6) address are 128-bits in length. A variety of IPv6 address types use specific ranges of IPv6 addresses. For example:

- **Global Unicast:** Addresses are in the range **2000::/3**. This type of address refers to a single device, and it can be routed over the Internet.
- **Multicast:** Addresses are in the range **FF00::/8**. This type of address refers to a multicast group.
- **Link Local:** Addresses are in the range **FEC0::/10**. This type of address is only valid on a local link (e.g. a link between two routers).
- **Unique Local:** Addresses are in the range **FC00::/7**. This is a private IP address, which cannot be routed on the Internet.
- **Loopback:** Address is **::1**. This address is used by a device to refer to itself. The loopback address is commonly called the localhost address.
- **Unspecified:** Address is **::**. This address is used when the source IPv6 address in a packet is not specified (i.e. all 128 bits are set to 0s).
- **Solicited-Node Multicast:** Addresses are in the range **FF02::1::FF00:0/104**. This address is a multicast IPv6 address, and it corresponds to a device's IPv6 address. Specifically, the right-most 24 bits are the 24 right-most bits in the device's IPv6 address.

Question #2

You are working with an enterprise router connecting out to two Internet Service Providers (ISPs). The router has a single link to each ISP. What type of topology is described by this scenario?

a. Single Homed
b. Dual Homed
c. Single Multihomed
d. Dual Multihomed

Answer: c

The scenario described in the question fits the definition of a **Single Multihomed** topology. The definitions of all four topologies are as follows:

- **Single Homed:** One link to one ISP.
- D**ual Homed:** Two or more links to one ISP.
- **Single Multihomed:** One link per ISP to two or more ISPs.
- **Dual Multihomed:** Two or more links per ISP to two or more ISPs.

Question #3

What is the range of assignable IP addresses for a subnet containing an IP address of 172.16.1.10 /19?

a. 172.16.0.1 – 172.16.31.254
b. 172.16.0.1 – 172.16.63.254
c. 172.16.0.0 – 172.16.31.255
d. 172.16.0.1 – 172.16.31.255
e. 172.16.0.0 – 172.16.63.254

Answer: a

To determine the subnets, assignable IP address ranges, and directed broadcast addresses created by the 19-bit subnet mask we perform the following steps:

Step #1: Identify the interesting octet (i.e. the octet that contains the first zero in the binary subnet mask).

In this question, we have a 19-bit subnet mask, which is written in binary as:
11111111 11111111 11100000 00000000

The interesting octet is the third octet, because the third octet (i.e. 11100000) is the first octet to contain a 0 in the binary.

Step #2: Identify the decimal value in the interesting octet of the subnet mask.

A 19-bit subnet mask can be written in dotted decimal notation as:
255.255.224.0

Since the third octet is the interesting octet, the decimal value in the interesting octet is 224.

Step #3: Determine the block size by subtracting the decimal value of the interesting octet from 256.

Block Size = 256 – 224 = 32

Step #4: Determine the subnets by counting by the block size in the interesting octet, starting at 0.

Placing a zero in the first interesting octet identifies the first subnet as:
172.16.0.0 /19

We then count by the block size (of 32) in the interesting octet (the third octet in this question) to determine the remaining subnets:

172.16.32.0 /19
172.16.64.0 /19
172.16.96.0 /19
172.16.128.0 /19
172.16.160.0 /19
172.16.192.0 /19
172.16.224.0 /19

Step #5: Identify the subnet address, the directed broadcast address, and the usable range of addresses.

Looking through the subnets created by the 19-bit subnet mask reveals that the IP address of 172.16.1.10 resides in the 172.16.0.0 /19 subnet.

The directed broadcast address, where all host bits are set to a 1, is 1 less than the next subnet address.

The next subnet address is 172.16.32.0. So, the directed broadcast address for the 172.16.0.0 /19 subnet is 1 less than 172.16.32.0, which is:
172.16.31.255

The usable IP addresses are all the IP addresses between the subnet address and the directed broadcast address. Therefore, in this example, the assignable IP address range for the 172.16.0.0 /19 network is:
172.16.0.1 – 172.16.31.254

Question #4

Consider the following topology and configuration for router R1.

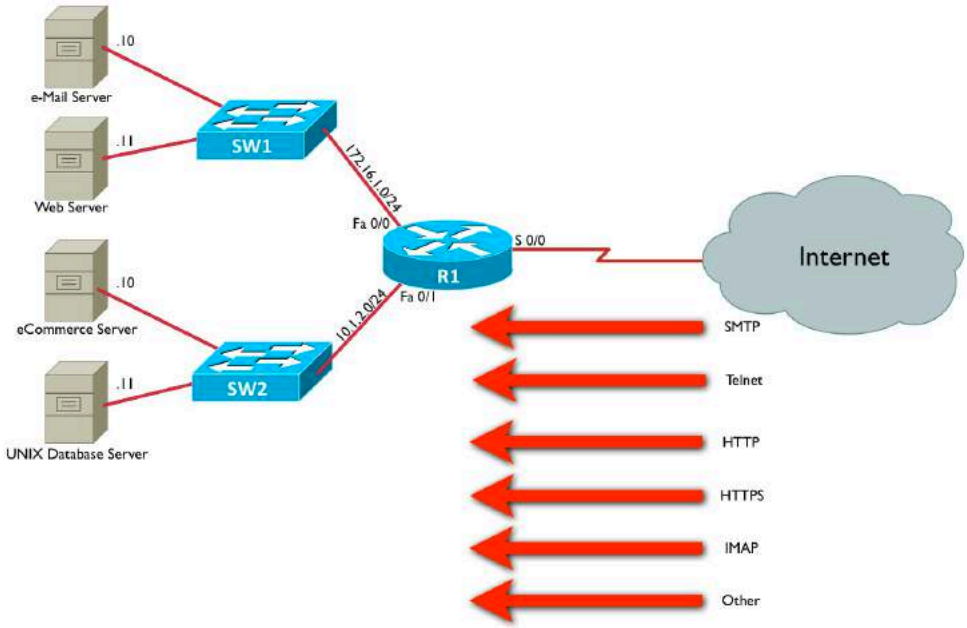


```
ip access-list extended TEST
 permit tcp any host 172.16.1.10 eq smtp
 deny ip any 172.16.1.0 0.0.0.255
 permit ip any 10.1.2.0 0.0.0.255
!
interface Serial 0/0
 ip access-group TEST in
```

What traffic from the Internet will be allowed to pass through router R1? (Choose all that apply.)

a. SMTP traffic destined for the e-mail server
b. HTTP traffic destined for the web server
c. HTTPS traffic destined for the eCommerce server
d. Telnet traffic destined for the UNIX database server

Answer: a, c, and d

Access control lists (ACLs) are processed top-down. Therefore, if a packet matches an access control entry (ACE) listed higher in an ACL, it will never be evaluated by ACEs residing lower in the ACL.

In this instance, the **permit tcp any host 172.16.1.10 eq smtp** command is allowing traffic from any source to send traffic to the e-Mail server (172.16.1.10) on port 25 (SMTP). Even though the next command, **deny ip any 172.16.1.0 0.0.0.255**, says that all traffic coming into the subnet of 172.16.1.0 /24 should be denied, that command will not deny SMTP traffic destined for the e-Mail server, because that traffic has already been matched and permitted. The **deny ip any**

172.16.1.0 0.0.0.255 command does, however, deny any other traffic destined for the 172.16.1.0 /24 subnet. Therefore, no Internet traffic destined for the web server (172.16.1.11 /24) will be permitted through router R1.

The **permit ip any 10.1.2.0 0.0.0.255** command allows all IP Internet traffic destined for the 10.1.2.0 /24 subnet to pass through router R1. Therefore, Internet traffic can reach the eCommerce server (10.1.2.10 /24) and the UNIX database server (10.1.2.11 /24).

Question #5

Which of the following is a set of wiring standards that describe the color coding of wires in RJ-45 connections?
a. CSMA/CD
b. Multimode
c. EIA/TIA 568
d. RS-232

Answer: c

The term EIA/TIA refers to two standards bodies, the Electronic Industries Alliance (EIA) and the Telecommunications Industry Association (TIA).

The EIA/TIA 568 standards specify the color coding of the eight wires in an unshielded twisted pair (UTP) Ethernet cable.

One of these standards is the T568A standard. However, it has largely been superseded by the newer T568B standard.

Question #6

You are configuring Network Address Translation (NAT). A PC on the inside of your network has an IP address of 192.168.1.10. However, since that IP address is a private IP address, it needs to be translated into an address that is routable on the public Internet. What NAT terminology is used to describe the PC's 192.168.1.10 IP address?

a. inside global address
b. outside local address
c. inside local address
d. outside global address

Answer: c
An IP address used in a NAT configuration could be one of the following four types of addresses:

(1) **Inside Local Address:** This is the address assigned to the device on the inside of the network, and the address is typically not routable on the public Internet.

(2) **Inside Global Address:** This is a publicly routable IP address that represents a device on the inside of a network.

(3) **Outside Local Address:** This type of NAT address is rarely used. It is a private IP address assigned to a device outside of a network. For example, let's say that a company had two sites, Site A and Site B, that communicate over the public Internet. Within each site, the company uses private IP addressing (e.g. the 192.168.1.0 /24 address space for Site A and the 172.16.1.0 /24 address space for Site B). If IP address 192.168.1.100 at Site A wants to communicate with a device at Site B with an IP address of 172.16.1.100, from the perspective of the device at Site A, 172.16.1.100 would be an outside local address.

(4) **Outside Global Address:** This is a publicly routable IP address assigned to a device outside a network. For example, if a device inside a network attempts to communicate with a web server on the Internet, the web server's IP address would be an outside global address.

As a memory aid, remember that "inside" always refers to a device on the inside of a network, while "outside" always refers to a device outside of a network. Also, "local" refers to an IP address not routable on the public Internet. Remember that "loco" is the Spanish word meaning crazy. That's what a local address is. It's a crazy made up address (i.e. not routable on the public Internet). The "g" in "global" can be used to remind you of the "g" in "good," because a global address is a good address (i.e. routable on the public Internet).

Question #7

What is the default Holdtime for HSRP?
a. 1 second
b. 3 seconds
c. 5 seconds
d. 10 seconds

Answer: d

By default, **Hot Standby Router Protocol** (HSRP) has a default Holdtime timer value of 10 seconds. The default Hello timer value is 3 seconds. Also, there is a requirement that the value of the Holdtime timer be at least three times the value of the Hello timer.

Question #8

In the following topology, how many collision domains are represented?

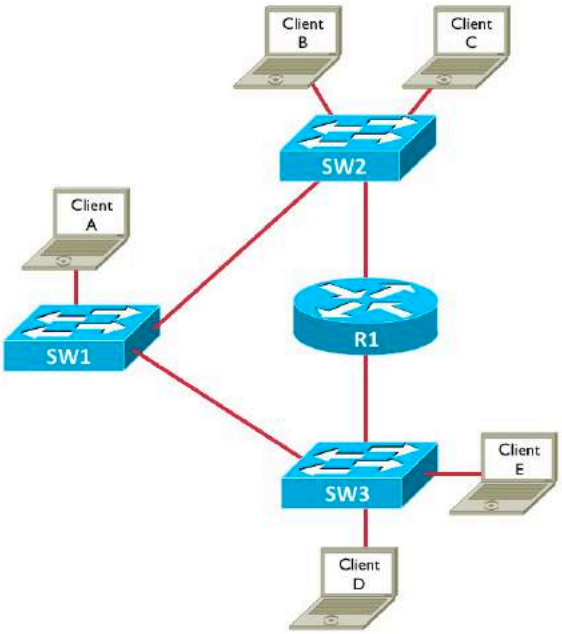

a. 3
b. 4
c. 5
d. 9

Answer: d

Each port on a switch is in its own collision domain. Similarly, each port on a router is in its own collision domain. Therefore, in the topology shown there are a total of 9 collision domains, as illustrated in the following figure:

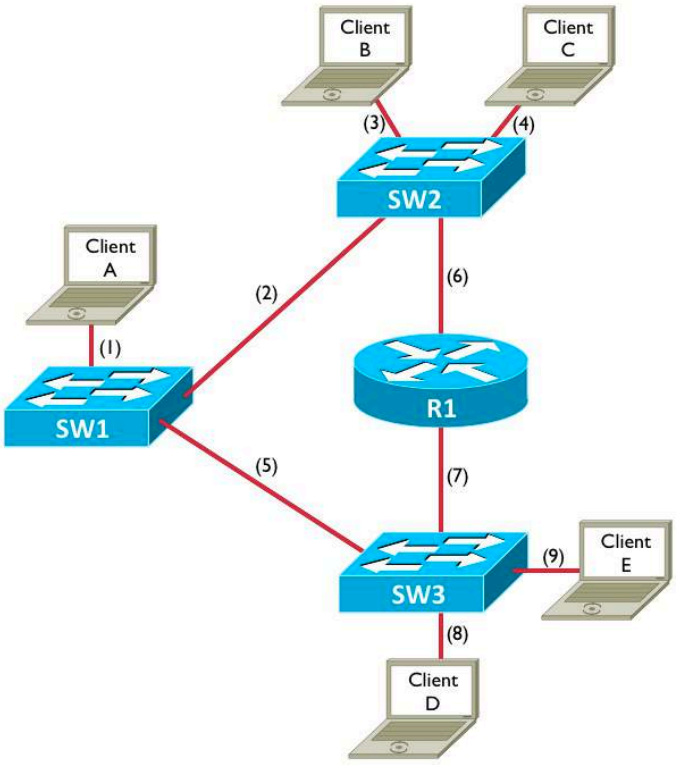

Question #9

Consider the following port security configuration:

```
Switch(config)# int gig 1/0/1
Switch(config-if)# switchport mode access
Switch(config-if)# switchport port-security
Switch(config-if)# switchport port-security maximum 3
Switch(config-if)# switchport port-security mac-address sticky
Switch(config-if)# switchport port-security violation shutdown
```

What does the **shutdown** option in the bottom command do?
a. The **shutdown** option causes the port to go into an err-disable state if a port security violation occurs.
b. The **shutdown** option causes the port to be administratively shutdown if a port security violation occurs.
c. The **shutdown** option disables port security on this port.
d. The **shutdown** option drops traffic from the fourth MAC address seen on this port, while permitting traffic from the first three MAC addresses learned on this port. However, the security violation counter does not get incremented.
e. The **shutdown** option drops traffic from the fourth MAC address seen on this port, while permitting traffic from the first three MAC addresses learned on this port. Also, the security violation counter is incremented.

Answer: a

The **switchport mode access** command places the port in access mode. So, the port carries traffic for only one VLAN. A port must be in access mode in order to enable port security on that port.

The **switchport port-security** command enables port security on the port.

The **switchport port-security maximum 3** command specifies the maximum number of allowed MAC addresses learned off of a port as three. If an additional MAC address is seen off of the port, a port violation occurs.

The **switchport port-security mac-address sticky** command allows the maximum number of MAC addresses learned off of a port to be entered into a switch's running configuration. You should copy the running configuration to the startup configuration if you want the learned MAC addresses to be retained after a switch reboot. Alternately, you could specify one or more MAC addresses allowed off of an interface with the command **switchport port-security mac-address *MAC_Address***.

The **switchport port-security violation shutdown** command says that the port will be placed into an err-disable state if a port violation occurs. To bring a port out of err-disable state (after clearing the condition that caused the violation), you can administratively shutdown the port (with the **shutdown** command) and then administratively bring the port back up (with the **no shutdown** command). In addition to the **shutdown** option, the **switchport port-security violation** command also includes the **protect** and **restrict** options. The **protect** option drops traffic from any unpermitted MAC addresses, while permitting traffic from the permitted MAC addresses learned on this port. However, the security violation counter does not get incremented. The **restrict** option also blocks traffic from unpermitted MAC addresses, while allowing traffic from permitted MAC addresses. However, the **restrict** option does increment the security violation counter. A port's security violation counter can be displayed with the **show port-security command**.

Question #10

You connect a laptop to an available port on a Cisco Catalyst switch. By default, how long does it take that port to start forwarding traffic?
a. 0 seconds
b. 15 seconds
c. 30 seconds
d. 50 seconds

Answer: c

By default, Spanning Tree Protocol (STP) is enabled on a switch port. Even though it takes (by default) 50 seconds for a Blocking switch port to transition to Forwarding when the switch's Root Port is no longer the best port to get back to the Root Bridge, in this question, the port was available. Therefore, it did not need to spend 20 seconds the Blocking State. Instead, it spends 15 seconds in the Listening State and 15 seconds in the Learning State, for a total of 30 seconds.

Question #11

Consider the output of the **show controllers** commands issued on routers R1 and R2. Of these two routers, which one should have a clock rate set for the specified serial interface, and how do you know from the output?

```
R1#show controllers serial 0/2/0
Interface Serial0/2/0
Hardware is GT96K
DTE V.35idb at 0x47FFCA4C, driver data structure at 0x48004290
wic_info 0x480048BC
Physical Port 5, SCC Num 5
MPSC Registers:
MMCR_L=0x000304C0, MMCR_H=0x00000000, MPCR=0x00000100
CHR1=0x00FE007E, CHR2=0x00000000, CHR3=0x0000064A, CHR4=0x00000000
CHR5=0x00000000, CHR6=0x00000000, CHR7=0x00000000, CHR8=0x00000000
CHR9=0x00000000, CHR10=0x00003008
SDMA Registers:

... OUTPUT OMITTED ...

R2#show controllers serial 1/0
CD2430 Slot 1, Port 0, Controller 0, Channel 0, Revision 19
Channel mode is synchronous serial
idb 0x466767F0, buffer size 1524, V.35 DCE cable, clockrate 128000

Global registers
  rpilr 0x2, rir 0x2, risr 0x0, rfoc 0x0, rdr 0x0
  tpilr 0x1, tir 0x0, tisr 0x68, tftc 0x0, tdr 0x90
  mpilr 0x3, mir 0x1, misr 0x60
  bercnt 0xFF, stk 0x0
Per-channel registers for channel 0

... OUTPUT OMITTED ...
```

a. R1 should have the clock rate set, because the DTE end of the serial cable connects to R1.
b. R1 should have the clock rate set, because the DCE end of the serial cable connects to R1.
c. R2 should have the clock rate set, because the DTE end of the serial cable connects to R2.
d. R2 should have the clock rate set, because the DCE end of the serial cable connects to R2.

Answer: d

A serial cable has a DTE (Data Terminal Equipment) connector at one end of the cable and a DCE (Data Communications Equipment) connector at the other end of the cable. A router's serial interface connected to the DCE side of a cable should provide clocking for the circuit. (**NOTE**: Clocking is needed, because the routers are doing synchronous communication over the serial link, and the clocking allows each router to determine when one bit stops and another bit starts.)

In the output of the **show controllers** command, router R2 shows the text:

```
V.35 DCE cable, clockrate 128000
```

Not only does the **`clockrate`** keyword indicate that clocking should be set on router R2's serial interface, the reference to the **`DCE cable`** also indicates that this interface should be doing clocking.

Question #12

A router resides at what layer of the TCP/IP Model?

a. Application
b. Transport
c. Internet
d. Network Interface
e. Session

Answer: c

At its essence, a router makes forwarding decisions based on a Layer 3 (i.e. Layer 3 of the OSI Model) address (e.g. an IP address). Therefore, a router resides at the Internet layer of the TCP/IP Model, which corresponds to the Network Layer (i.e. Layer 3) of the OSI Model.

The TCP/IP Model contains four layers (listed from top to bottom): Application, Transport, Internet, and Network Interface. The OSI Model contains seven layers (listed from top to bottom): Application, Presentation, Session, Transport, Network, Data Link, and Physical.

The Application Layer of the TCP/IP Model maps to the Session, Presentation, and Application Layers of the OSI Model. Therefore, this layer contains protocols used to setup, maintain, and tear down a session (e.g. the *Session Initiation Protocol (SIP)*), in addition to data formatting (e.g. ASCII), and network services (e.g. e-mail).

The Transport Layer of the TCP/IP Model maps to the Transport Layer of the OSI Model, and includes protocols such as TCP and UDP.

The Internet Layer of the TCP/IP Model maps to the Network Layer of the OSI Model. Routers reside at this layer, as well as IP (Internet Protocol).

The Network Interface Layer of the TCP/IP Model maps to the Physical and Data Link Layers of the OSI Model. Therefore, Ethernet hubs and Ethernet switches reside at the Network Interface Layer. Note that the Network Interface Layer is also referred to as the *Network Access Layer* in some literature.

Question #13

Which of the following is not a Rapid Spanning Tree Protocol (RSTP) port state?
a. Listening
b. Discarding
c. Learning
d. Forwarding

Answer: a

White traditional IEEE 802.1D Spanning Tree Protocol (STP) does have a Listening state (where a port remains for 15 seconds), Rapid Spanning Tree Protocol (RSTP) does not have this state.

Question #14

In the diagram below, Client A is sending a packet to Host 1. As the packet is coming into the Fa 0/0 interface on router R2, what is the source IP address in the packet's header?

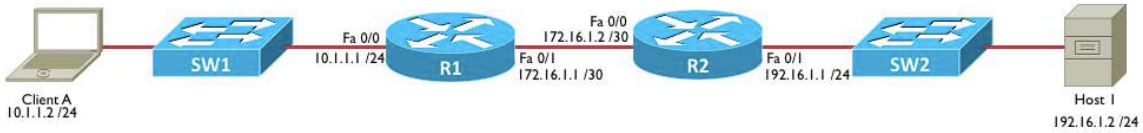

a. 10.1.1.1
b. 172.16.1.2
c. 192.16.1.1
d. 10.1.1.2
e. 172.16.1.1
f. 192.16.1.2

Answer: d

As the frame is being sent from Client A to Host 1, the source and destination IP addresses never change. The IP address of Client A is always the source IP

address, and the IP address of Host 1 is always the destination IP address, unless *Network Address Translation (NAT)* is being performed.

Question #15

In the diagram below, Client A is sending a packet to Host 1. Which devices will make a packet forwarding decision based on a destination IP address? (Choose 2.)

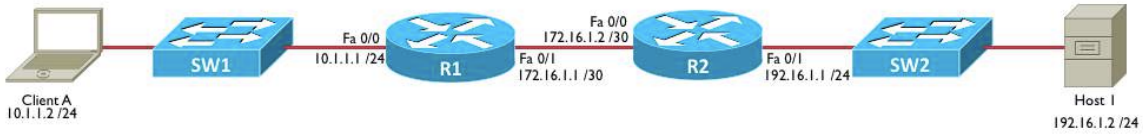

a. switch SW1
b. router R1
c. router R2
d. switch SW2

Answer: b and c

Traditional Ethernet switches are considered to be Layer 2 (i.e. Data Link Layer) devices, and they make forwarding decisions based on Layer 2 MAC (Media Access Control) addresses.

Routers, however, are considered to be Layer 3 (i.e. Network Layer) devices, and they can make forwarding decisions based on Layer 3 IP (Internet Protocol) addresses.

Question #16

Which of the following is considered to be a reliable Transport Layer protocol?

a. IP
b. UDP
c. TCP
d. ICMP
e. PPP

Answer: c

IP (Internet Protocol) and ICMP (Internet Control Message Protocol) are Network Layer (i.e. Layer 3) protocols.

PPP (Point-to-Point Protocol) is a Data Link Layer (i.e. Layer 2) protocol.

Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) are both Transport Layer (i.e. Layer 4) protocols. However, TCP uses acknowledgements to confirm receipt of data, while UDP does not confirm receipt of data. Therefore, TCP is considered to be a reliable, connection-oriented protocol, while UDP is considered to be an unreliable, connectionless protocol.

Question #17

You are configuring eBGP on router R1 seen in the figure. You want to form a BGP neighborship with router R2. Which command would you issue from within BGP router configuration mode on router R1?

a. `neighbor 198.51.100.2 remote-as 64495`
b. `neighbor 198.51.100.2 remote-as 64500`
c. `neighbor 198.51.100.1 remote-as 64495`
d. `neighbor 198.51.100.1 remote-as 64500`

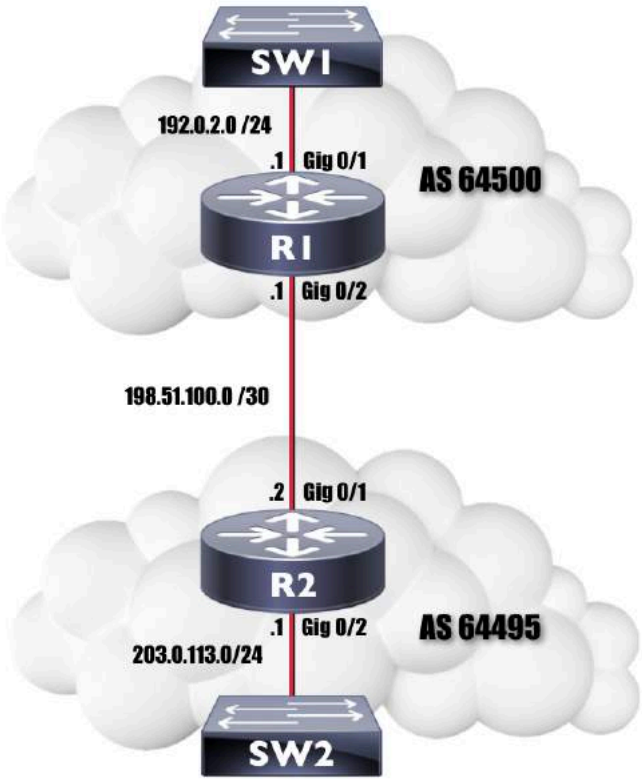

Answer: a

The BGP command used to form a neighborship is:

neighbor *remote_IP* **remote-as** *remote_AS*

In this question, from the perspective of router R1, the remote IP address is an IP address on router R2 (e.g. 198.51.100.2). Also, the remote AS, from the perspective of router R1 is 64495 (i.e. router R2's AS). Therefore, the command issued from BGP router configuration mode on router R1 sould be:

```
neighbor 198.51.100.2 remote-as 64495
```

Question #18

You are assigning IP addresses to hosts in the 192.168.4.0 /26 subnet. Which two of the following IP addresses are assignable IP addresses that reside in that subnet?

a. 192.168.4.0
b. 192.168.4.63
c. 192.168.4.62
d. 192.168.4.32
e. 192.168.4.64

Answer: c and d

To determine subnets and usable address ranges created by the 26-bit subnet mask we perform the following steps:

Step #1: Identify the interesting octet (i.e. the octet that contains the first zero in the binary subnet mask).

In this question, we have a 26-bit subnet mask, which is written in binary as:
11111111 11111111 11111111 11000000

The interesting octet is the forth octet, because the forth octet (i.e. 11000000) is the first octet to contain a 0 in the binary.

Step #2: Identify the decimal value in the interesting octet of the subnet mask.

A 26-bit subnet mask can be written in dotted decimal notation as:
255.255.255.192

Since the forth octet is the interesting octet, the decimal value in the interesting octet is 192.

Step #3: Determine the block size by subtracting the decimal value of the interesting octet from 256.

Block Size = 256 – 192 = 64

Step #4: Determine the subnets by counting by the block size in the interesting octet, starting at 0.

Placing a zero in the first interesting octet identifies the first subnet as:
192.168.4.0 /26

We then count by the block size (of 64) in the interesting octet (the forth octet in this question) to determine the remaining subnets:

192.168.4.64 /26
192.168.4.128 /26
192.168.4.192 /26

Step #5:
This question is asking about the 192.168.4.0 /26 subnet. From the above list of subnets, we can determine that the assignable range of IP addresses for this subnet is **192.168.4.1 – 192.168.4.62**. We can also determine that **192.168.4.0 is the network address**, and **192.168.4.63 is the directed broadcast address**.

From the assignable range of IP addresses we have calculated, we can determine that the two assignable IP addresses given as options in this question are: **192.168.4.62** and **192.168.4.32**.

Question #19

Which of the following are security best practices when setting up a Cisco router? (Choose all the apply.)

a. Enable the password encryption service.
b. Enable Telnet for remote access.
c. Create a welcome.
d. Use the **enable password** command to restrict access to the router's privileged mode.
e. Issue the **exec-timeout 0 0** command to log out any idle sessions.

Answer: a

While enabling the password encryption service does not provide strong encryption for line passwords (e.g. passwords for the con or vty lines), it does prevent someone from "shoulder surfing," and seeing an unencrypted password displayed on-screen.

While Telnet can provide remote access to a router, Telnet sends data (e.g. passwords) in clear text. Therefore, for security purposes, Cisco recommends that Secure Shell (SSH) be used to remotely connect to a router, as opposed to Telnet.

Cisco routers allow you to configure messages to users in the form of banners. For example, you can create a login banner or a message of the day (MOTD) banner. However, these banners should not offer a “welcome” message, because if a malicious user attempted to gain access to your router, they could try to defend their actions by pointing out that they were welcomed into the system.

Older versions of Cisco IOS used the **enable password** command to configure the privileged mode password. However, the **enable password** command, by default, displays a password in clear text in a router’s configuration. A much more secure approach is to use the **enable secret** command to set the password required for privileged mode access. When the **enable secret** command is used, the privileged mode password appears in a router’s configuration as a hashed value of the password.

The **exec-timeout *minutes seconds*** command can be used to log a user out of the router if the user has been idle for a specified amount of time. While it is a good security practice to log a user out after a period of inactivity, if you specify a *zero* for both the ***minutes*** and ***seconds*** values in the command (i.e. **exec-timeout 0 0**), the connection is never timed out.

Question #20

A host in your network has been assigned an IP address of 192.168.181.182 /25. What is the subnet to which the host belongs?

a. 192.168.181.128 /25
b. 192.168.181.0 /25
c. 192.168.181.176 /25
d. 192.168.181.192 /25
e. 192.168.181.160 /25

Answer: a

To determine subnets and usable address ranges created by the 25-bit subnet mask we perform the following steps:

Step #1: Identify the interesting octet (i.e. the octet that contains the first zero in the binary subnet mask).

In this question, we have a 25-bit subnet mask, which is written in binary as:
11111111 11111111 11111111 10000000

The interesting octet is the forth octet, because the forth octet (i.e. 10000000) is the first octet to contain a 0 in the binary.

Step #2: Identify the decimal value in the interesting octet of the subnet mask.

A 25-bit subnet mask can be written in dotted decimal notation as:
255.255.255.128

Since the forth octet is the interesting octet, the decimal value in the interesting octet is 128.

Step #3: Determine the block size by subtracting the decimal value of the interesting octet from 256.

Block Size = 256 – 128 = 128

Step #4: Determine the subnets by counting by the block size in the interesting octet, starting at 0.

Placing a zero in the first interesting octet identifies the first subnet as:
192.168.181.0 /25

We then count by the block size (of 128) in the interesting octet (the forth octet in this question) to determine the remaining subnets, or in this case just a single additional subnet.

192.168.181.128 /25

Now that we have our two subnets identified, we can determine the subnet in which the IP address of 192.168.181.182 resides.

Since the usable range of IP addresses for the 192.168.181.128 /25 network is 192.168.181.129 – 192.168.181.254 (because 192.168.181.128 is the network address, and 192.168.181.255 is the directed broadcast address), and since 192.168.181.182 is in that range, the subnet to which 192.168.181.182 /25 belongs is:
192.168.181.128 /25

Question #21

The PING utility uses which of the following protocols to check for network connectivity to an IP address?

a. UDP
b. TCP
c. IGMP
d. ICMP
e. OSPF

Answer: d

UDP (User Datagram Protocol) and **TCP (Transmission Control Protocol)** are Layer 4 (i.e. Transport Layer) protocols used for transmitting data.

IGMP (Internet Group Management Protocol) is a multicast protocol. For example, if a PC wants to join a multicast group, it can send an ICMP Join message to its upstream router.

OSPF (Open Shortest Path First) is a routing protocol, which routers can use to exchange IP route information.

ICMP (Internet Control Message Protocol) is a Layer 3 (i.e. Network Layer) protocol. ICMP is used by the PING utility to check for network connectivity. For example, if a PC Pings a remote host, the PC sends an **ICMP Echo Request** message to that remote host. If the PC receives an **ICMP Echo Reply** message from that host, the PC can conclude the remote host is reachable over the network.

The **TRACEROUTE (or TRACERT)** utility can trace the path from a source IP address to a destination IP address through a series of routers, and it can also use ICMP. Although there are different implementations of the TRACEROUTE (or TRACERT) feature, in general, when a device issues the **traceroute** or **tracert** command along with an IP address, a packet (which could be a UDP-based segment) is sent to that destination IP address. However, the Time-to-Live (TTL) value in the IP header is set to a one. Therefore, when the packet enters the first router along the path to the destination IP address, the TTL value is decremented to zero, and the packet is discarded. However, the router that discarded the packet can let the sender know that the packet was discarded by sending an **ICMP Time to Live Exceeded In Transit** message to the sender. The sender then resends a packet to the destination IP address. Although, this time, the TTL is set to a two, and this process continues until the source IP address eventually reaches the destination IP address. This approach to incrementing the TTL value by one each time a packet is sent allows each subsequent packet to reach one router further in the network than the previous packet, allowing the sender to get information (e.g. round-trip delay) from each router along the path to the destination network. This can help the sender determine where excessive delay might exist in the path, which can be valuable information for troubleshooting.

Question #22

On an IEEE 802.1Q trunk, what do we call the VLAN that is not tagged by 802.1Q?
a. Access VLAN

b. Voice VLAN
c. Trunk VLAN
d. Native VLAN
e. Passive VLAN

Answer: d

In an IEEE 802.1Q trunk, frames in the **Native VLAN** do not receive any tagging by the trunk. However, frames in other VLANs receive four Bytes of tagging. As a result, it's very important for switches at each end of a trunk to agree on the Native VLAN. Otherwise, frames belonging to one switch's Native VLAN could "hop" to a different VLAN (i.e. the Native VLAN of the other switch).

Question #23

You wish to Telnet into a Layer 2 Cisco Catalyst switch. Which of the following must be configured on the switch before you can Telnet into the switch?

a. You must issue the **no switchport** command on one of the interfaces.
b. You must assign an IP address to the switch.
c. You must enable IP routing on the switch.
d. You must issue the **service telnet** global configuration command.
e. You must configure a *router-on-a-stick* topology to provide a Layer 3 path to the switch.

Answer: b

Even though a Layer 2 Ethernet switch does not need an IP address for routing traffic, it does need an IP address for remote connectivity into the switch (e.g. connecting via Telnet or SSH).

The **no switchport** command is used on a Layer 3 switch to convert a Layer 2 switched port into a Layer 3 routed port. So, this command is not needed for Telnet connectivity.

There is no globally enabled Telnet service.

Neither IP routing nor a router-on-a-stick topology need to be configured to support Telnet connectivity into a Layer 2 switch. Instead, a Layer 2 switch can be configured with an IP address and a default gateway (similar to a PC) in order to have reachability to remote networks.

Question #24

You are working with a Class B network with the private IP address of 172.16.0.0 /16. You need to maximize the number of broadcast domains, where each

broadcast domain can accommodate 1000 hosts. What subnet mask should you use?

a. /22
b. /23
c. /24
d. /25
e. /26

Answer: a

In addition to testing your knowledge of subnetting, this question is also making sure you understand that a subnet is a broadcast domain. This should not be confused with a collision domain (i.e. each port on a switch is in its own collision domain).

To determine how many host bits are required to support 1000 hosts, we can create a table from the following formula:

Number of Hosts = 2^h – 2, where *h* is the number of host bits

From this formula, we can create the following table:

1 Host Bit => 0 Hosts
2 Host Bits => 2 Hosts
3 Host Bits => 6 Hosts
4 Host Bits => 14 Hosts
5 Host Bits => 30 Hosts
6 Host Bits => 62 Hosts
7 Host Bits => 126 Hosts
8 Host Bits => 254 Hosts
9 Host Bits => 210 Hosts
10 Host Bits => 1022 Hosts

This table tells us that a subnet with 10 host bits will accommodate the requirement of 1000 hosts. If we have 10 host bits, then we have a **22-bit subnet** mask (i.e. 32 – 10 = 22). Also, by not using more host bits than we need, we are maximizing the number of subnets that can be created.

Question #25

CDP (Cisco Discovery Protocol) resides at which layer of the OSI Model?

a. Layer 1
b. Layer 2
c. Layer 3

d. Layer 4
e. Layer 5

Answer: b

CDP is a Cisco-proprietary protocol that can be used to discover (and collect information about) Layer 2 adjacent (i.e. directly attached) Cisco devices. Since CDP is a Layer 2 protocol, Cisco devices can discover one another even if they do not have any IP addresses configured. Devices advertising themselves via CDP (which can be administratively disabled) send their advertisements to the Layer 2 multicast address of 01-00-CC-CC-CC.

CDP is also frequently used with Cisco IP telephony. For example, CDP can be used by a Cisco IP Phone to tell an attached Cisco Catalyst switch how much inline power it needs (which the switch can provide via Power over Ethernet (PoE)). Also, the attached Cisco Catalyst switch can use CDP to tell the IP phone to which VLAN the IP phone belongs.

A Cisco Catalyst switch can also use CDP for Quality of Service (QoS) purposes. Specifically, a Cisco Catalyst switch can be configured to trust QoS markings, if and only if those markings are coming from a Cisco IP Phone. CDP is the protocol used by a Cisco Catalyst switch to confirm the attached device is a Cisco IP Phone.

Question #26

Which of the following technologies allows VPN tunnels to be setup and torn down on an as-needed basis?
a. IPsec
b. DMVPN
c. GRE
d. PPP

Answer: b

Dynamic Multipoint VPN (DMVPN) is a technology that allows Virtual Private Network (VPN) tunnels to be setup and torn down between locations on an as-needed basis. A couple of the underlying technologies used by DMVPNs are:

- **Multipoint GRE (mGRE):** Allows a single router interface to have multiple GRE tunnels.
- **Next Hop Resolution Protocol (NHRP):** Allows an interface configured for mGRE to discover the address of a device at the far end of a tunnel.

Question #27

Typically, HTTP uses what Transport Layer protocol and what port number?

a. UDP port 80
b. TCP port 443
c. UDP port 443
d. TCP port 80
e. TCP port 143

Answer: d

Hypertext Transfer Protocol (HTTP) typically uses TCP port 80 and is used to send messages to a web server.

Hypertext Transfer Protocol over TLS/SSL (HTTPS) typically uses TCP port 443 and is a secure version of HTTP.

Internet Message Access Protocol (IMAP) typically uses TCP port 143 and is used to retrieve e-mail from an e-mail server.

Question #28

BPDUGuard should be only be configured on ports that have which feature enabled?
a. RootFast
b. PortFast
c. RootGuard
d. UplinkFast

Answer: b

BPDUGuard is a Spanning Tree Protocol (STP) feature that causes a port to go into an Errdisable state if a Bridge Protocol Data Unit (BPDU) arrives on a port configured for BPDUGuard. This can help prevent a Layer 2 topological loop from forming over a port configured for PortFast, which is an STP feature that bypasses the STP Listening and Learning states on a port when a device (e.g. a PC) connects to that port. Therefore, you only need to configure BPDUGuard on ports that have PortFast enabled.

Question #29

A PC has just booted up and wants to communicate with a host on a remote subnet. The PC knows the IP addresses of its default gateway and the remote host. However, in order to properly construct a frame, the PC needs an appropriate destination MAC address. What MAC address does the PC need to learn, and what protocol will the PC use to learn that MAC address?

a. the MAC address of the PC's default gateway, learned via ARP
b. the MAC address of the remote host, learned via DNS
c. the MAC address of the PC's default gateway, learned via DNS
d. the MAC address of the remote host, learned via ARP

Answer: a

Source and destination MAC addresses are rewritten at each router hop along the path from the source to the destination (if the router interface is Ethernet-based). Therefore, the frame sent by the PC should have a destination MAC address of the PC's default gateway (i.e. the PC's next-hop router). However, the source and destination IP addresses do not change (unless some other service, such as Network Address Translation (NAT), is in use).

The PC can learn the MAC address of its default gateway using ARP (Address Resolution Protocol).

While ARP provides IP address to MAC address translation, DNS (Domain Name System) provides domain name to IP address translation.

Question #30

What is the directed broadcast address of a subnet containing an IP address of 172.16.1.10 /19?

a. 172.16.15.255
b. 172.16.31.255
c. 172.16.255.255
d. 172.16.95.255
e. 172.16.0.255

Answer: b

To determine the subnets, assignable IP address ranges, and directed broadcast addresses created by the 19-bit subnet mask we perform the following steps:

Step #1: Identify the interesting octet (i.e. the octet that contains the first zero in the binary subnet mask).

In this question, we have a 19-bit subnet mask, which is written in binary as:
11111111 11111111 11100000 00000000

The interesting octet is the third octet, because the third octet (i.e. 11100000) is the first octet to contain a 0 in the binary.

Step #2: Identify the decimal value in the interesting octet of the subnet mask.

A 19-bit subnet mask can be written in dotted decimal notation as:
255.255.224.0

Since the third octet is the interesting octet, the decimal value in the interesting octet is 224.

Step #3: Determine the block size by subtracting the decimal value of the interesting octet from 256.

Block Size = 256 – 224 = 32

Step #4: Determine the subnets by counting by the block size in the interesting octet, starting at 0.

Placing a zero in the first interesting octet identifies the first subnet as:
172.16.0.0 /19

We then count by the block size (of 32) in the interesting octet (the third octet in this question) to determine the remaining subnets:

172.16.32.0 /19
172.16.64.0 /19
172.16.96.0 /19
172.16.128.0 /19
172.16.160.0 /19
172.16.192.0 /19
172.16.224.0 /19

Step #5: Identify the subnet address, the directed broadcast address, and the usable range of addresses.

Looking through the subnets created by the 19-bit subnet mask reveals that the IP address of 172.16.1.10 resides in the 172.16.0.0 /19 subnet.

The directed broadcast address, where all host bits are set to a 1, is 1 less than the next subnet address.

The next subnet address is 172.16.32.0. So, the directed broadcast address for the 172.16.0.0 /19 subnet is 1 less than 172.16.32.0, which is:
172.16.31.255

The usable IP addresses are all the IP addresses between the subnet address and the directed broadcast address. Therefore, in this example, the assignable IP address range for the 172.16.0.0 /19 network is:
172.16.0.1 – 172.16.31.254

Question #31

Of the following protocols, select the ones that are UDP-based. (Choose 2)

a. SSH
b. NTP
c. IMAP4
d. DHCP
e. POP3

Answer: b and d

Secure Shell (SSH) is a TCP-based protocol (port 22) used to securely establish a remote console connection with a host (e.g. a router, switch, or server).

Network Time Protocol (NTP) is a UDP-based protocol (port 123) used for time synchronization.

Internet Message Access Protocol version 4 (IMAP4) is a TCP-based protocol (port 143) used by an e-mail client to retrieve messages from an e-mail server.

Dynamic Host Configuration Protocol (DHCP) is a UDP-based protocol (port 67) used by a network client to automatically obtain IP address information from a DHCP server.

Post Office Protocol version 3 (POP3) is a TCP-based protocol (port 110) used by an e-mail client to retrieve messages from an e-mail server.

Other common protocols that are TCP-based include: FTP (ports 20 and 21), SFTP (port 22), SCP (port 22), Telnet (port 23), SMTP (port 25), HTTP (port 80), NNTP (port 119), LDAP (port 389), HTTPS (port 443), and RDP (port 3389).

Other common protocols that are UDP-based include: TFTP (port 69), SNMP (161), and SNTP (port 123).

Note that DNS (port 53) and RTSP (port 554) can use either TCP or UDP as their Layer 4 transport protocol.

Question #32

Consider the following topology. Client F has just booted up, and its MAC address has not been learned by switch SW1 (as seen in the output of **the show mac address-table** command). If Client A sends a frame destined for Client F's MAC address, what will switch SW1 do with the frame? (Assume all switch ports are assigned to VLAN 100.)

a. SW1 will flood the frame out of all of its ports.
b. SW1 will flood the frame out of all of its ports, other than the port on which the frame was received.
c. SW1 will drop the frame, because there is no entry for the destination MAC address in switch SW1's MAC address table.
d. SW1 will forward the frame out of port Gi 1/0/2 only.
e. SW1 will forward the frame out of port Gi 1/0/1 only.

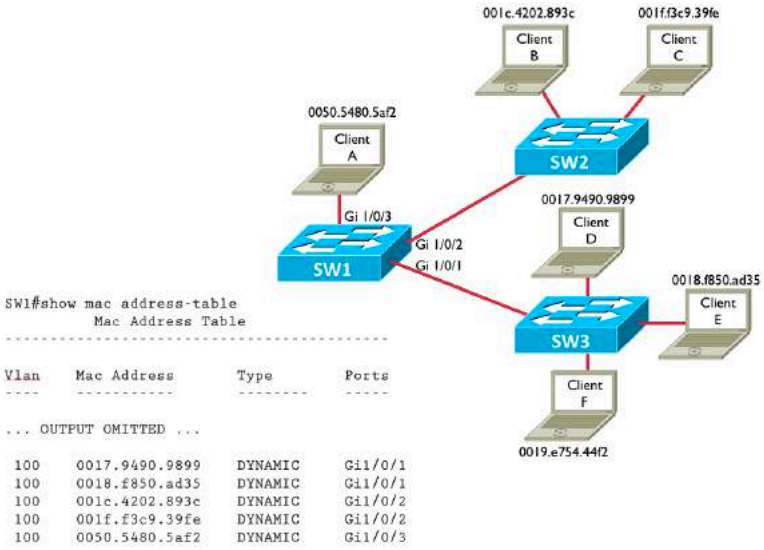

Answer: b

Layer 2 switches forward frames based on destination MAC addresses. These addresses can be statically configured in a switch, or they can be dynamically learned. In this question, switch SW1 has learned all MAC addresses on the network, other than the MAC address of Client F.

When a switch receives a unicast frame destined for MAC address known to the switch, the switch only forwards that frame out of the port off of which the destination MAC address has been learned.

However, if a unicast frame is destined for a MAC address not known by a switch, the switch will flood the frame out all of the switch ports (belonging to the same VLAN as the source port), other than the port on which the frame was received. This flooding behavior not only applies to unknown unicast traffic, but also to broadcast traffic, and in some instances, multicast traffic.

Question #33

A customer is using a Class C network of 192.168.10.0 subnetted with a 28-bit subnet mask. How many subnets can be created by using this subnet mask?

a. 32
b. 16
c. 30
d. 8
e. 14

Answer: b

The subnet in this question is a Class C network, because there is a 192 in the first octet. A class C network has a *natural mask* of 24 bits. However, this network has a 28-bit subnet mask. Therefore, we have 4 *borrowed bits*, which are network bits added to a network's natural mask (i.e. 28 – 24 = 4). The number of subnets can be calculated as follows:

Number of Subnets = 2^s, where *s* is the number of borrowed bits.

Therefore, in this question, the number of created subnets is 16:

Number of Subnets = $2^4 = 16$

Question #34

A PC requires which of the following configuration parameters in order to Telnet to a host on a remote subnet? (Choose 3.)

a. IP address
b. subnet mask
c. routing protocol
d. VLAN assignment
e. default gateway

Answer: a, b, and e

A PC typically gets its IP configuration from a DHCP server. This configuration includes such information as the PC's IP address, default gateway, subnet mask, and DNS server, among other parameters.

Some of these parameters are required for the PC to communicate off of its local subnet. For example, a PC needs the IP address of its default gateway, because the default gateway (hopefully) knows how to get to a specified destination network.

However, the PC itself does not need a routing protocol. Rather, the PC simply points to a router, which might be running a routing protocol.

Similarly, the PC does not need a VLAN assignment. Although a PC might belong to a certain VLAN, the VLAN assignment is determined by the switchport into which the PC connects.

Question #35

Which of the follow are dynamic routing protocols? (Choose 2.)
a. ICMP
b. OSPF
c. SNMP
d. ARP
e. EIGRP

Answer: b and e

ICMP (Internet Control Message Protocol) is used by the **ping** command to determine if a given IP address is reachable on the network. This is done by sending out an ICMP Echo Request message and, in response, receiving an ICMP Echo Reply. While ICMP is used for other purposes, its use by the **ping** command is probably the most well-known.

OSPF (Open Shortest Path First) is a distance-vector dynamic routing protocol that uses the *Dijkstra Shortest Path First (SPF) Algorithm* to determine the best path to a remote network.

SNMP (Simple Network Management Protocol) is used to monitor the status of network devices. This monitoring is performed by a Network Management Server (NMS). SNMP also allows the monitored devices to send notifications (called *traps*) to an NMS if specific events occur.

ARP (Address Resolution Protocol) allows a network host to request the MAC address corresponding to a known IP address. Typically, a PC will send out an ARP broadcast after being powered on, so that the PC can learn the MAC address corresponding to the IP address of the PC's default gateway.

EIGRP (Enhanced Interior Gateway Routing Protocol) is an advanced distance-vector dynamic routing protocol that uses the *Diffusing Update Algorithm (DUAL)* to determine the best path to a remote network.

Question #36

Identify the three layers of Cisco's traditional hierarchical network design model. (Choose 3.)

a. Access
b. Cross-connect

c. Distribution
d. Uplink
e. Core

Answer: a, c, and e

The three layers of Cisco's hierarchical network design model include the following layers:

Access Layer: This layer is where network hosts (e.g. PCs, servers, or printers) connect into the network. Access Layer components often include switches (Layer 2 and/or Layer 3 switches) located in a wiring closet. Therefore, this is where port security is typically configured.

Distribution Layer: This layer aggregates connections from multiple Access Layer switches. Distribution Layer switches or routers are often configured to provide redundant connections to downstream Access Layer switches and upstream Core Layer switches. Network security (e.g. Access Control Lists) and Wide Area Network (WAN) connectivity often occur at this layer.

Core Layer: This layer, which is often referred to as the *backbone network*, is primarily concerned with high-speed connectivity between Distribution Layer switches and/or routers. Therefore, the Core Layer is usually composed of high-end multilayer switches.

Question #37

Which of the following devices can be used to interconnect broadcast domains? (Choose 2.)

a. router
b. Ethernet hub
c. Layer 2 switch
d. Ethernet bridge
e. Layer 3 switch

Answer: a and e

Both routers and Layer 3 switches (also known as *multilayer switches*) can route between different subnets. Each subnet is a broadcast domain. Therefore, routers and Layer 3 switches can interconnect broadcast domains.

All ports on an Ethernet hub belong to the same broadcast domain.

An Ethernet bridge is an older device (which makes Layer 2 forwarding decisions in software, as opposed to hardware). Typically, a bridge is used to split up

collision domains and only contains a few ports, all belonging to the same broadcast domain.

Although you can assign ports on a Layer 2 switch to different VLANs (and therefore to different broadcast domains), a Layer 2 switch cannot independently interconnect those broadcast domains.

Question #38

ARP can perform which of the following functions?

a. ARP allows a network host to learn the MAC address corresponding to a known IP address.
b. ARP allows a network host to dynamically obtain an IP address.
c. ARP allows a network host to learn an IP address corresponding to a known domain name.
d. ARP allows a network host to learn the IP address corresponding to a known MAC address.

Answer: a

When a network host knows an IP address (e.g. the IP address of the host's default gateway), but does not know a MAC address corresponding to that IP address, the host can send out an *Address Resolution Protocol (ARP) request*. (**NOTE:** The ARP request is sent as a broadcast.) The device on the network to which the known IP address is assigned receives the ARP request and sends back an *ARP reply* containing its MAC address.

A network host can use *Dynamic Host Configuration Protocol (DHCP)* to dynamically obtain an IP address.

A network host can use the *Domain Name System (DNS)* service to learn the IP address corresponding to a known domain name.

Reverse ARP (RARP) is rarely used. However, several years ago, RARP was used by a network host to learn the IP address corresponding to a known MAC address.

Question #39

You're viewing the running configuration on a Cisco router and notice the **enable secret** command. The command begins with **enable secret**, followed by the number **9**, followed by a hash digest. What does the **9** indicate?
a. The password is encrypted using AES.
b. The password is hashed using Message Digest 5 (MD5).

c. The password is hashed using Password-Based Key Derivation Function 2 (PBKDF2).
d. The password is hashed using SCRYPT.

Answer: d

The **enable secret** command shown in a router's running configuration is followed by a number (indicating the type of hash), followed by the hash digest. Note that the string is a hash digest, not an encrypted string. A type of 5 indicates MD5 hashing. A type of 8 indicates PBKDF2 hashing, and a type of 9 indicates SCRYPT hashing. The default for Cisco routers is type 5 (i.e. MD5 hashing).

Question #40

An Ethernet hub has 12 Ethernet ports. How many collision domains and how many broadcast domains exist on the Ethernet hub?

a. 12 Collision Domains and 12 Broadcast Domains
b. 1 Collision Domain and 12 Broadcast Domains
c. 12 Collision Domains and 1 Broadcast Domain
d. 1 Collision Domain and 1 Broadcast Domain

Answer: d

When multiple devices are connected to an Ethernet hub, all of the devices belong to the same collision domain, because only one of the devices should transmit on the shared segment at any one time. All devices connected to an Ethernet hub are also part of the same broadcast domain, because a broadcast arriving on one port will be flooded out of all other ports on the Ethernet hub.

When multiple devices are connected to an Ethernet switch, each of the devices belongs to their own collision domain. Since collisions are eliminated, Carrier Sense Multiple Access with Collision Detection (CSMA/CD) is not needed on a switch port, and each switch port can run in full-duplex mode (i.e. where a port can transmit and receive data simultaneously), provided the device attached to the port supports full-duplex mode. All devices connected to an Ethernet switch are part of the same broadcast domain, because a broadcast arriving on one port will be flooded out of all other ports on an Ethernet switch.

When multiple devices are connected to a router, each of the devices belongs to its own collision domain and its own broadcast domain.

Question #41

A router performs which of the following tasks? (Choose the 3 best answers.)

a. packet switching
b. rewriting the MAC addresses of Ethernet frames
c. determining the full path to a destination network
d. determining the next hop to a destination network
e. rewriting the IP addresses of IP packets

Answer: a, b, and d

The process of a router receiving a packet on one interface (i.e. the *ingress* interface) and forwarding that packet out of another interface (i.e. the *egress* interface) is called *packet switching*. In this context, the term switching should not be confused with the switching performed by a Layer 2 Ethernet switch (i.e. forwarding frames based on destination MAC addresses).

A router rewrites source and destination MAC address information for frames coming into or going out of an Ethernet interface. However, the router does not typically rewrite source and destination IP address (unless a service such as Network Address Translation (NAT)) is being used).

Also, if you look at a router's IP routing table, you'll notice that it does not contain information about the full path a packet must take to reach a destination network. Rather, the IP routing table contains information about the next hop a packet should take on its way to a destination network.

Question #42

Given the following topology, what would be the most efficient way of configuring routing on the BR1 router, such that devices at the BR1 site would have connectivity to the HQ site and the Internet?

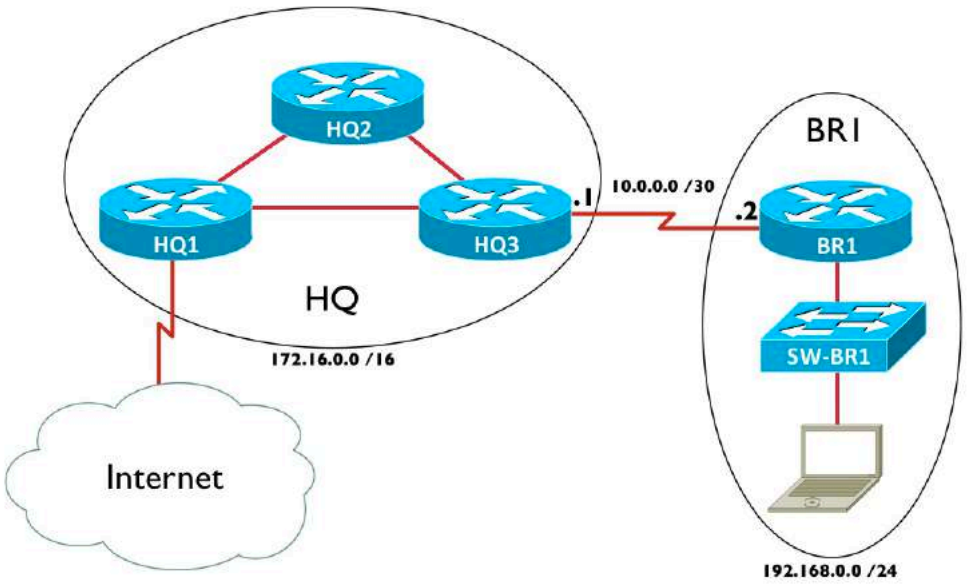

a. Create a static route to the 172.16.0.0 /16 network.
b. Create a default static route.
c. Configure OSPFv3.
d. Configure EIGRP.
e. Configure BGP.

Answer: b

The network at the BR1 site is referred to as a *stub network* (not to be confused with an OSPF stub area), because there is only one way in or out of the network (i.e. via the link to the HQ3 router at the HQ location). Since any non-local network (i.e. any network other than 192.168.0.0 /24) would only be available via the HQ3 router, running a routing protocol on router BR1 (e.g. RIP, OSPF, EIGRP, or BGP) is unnecessary.

Instead, a *default static route* can be configured. A default static route is a statically configured route that tells traffic destined for any unknown IP network to be sent to a specified next-hop address. In this case, the following command would be given to configure a default static route.

BR1(config)# **ip route 0.0.0.0 0.0.0.0 10.0.0.1**

Question #43

Which of the following are components of an Ethernet frame? (Choose 3.)
a. Preamble
b. ToS Byte
c. SOF
d. DF bit
e. FCS

Answer: a, c, and e

The Type of Service (ToS) Byte and Do-Not-Fragment (DF) bit are parts of an IPv4 header. An Ethernet frame contains the following fields: Preamble, Start-of-Frame Delimiter (SOF, or SFD in some literature), Destination MAC Address, Source MAC Address, Type (a.k.a. EtherType), Data and Pad, and Frame Check Sequence (FCS).

Question #44

To what address does EIGRP for IPv6 send updates?
a. 224.0.0.10
b. FF02::A
c. ::1
d. FE80::10

Answer: b

While EIGRP for IPv4 sends messages to a multicast address of 224.0.0.10, EIGRP for IPv6 sends messages to a multicast address of FF02::A.

Question #45

Your router has more than one routing information source that is telling the router how to reach the 10.10.10.0 /24 network. One source is a statically configured route. Another source is the OSPF routing protocol, and yet another source is the EIGRP routing protocol. What determines which routing information source the router will use?

a. Static routes are always trusted over dynamically learned routes.
b. The route with the lowest metric will be used.
c. The information source with the lowest administrative distance will be used.
d. The route will load-balance among the three available routes.

Answer: c

The *administrative distance* (AD) is the believability of a route, and the lower the AD, the more believable the route. By default, a statically configured route has an AD of 1. However, static routes are not always trusted over dynamically learned routes. (**NOTE:** The AD of EIGRP-learned routes is 90, and the AD of OSPF-learned routes is 110.)

For example, let's say that you have a static route configured as a backup route. You do not want to use this backup route unless your primary route fails. To make your statically configured route less believable than a dynamically learned route, you can specify an AD value, as part of your static route configuration, that is higher than the AD value of your dynamic routing protocol.

The syntax of a statically configured route to the 10.10.10.0 /24 network, where the statically configured route has an AD of 150 and a next-hop IP address of 192.168.1.1, is:

Router(config)# **ip route 10.10.10.0 255.255.255.0 192.168.1.1 150**

A *metric* is a measurement used by a routing protocol to measure the distance to reach known destination networks. Since different dynamic routing protocols use different calculations to determine the metric of a route, metrics between routing protocols (e.g. OSPF and EIGRP) cannot be compared.

When a router has multiple sources of information specifying how to reach a destination network, the router takes the route from the routing information

source with the lowest AD and injects that route into the router's IP routing table. No load-balancing occurs between different route information sources.

Question #46

You are sitting at the privileged mode command line prompt on RouterA, and you want to remotely connect to RouterB (which has an IP address of 172.16.1.1). What command would you issue on RouterA?

a. **ssh 172.16.1.1**
b. **traceroute 172.16.1.1**
c. **ping 172.16.1.1**
d. **telnet 172.16.1.1**

Answer: d

Cisco IOS does not support the **ssh** command from the command line interface (CLI).

The **traceroute** command is used to gather information about each router hop along the way to a specified destination address.

The **ping** command is used to check IP connectivity with a specified destination address.

The **telnet** command can be used to remotely connect to another router's VTY (i.e. Virtual TTY) line, if that router's VTY lines support Telnet as a valid transport protocol.

Question #47

You are configuring an EtherChannel connection between two Cisco Catalyst switches, using PAgP for dynamic EtherChannel formation. Ports on one side of the EtherChannel are configured to the Auto PAgP channel mode. To which PAgP channel mode must the ports at the other side of the EtherChannel be configured?
a. On
b. Auto
c. Desirable
d. Active
e. Passive

Answer: c

You can use either **Port Aggregation Protocol** (PAgP) or **Link Aggregation Control Protocol** (LACP) to dynamically form an EtherChannel (i.e. a bundle of

physical links configured to act as a single logical link between two Ethernet switches). PAgP is Cisco-proprietary, while LACP is an industry standard (i.e. IEEE 802.3ad). PAgP has three channel modes:

- **On:** Configures ports to be an EtherChannel without sending or receiving any PAgP frames.
- **Auto:** Configures ports to become an EtherChannel if they receive PAgP frames from the far end. However, this channel mode does not initiate the sending of PAgP frames.
- **Desirable:** Configures ports to initiate the negotiation of an EtherChannel by sending out PAgP frames to the far end, and an EtherChannel will form if the far end is configured for either the Auto or Desirable channel mode.

In this question, one end of the EtherChannel was configured for the Auto mode, which requires the far end to initiate the formation of an EtherChannel. Since neither the On nor the Auto channel modes would initiate the sending of PAgP frames, the ports at the far end must be configured for the Desirable mode.

LACP has its own set of channel modes, including:

- **On:** Configures ports to be an EtherChannel without sending or receiving any LACP frames.
- **Passive:** Configures ports to become an EtherChannel if they receive LACP frames from the far end. However, this channel mode does not initiate the sending of LACP frames.
- **Active:** Configures ports to initiate the negotiation of an EtherChannel by sending out LACP frames to the far end, and an EtherChannel will form if the far end is configured for either the Passive or Active channel mode.

Question #48

Consider the following topology. You issue the **show cdp neighbors** command on router R1. Which of the following devices will appear in the output of the command? (Choose all that apply.)

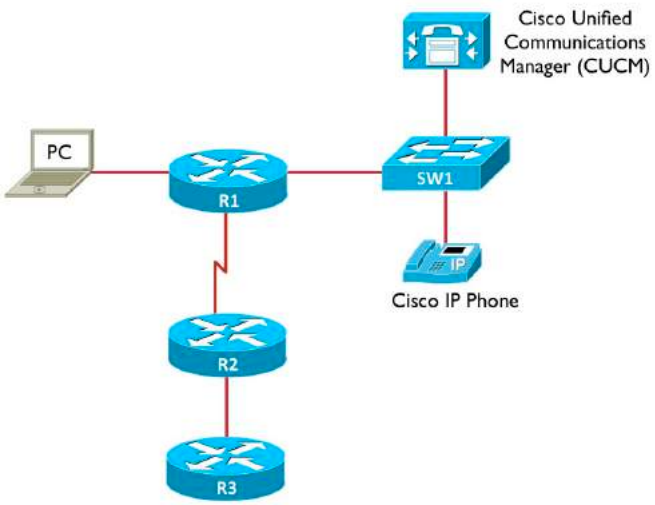

a. PC
b. CUCM server
c. Cisco IP Phone
d. R2
e. R3
f. SW1

Answer: d and f

The **show cdp neighbors** command issued on a Cisco router or switch displays a listing of Layer 2 adjacent Cisco devices enabled for CDP. These devices could be, as a few examples, Cisco routers, Cisco switches, Cisco IP Phones, or Cisco Unified Communications Manager servers.

In this topology, we are issuing the **show cdp neighbors** command on router R1. It's only Layer 2 adjacencies are the PC, router R2, and switch SW1. Since a PC is not a Cisco device enabled for CDP, it will now show up in the output of the **show cdp neighbors** command. However, router R2 and switch SW1 will.

Question #49

The following output was generated from what command?

```
Global values:
  Internet Protocol routing is enabled
FastEthernet0/0 is up, line protocol is up
  Internet address is 10.1.1.1/24
FastEthernet0/1 is up, line protocol is up
  Internet address is 4.4.4.4/24
FastEthernet1/0 is administratively down, line protocol is down
FastEthernet1/1 is administratively down, line protocol is down
```

a. **show cdp neighbor**
b. **show protocols**

c. **show version**
d. **show ip interface brief**
e. **show interfaces**
f. **show running-config**

Answer: b

The **show cdp neighbor** command displays information about other Cisco devices (enabled for *Cisco Discovery Protocol (CDP)*) that are Layer 2 adjacent (e.g. directly connected, or connected via a Layer 1 hub).

The **show protocols** command displays information about the Layer 1 (i.e. Physical) and Layer 2 (i.e. Data Link) status of a router's interfaces. Additionally it shows any IP address assigned to those interfaces.

The **show version** command displays a collection of information, including such things as how long the router has been up, the amount of memory in the router, the interfaces in the router, and the version of Cisco IOS currently running.

The **show ip interface brief** command displays summary information about a router's interfaces, including any IP addresses assigned to interfaces and the Layer 1/Layer 2 interface status.

The **show interfaces** command displays detailed information about a router's interfaces, including such information as the interface's Layer 1/Layer 2 status, MAC address (for an Ethernet interface), MTU (Maximum Transmission Unit) size, bandwidth of the interface, delay of the interface, reliability of the interface, and duplex of the interface.

The **show running-config** command displays a router's configuration, currently running in the router's memory. If the router looses power or is rebooted, the running-config is not preserved. However, a router's startup-config, which resides in the routers non-volatile RAM (NVRAM) is preserved. Therefore, if you make a change to the router's running-configuration and want that configuration to be used, even after a router reboot, you should copy the contents of the running-config to the startup-config. This is accomplished with the command: **copy running-config startup-config**.

Question #50

The default OSPF cost for a 1 Gbps link is 1. What is the default cost for a 100 Mbps link?
a. 1
b. 10
c. 100
d. 1000

Answer: a

By default, OSPF has a Reference Bandwidth of 100 Mbps, and cost is calculated with the formula: **Cost = Reference_BW/Interface_BW**
For a 1 Gbps interface, the cost = 100 Mbps/1000 Mbps = 0.1.

However, we have to have an integer value for the cost. So, 0.1 gets rounded up to 1. For a 100 Mbps interface, the cost = 100 Mbps/100 Mbps = 1.

Therefore, in networks containing link speeds of 100 Mbps and greater, it's often a good decision to change the reference bandwidth to a higher value (e.g. 100,000 Mbps).

Question #51

Which of the follow protocols is used to determine the MAC address for a known IP address?
a. DNS
b. DHCP
c. WINS
d. ARP
e. ICMP

Answer: d

DNS (Domain Name System) provides domain name to IP address translation.

WINS (Windows Internet Name Service) provides NetBIOS name to IP address translation.

ARP (Address Resolution Protocol) allows a network host to request the MAC address corresponding to a known IP address. Typically, a PC will send out an ARP broadcast after being powered on, so that the PC can learn the MAC address corresponding to the IP address of the PC's default gateway.

ICMP (Internet Control Message Protocol) is used by the **ping** command to determine if a given IP address is reachable on the network. This is done by sending out an ICMP Echo Request message and, in response, receiving an ICMP Echo Reply. While ICMP is used for other purposes, its use by the **ping** command is probably the most well-known.

DHCP (Dynamic Host Configuration Protocol) allows a network host to request and dynamically obtain IP address information (e.g. IP address, subnet mask, default gateway, DNS server IP address, WINS server IP address) from a DHCP server. DHCP is such an efficient method of distributing IP addresses to network hosts, Cisco recommends its use even on small networks.

Question #52

A traditional Ethernet switch operates at which layer of the OSI Model?

a. Transport
b. Data Link
c. Network
d. Physical
e. Session

Answer: b

The OSI Model has seven layers:

Layer 1: Physical
Layer 2: Data Link
Layer 3: Network
Layer 4: Transport
Layer 5: Session
Layer 6: Presentation
Layer 7: Application

An Ethernet hub, which does not make any forwarding decisions, but just takes bits coming in on one port and sends them out all other ports, operates at the Physical Layer of the OSI Model.

An Ethernet switch, which can make forwarding decisions based on destination MAC addresses, operates at the Data Link Layer.

A router, which can make forwarding decisions based on network addresses (e.g. IP addresses), operates at the Network Layer.

Question #53

You are in interface configuration mode of a Cisco router, and you want to assign an IP address of 172.16.1.1 /24 to the interface. Which of the following is the command you should enter?

a. Router1(config-if)# **ip address 172.16.1.1 /24**
b. Router1(config-if)# **ip address 172.16.1.1 0.0.0.255**
c. Router1(config-if)# **ip address 172.16.1.1 255.255.255.0**
d. Router1(config-if)# **ip address 172.16.1.1 classful**

Answer: c

When configuring an IP address on a router interface, the subnet mask is represented in dotted decimal notation (e.g. **255.255.255.0).** The use of a wildcard mask (e.g. **0.0.0.255**) or prefix notation (e.g. **/24**) is not supported. Also, the **classful** keyword is not supported.

Question #54

What file transfer protocol uses a connectionless Layer 4 transport protocol and does not required user authentication?

a. TFTP
b. SFTP
c. FTP
d. SSH
e. Telnet

Answer: a

Trivial File Transfer Protocol (TFTP) can be used to transfer files between two networked devices. For example, TFTP is commonly used to download files to a router's flash storage.

TFTP uses UDP, a connectionless protocol, as its Layer 4 transport protocol. Another unique characteristic of TFTP is that a user does not have to authentication before performing a file transfer.

Secure File Transfer Protocol (SFTP) is an encrypted version of **File Transfer Protocol (FTP)**. Both SFTP and FTP uses TCP, a connection-oriented protocol, as their Layer 4 transport mechanism.

Secure Shell (SSH) and Telnet are protocols used to access a remote host (e.g. a UNIX host, a router, or a switch). They both use TCP as their Layer 4 transport protocol. Cisco recommends that SSH be used instead of Telnet, because SSH encrypts data, while Telnet sends data in clear text.

Question #55

A Layer 2 Ethernet switch with 12 ports, where all ports belong to the same VLAN, has how many collision domains and how many broadcast domains?

a. 12 collision domains and 12 broadcast domains
b. 1 collision domain and 1 broadcast domain
c. 12 collision domains and 1 broadcast domain
d. 1 collision domain and 12 broadcast domains

Answer: c

When multiple devices are connected to an Ethernet hub, all of the devices belong to the same collision domain, because only one of the devices should transmit on the shared segment at any one time. All devices connected to an Ethernet hub are also part of the same broadcast domain, because a broadcast arriving on one port will be flooded out of all other ports on the Ethernet hub.

When multiple devices are connected to an Ethernet switch, each of the devices belongs to their own collision domain. Since collisions are eliminated, Carrier Sense Multiple Access with Collision Detection (CSMA/CD) is not needed on a switch port, and each switch port can run in full-duplex mode (i.e. where a port can transmit and receive data simultaneously), provided the device attached to the port supports full-duplex mode. All devices connected to an Ethernet switch are part of the same broadcast domain, because a broadcast arriving on one port will be flooded out of all other ports on an Ethernet switch.

When multiple devices are connected to a router, each of the devices belongs to their own collision domain and its own broadcast domain.

Question #56

RIPv2 advertisements are sent to what multicast IP address?
a. 224.0.0.5
b. 224.0.0.6
c. 224.0.0.9
d. 224.0.0.10

Answer: c

OSPFv2 uses multicast IP addresses of 224.0.0.5 (all OSPF routers) and 224.0.0.6 (all Designated Routers). RIPv2 advertises to 224.0.0.9, and EIGRP for IPv4 advertises to 224.0.0.10.

Question #57

In the diagram below, Client A is sending a packet to Host 1. As the frame is coming into the Fa 0/0 interface on router R2, what is the destination MAC address in the frame's header?

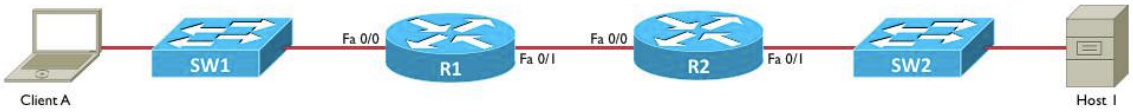

a. Host 1's MAC address
b. Client A's MAC address
c. Router R1's Fa 0/0 MAC address

d. Router R2's Fa 0/0 MAC address
e. Router R2's Fa 0/1 MAC address

Answer: d

As the frame is being sent from Client A to Host 1, the source and destination IP addresses never change. The IP address of Client A is always the source IP address, and the IP address of Host 1 is always the destination IP address.

However, Layer 2 MAC addresses are rewritten at every router hop. In this example, as the packet is coming into router R1's Fa 0/0 interface, the source MAC address is the MAC address of Client A, and the destination MAC address is the MAC address of the Fa 0/0 interface on router R1.

However, the source and destination MAC addresses are rewritten as router R1 sends the frame out to router R2. Specifically, the frame leaving router R1's Fa 0/1 interface and entering router R2's Fa 0/0 interface has a source MAC address of router R1's Fa 0/1 interface, and it has a destination MAC address of router R2's Fa 0/0 interface.

Question #58

You are connected to the console line of RouterA. From there, you connect to RouterB via Telnet. Without terminating the Telnet session, what key sequence could you enter to return to the RouterA prompt?
a. **<CTRL-SHIFT-6> x**
b. **<CTRL-Break>**
c. **<CTRL-ALT-DELETE>**
d. **<CTRL-ALT-x>**

Answer: a

If you are Telnetting from a Cisco router to another Cisco router, you can suspend the Telnet session and return to the original router using a key sequence of **<CTRL-SHIFT-6> x**. To resume the Telnet session, you can press **<ENTER>** or enter the **resume** command.

Question #59

Which three of the following are components of a network secured using IEEE 802.1x? (Choose 3.)
a. Encryption Server
b. Supplicant
c. Authorization Server
d. Key Manager
e. Authenticator

f. Authentication Server

Answer: b, e, and f

The three primary components of an IEEE 802.1x configuration are:

- **Supplicant:** The device wanting to gain access to the network.
- **Authenticator:** The device with which the supplicant communicates (e.g. an Ethernet switch or a wireless access point (AP)).
- **Authentication Server:** A server (e.g. a RADIUS server) that checks the credentials of the supplicant, and tells the authenticator if the supplicant should be allowed on the network.

Question #60

What protocol allows multiple hosts to dynamically obtain IP addresses from a server?
a. DNS
b. DHCP
c. WINS
d. ARP
e. ICMP

Answer: b

DNS (Domain Name System) provides domain name to IP address translation.

WINS (Windows Internet Name Service) provides NetBIOS name to IP address translation.

ARP (Address Resolution Protocol) allows a network host to request the MAC address corresponding to a known IP address. Typically, a PC will send out an ARP broadcast after being powered on, so that the PC can learn the MAC address corresponding to the IP address of the PC's default gateway.

ICMP (Internet Control Message Protocol) is used by the **ping** command to determine if a given IP address is reachable on the network. This is done by sending out an ICMP Echo Request message and, in response, receiving an ICMP Echo Reply. While ICMP is used for other purposes, its use by the **ping** command is probably the most well-known.

DHCP (Dynamic Host Configuration Protocol) allows a network host to request and dynamically obtain IP address information (e.g. IP address, subnet mask, default gateway, DNS server IP address, WINS server IP address) from a DHCP server. DHCP is such an efficient method of distributing IP addresses to network hosts, Cisco recommends its use even on small networks.

Question #61

Which of the following are true regarding CSMA/CD? (Choose 3.)

a. CSMA/CD is an Ethernet technology.
b. CSMA/CD is a wireless technology.
c. CSMA/CD should run on every port of an Ethernet switch.
d. CSMA/CD is used for half-duplex connections.
e. CSMA/CD is used by devices connecting to Ethernet hubs.

Answer: a, d, e

Carrier Sense Multiple Access with Collision Detection (CSMA/CD) is an Ethernet technology that allows a device connected to an Ethernet segment to listen to a network segment for a brief period of time to determine if there is currently a frame on the segment. If no frame is detected, the device can transmit its frame. However, if that device transmits its frame at approximately the same time another device transmits a frame, a collision could result. A collision corrupts both frames. Fortunately, a device running CSMA/CD can detect that a collision occurred and retransmit the frame after waiting for a random *backoff time*.

Since CSMA/CD is an Ethernet technology, while Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) is usually seen in wireless networks.

CSMA/CD does not need to be run by devices operating in full-duplex mode. Therefore, when a device connects into an Ethernet switch port, if the switch port is configured for full-duplex operation (i.e. is able to simultaneously transmit and receive), there is no need for it to run CSMA/CD on that port, because on that Ethernet segment, there are only two devices (i.e. the switch port and the device connected to the switch port). Therefore, no collisions are possible.

Ports on an Ethernet switch, however, run in half-duplex mode (i.e. can only transmit or receive at any one time). In fact, all ports on an Ethernet hub are part of the same network segment. Therefore, collisions are possible, and devices connecting to an Ethernet hub should run CSMA/CD.

Question #62

The following output was generated from what command?

```
FastEthernet0/1 is up, line protocol is up
  Hardware is i82543 (Livengood), address is ca03.1af0.0006 (bia
ca03.1af0.0006)
  Internet address is 4.4.4.4/24
  MTU 1500 bytes, BW 100000 Kbit/sec, DLY 100 usec,
     reliability 255/255, txload 1/255, rxload 1/255
```

```
Encapsulation ARPA, loopback not set
Keepalive set (10 sec)
Full-duplex, 100Mb/s, 100BaseTX/FX
ARP type: ARPA, ARP Timeout 04:00:00
Last input 00:00:00, output 00:00:03, output hang never
Last clearing of "show interface" counters never
Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: fifo
Output queue: 0/40 (size/max)
5 minute input rate 0 bits/sec, 0 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
   476 packets input, 41501 bytes
   Received 333 broadcasts (0 IP multicasts)
   0 runts, 0 giants, 0 throttles
   0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
   0 watchdog
   0 input packets with dribble condition detected
   481 packets output, 41819 bytes, 0 underruns
   0 output errors, 0 collisions, 1 interface resets
   0 unknown protocol drops
   0 babbles, 0 late collision, 0 deferred
   0 lost carrier, 0 no carrier
   0 output buffer failures, 0 output buffers swapped out
```

a. **show cdp neighbor**
b. **show protocols**
c. **show version**
d. **show ip interface brief**
e. **show interfaces**
f. **show running-config**

Answer: e

The **show cdp neighbor** command displays information about other Cisco devices (enabled for *Cisco Discovery Protocol (CDP)*) that are Layer 2 adjacent (e.g. directly connected, or connected via a Layer 1 hub).

The **show protocols** command displays information about the Layer 1 (i.e. Physical) and Layer 2 (i.e. Data Link) status of a router's interfaces. Additionally it shows any IP address assigned to those interfaces.

The **show version** command displays a collection of information, including such things as how long the router has been up, the amount of memory in the router, the interfaces in the router, and the version of Cisco IOS currently running.

The **show ip interface brief** command displays summary information about a router's interfaces, including any IP addresses assigned to interfaces and the Layer 1/Layer 2 interface status.

The **show interfaces** command displays detailed information about a router's interfaces, including such information as the interface's Layer 1/Layer 2 status, MAC address (for an Ethernet interface), MTU (Maximum Transmission Unit)

size, bandwidth of the interface, delay of the interface, reliability of the interface, and duplex of the interface.

The **show running-config** command displays a router's configuration, currently running in the router's memory. If the router looses power or is rebooted, the running-config is not preserved. However, a router's startup-config, which resides in the routers non-volatile RAM (NVRAM) is preserved. Therefore, if you make a change to the router's running-configuration and want that configuration to be used, even after a router reboot, you should copy the contents of the running-config to the startup-config. This is accomplished with the command: **copy running-config startup-config**.

Question #63

You connect a laptop to an available port on a Cisco Catalyst switch. By default, how long does it take that port to start forwarding traffic?
a. 0 seconds
b. 15 seconds
c. 30 seconds
d. 50 seconds

Answer: c

By default, Spanning Tree Protocol (STP) is enabled on a switch port. Even though it takes (by default) 50 seconds for a Blocking switch port to transition to Forwarding when the switch's Root Port is no longer the best port to get back to the Root Bridge, in this question, the port was available. Therefore, it did not need to spend 20 seconds the Blocking State. Instead, it spends 15 seconds in the Listening State and 15 seconds in the Learning State, for a total of 30 seconds.

Question #64

What is another term for Port Address Translation?

a. static NAT
b. dynamic NAT
c. NAT overloading
d. NAT pooling

Answer: c

Network Address Translation (NAT) allows private IP addressed to be used inside of a network and have those addresses translated into publicly routable IP addresses.

Static NAT is a configuration where an administrator has statically configured which private IP addresses get mapped to which public IP addresses.

Dynamic NAT is a configuration where the public IP address to which a private IP address is translated is dynamically selected from a pool of publicly routable IP addresses.

Port Address Translation (PAT) is also known as NAT Overloading, because with PAT, multiple private IP addresses can be mapped to a single publicly routable IP address. The router is able to distinguish between the incoming traffic flows to the same publicly routable IP address by keeping track of port number information.

Question #65

By default, a DHCP Discover message cannot pass through a router, because it is a broadcast packet. What interface configuration-mode command can cause the router to forward the DHCP Discover message to a target IP address or subnet?

a. Router(config-if)# **dhcp-relay *ip-address***
b. Router(config-if)# **ip helper-address *ip-address***
c. Router(config-if)# **ip discover-forward *ip-address***
d. Router(config-if)# **forward-bootp *ip-address***

Answer: b

Four Dynamic Host Configuration Protocol (DHCP) messages are used when a PC is allocating an IP address from a DHCP server. Those messages are: (1) Discover, (2) Offer, (3) Request, (4) Acknowledgement. The initial Discover message is a broadcast message. As a result, it is dropped by a router, by default. However, the **ip helper-address *ip-address*** interface configuration mode command can be used to tell a router's interface to forward the Discover broadcast to a specified destination IP address or subnet.

Question #66

What is the Administrative Distance (AD) for External EIGRP?
a. 90
b. 110
c. 120
d. 170

Answer: d

Even though EIGRP has an Administrative Distance (AD) of 90, if a route is redistributed into an EIGRP Autonomous System (AS), it's considered to be an "External EIGRP" route. An External EIGRP route has an AD of 170.

Question #67

You enter the following commands in a router:

```
Router(config)# enable secret Pa$$1
Router(config)# enable password Pa$$2
```

What password must you enter the next time you attempt to enter privileged mode on the router?

a. You must enter both passwords.
b. Pa$$1
c. Pa$$2
d. Since the passwords do not match, remote authentication is disabled.

Answer: b

Either the **enable password** or the **enable secret** command can be used to set the password required for privileged mode access. By default, the password entered as part of the **enable password** command appears as a clear text in a router's configuration. However, the password entered as part of the **enable secret** command appears as a hash value in the router's configuration. Since the **enable secret** password is therefore more secure, Cisco recommends using the **enable secret** command instead of the **enable password** command. If both passwords are entered in a router, as in this question, the **enable secret** value is used, and the **enable password** value is ignored.

Question #68

Your new employee tells you that they are unable to log into router R1 via Telnet. You examine the router's configuration and find the following configuration:

```
... OUTPUT OMITTED ...
line vty 0 4
 exec-timeout 0 0
 logging synchronous
 login
... OUTPUT OMITTED ...
```

Why is your new employee unable to log into router R1 via Telnet?

a. The logging synchronous command requires that the login be authenticated by a AAA server. Therefore, the issue must be with the AAA server.
b. The exec-timeout 0 0 command causes an instantaneous timeout whenever someone attempts to log into the router.
c. Due to SSH being more secure, Telnet access is disabled by default.
d. No password is configured for the VTY lines.

Answer: d

When configuring a line (e.g. a con, aux, or vty line) to authenticate users, you can enter the **login** command to require a login. However, the **login** command is ignored until a password is specified. In this example, no password is configured for the VTY lines. Since a password is required, but not set, the new employee cannot Telnet into the router.

Sometimes as you are entering a command in the Cisco IOS command line interface (CLI), a syslog message might pop up on screen in the middle of the command you were typing. To prevent this behavior, you can enter the **logging synchronous** command. After issuing this line configuration mode command, if a syslog message pops up in the middle of your command, you will be given a fresh line, without the syslog message, displaying what you have already typed and allowing you to complete your command, free of the distraction of the syslog command.

The **exec-timeout *min sec*** command can be used to specify a period of inactivity after which a user will be logged out of their connection to a router. However, the **exec-timeout 0 0** disables the timeout timer, thus preventing a user from being timed out.

Even though Telnet is less secure than SSH, Telnet, by default, is a supported transport protocol for connecting into a router. You can use the **transport-input *protocol*** command to specify which protocol, or protocols, to support for remote connections.

Question #69

You are examining a router's running configuration and notice that the password for the VTY lines is in clear text:

```
...OUTPUT OMITTED...
line vty 0 4
 exec-timeout 0 0
 password cisco
 logging synchronous
 login
...OUTPUT OMITTED...
```

You want the VTY line password to be encrypted in the running configuration, as follows:

```
line vty 0 4
 exec-timeout 0 0
 password 7 02050D480809
 logging synchronous
 login
```

What command should you enter to encrypt the VTY line password?

a. Router(config-line)# **enable password-encryption**
b. Router(config)# **enable password-encryption**
c. Router(config-line)# **service encryption**
d. Router(config)# **service password-encryption**

Answer: d

The **service password-encryption** command can be issued from global configuration mode to encrypt router passwords, including all of the passwords under the con, aux, and tty lines. Please note that this encryption is considered to be fairly weak, and can easily be broken by utilities on the Internet. However, it does help prevent a casual observer from “shoulder surfing” and seeing the password in clear text.

Question #70

Which type of DNS record is used to map a hostname to an IPv6 address?
a. A
b. AAAA
c. SOA
d. MX

Answer: b

Domain Name System (DNS) is a service used to resolve fully-qualified domain names (FQDNs) to IP addresses. When configuring a DNS server, there are a variety of DNS record types you can configure. Following are a few examples of these record types:

- **A**: An **address record** is used to map a hostname to an IPv4 address.
- **AAAA**: An **IPv6 address record** is used to map a hostname to an IPv6 address.
- **CNAME**: A **canonical name record** is an alias of an existing record, thus allowing multiple DNS records to map to the same IP address.

- **MX**: A **mail exchange record** maps a domain name to an e-mail server for that domain.
- **PTR**: A **pointer record** points to a canonical name and is commonly used when performing a reverse DNS lookup, which is used to determine what domain name is associated with a known IP address.
- **SOA**: A **start of authority record** provides authoritative information about a DNS zone, such as: e-mail contact information for the zone's administrator, the zone's primary name server, and various refresh timers.

Question #71

Consider the following port security configuration:

```
Switch(config)# int gig 1/0/1
Switch(config-if)# switchport mode access
Switch(config-if)# switchport port-security
Switch(config-if)# switchport port-security maximum 3
Switch(config-if)# switchport port-security mac-address sticky
Switch(config-if)# switchport port-security violation restrict
```

What does the **restrict** option in the bottom command do?

a. The **restrict** option causes the port to go into an err-disable state if a port security violation occurs.
b. The **restrict** option causes the port to be administratively shutdown if a port security violation occurs.
c. The **restrict** option disables port security on this port.
d. The **restrict** option drops traffic from the fourth MAC address seen on this port, while permitting traffic from the first three MAC addresses learned on this port. However, the security violation counter does not get incremented.
e. The **restrict** option drops traffic from the fourth MAC address seen on this port, while permitting traffic from the first three MAC addresses learned on this port. Also, the security violation counter is incremented.

Answer: e

The **switchport mode access** command places the port in access mode. So, the port carries traffic for only one VLAN. A port must be in access mode in order to enable port security on that port.

The **switchport port-security** command enables port security on the port.

The **switchport port-security maximum 3** command specifies the maximum number of allowed MAC addresses learned off of a port as three. If an additional MAC address is seen off of the port, a port violation occurs.

The **switchport port-security mac-address sticky** command allows the maximum number of MAC addresses learned off of a port to be entered into a

switch's running configuration. You should copy the running configuration to the startup configuration if you want the learned MAC addresses to be retained after a switch reboot. Alternately, you could specify one or more MAC addresses allowed off of an interface with the command **switchport port-security mac-address *MAC_Address***.

The **switchport port-security violation shutdown** command says that the port will be placed into an err-disable state if a port violation occurs. To bring a port out of err-disable state (after clearing the condition that caused the violation), you can administratively shutdown the port (with the **shutdown** command) and then administratively bring the port back up (with the **no shutdown** command).

The **switchport port-security violation protect** option drops traffic from any unpermitted MAC addresses, while permitting traffic from the permitted MAC addresses learned on this port. However, the security violation counter does not get incremented.

The **switchport port-security violation restrict** option also blocks traffic from unpermitted MAC addresses, while allowing traffic from permitted MAC addresses. However, the **restrict** option does increment the security violation counter.

A port's security violation counter can be displayed with the **show port-security command**.

Question #72

If Spanning Tree Protocol (STP) operation fails in a network with redundant links, what symptoms could result? (Choose 3.)

a. broadcast storm
b. MAC address table corruption
c. duplex mismatch
d. duplicate frames received by the intended receiver

Answer: a, b, and d

Spanning Tree Protocol (STP) allows a network to have a redundant Layer 2 topology, while preventing a Layer 2 topological loop. However, if STP were to fail, broadcast, multicast, and unknown unicast traffic could circulate through the redundant connection consuming network bandwidth and consuming CPU cycles of devices on the network segment. The underlying issue is that a Layer 2 frame, unlike a Layer 3 packet, does not have a Time-to-Live (TTL) field. As a result, frames can continuously circulate through the network without ever timing out.

One symptom that could result from such a scenario is a broadcast storm. Since a switch never learns a destination address of FFFF.FFFF.FFFF (i.e. the destination MAC address for a broadcast frame), a switch will flood that frame out all of its interfaces (in the same VLAN as the receiving port), other than the port the on which the frame was received. This can cause multiple copies of a frame to be endlessly circulating through a network.

This flooding behavior can cause a frame flooded out of one switch to be seen on a port on another switch that is not the port off of which the source device resides. As a result, that switch could incorrectly conclude that the frame's source MAC address resided off a certain port. This effect is known as *MAC address table corruption*.

When STP fails, a switch port that should be blocking traffic could be forwarding traffic. This could result in more than one flow of traffic from the sender to the receiver, which could cause the intended destination to receive more than one copy of each frame.

An STP issue would not cause a duplex mismatch. A duplex mismatch occurs when one side of a link is operating in full-duplex mode, while the other side of the link is operating in half-duplex mode. This could occur, because the devices at each end of the link failed to correctly negotiate what duplex they would use. Or, the devices at each end of the link could have had their duplex incorrectly configured. A commonly reported symptom of a duplex mismatch is slow network performance. Specifically, some traffic is passing. However, many frames are being dropped and have to be retransmitted (if the traffic is using a reliable transport protocol, such as TCP).

Question #73

You wish to remotely connect to a Cisco Catalyst 2960 switch. Which of the following parameters must be configured on the switch? (Choose 2.)
a. a default static route
b. an IP address
c. a default gateway
d. a routing protocol

Answer: b and c

A Cisco Catalyst 2960 switch is a Layer 2 switch. Therefore, it does not run a routing protocol and cannot be configured with a static route. Instead, much like a PC, the switch needs an IP address and a default gateway to communicate with a remote device. This management IP address for a Cisco Catalyst 2960 is typically configured under the switch's VLAN 1 interface.

Question #74

Which of the following features allows a Cisco Catalyst switch to create a copy of frames appearing on a switch port or in a VLAN, and send those copied frames out of a designated port?

a. SPAN
b. CEF
c. HSRP
d. VRRP

Answer: a

The **Switched Port Analyzer** (SPAN) feature allows a Cisco Catalyst switch to create a copy of frames appearing on a switch port or in a VLAN, and send those copied frames out of a designated port. A packet capture device could then be connected to that port, which can dramatically assist in your troubleshooting efforts.

Question #75

Given a subnet of 172.16.56.0 /21, identify which of the following IP addresses belong to this subnet. (Select 2.)

a. 172.16.54.129
b. 172.16.62.255
c. 172.16.61.0
d. 172.16.65.255
e. 172.16.64.1

Answer: b, c

To determine subnets and usable address ranges created by the 21-bit subnet mask we perform the following steps:

Step #1: Identify the interesting octet (i.e. the octet that contains the first zero in the binary subnet mask).

In this question, we have a 21-bit subnet mask, which is written in binary as:
11111111 11111111 11111000 00000000

The interesting octet is the third octet, because the third octet (i.e. 11111000) is the first octet to contain a 0 in the binary subnet mask.

Step #2: Identify the decimal value in the interesting octet of the subnet mask.

A 21-bit subnet mask can be written in dotted decimal notation as: 255.255.248.0

Since the third octet is the interesting octet, the decimal value in the interesting octet is 248.

Step #3: Determine the block size by subtracting the decimal value of the interesting octet from 256.

Block Size = 256 – 248 = 8

Step #4: Determine the subnets by counting by the block size in the interesting octet, starting at 0.

Placing a zero in the first interesting octet identifies the first subnet as:
172.16.0.0 /21

We then count by the block size (of 8) in the interesting octet (the third octet in this question) to determine the remaining subnets:

172.16.8.0 /21
172.16.16.0 /21
172.16.24.0 /21
172.16.32.0 /21
172.16.40.0 /21
172.16.48.0 /21
172.16.56.0 /21
172.16.64.0 /21
... SUBNETS OMITTED ...

We can stop counting after we pass the subnet we are being asked about. Specifically, in this question, we're being asked about 172.16.56.0 /21.

Step #5: Identify the subnet address, the directed broadcast address, and the usable range of addresses.

The subnet address, where all host bits are set to a 0, is given:
172.16.56.0 /24

The directed broadcast address, where all host bits are set to a 1, is 1 less than the next subnet address.

The next subnet address is 172.16.64.0. So, the directed broadcast address for the 172.16.54.0 /21 subnet is 1 less than 172.16.64.0, which is:
172.16.63.255

The usable IP addresses are all the IP addresses between the subnet address and the directed broadcast address. Therefore, in this example, the usable IP address range for the 172.16.56.0 /21 network is:

172.16.56.1 – 172.16.63.254

The only IP addresses in this question that reside in this range are:
172.16.62.255
172.16.61.0

WARNING: Many CCNA R&S candidates look at IP addresses like these and immediately assume they are not usable IP addresses, because they have a 0 or a 255 in the forth octet. They argue that 172.16.61.0 is a subnet address and that 172.16.62.255 is a directed broadcast address.

While that would only be true of the subnet mask were 24-bits, remember that, by definition, a subnet address has all of its host bits set to a 0, and a directed broadcast address has all of its host bits set to a 1. In this question, we have 11 host bits (i.e. 32 – 21 = 11), not 8 host bits. So, 172.16.62.255 and 172.16.61.0 are actually usable IP addresses.

Question #76

What is the subnet address of the IP address 192.168.5.55 with a subnet mask of 255.255.255.224?
a. 192.168.5.0 /27
b. 192.168.5.16 /27
c. 192.168.5.32 /27
d. 192.168.5.48 /27
e. 192.168.5.64 /27

Answer: c

To determine subnets and usable address ranges created by the 27-bit subnet mask we perform the following steps:

Step #1: Identify the interesting octet (i.e. the octet that contains the first zero in the binary subnet mask).

In this question, we have a 27-bit subnet mask, which is written in binary as:
11111111 11111111 11111111 11100000

The interesting octet is the forth octet, because the forth octet (i.e. 11100000) is the first octet to contain a 0 in the binary.

Step #2: Identify the decimal value in the interesting octet of the subnet mask.

A 27-bit subnet mask can be written in dotted decimal notation as:
255.255.255.224

Since the forth octet is the interesting octet, the decimal value in the interesting octet is 224.

Step #3: Determine the block size by subtracting the decimal value of the interesting octet from 256.

Block Size = 256 – 224 = 32

Step #4: Determine the subnets by counting by the block size in the interesting octet, starting at 0.

Placing a zero in the first interesting octet identifies the first subnet as:
192.168.5.0 /27

We then count by the block size (of 32) in the interesting octet (the forth octet in this question) to determine the remaining subnets:

192.168.5.32 /27
192.168.5.64 /27
192.168.5.96 /27
192.168.5.128 /27
192.168.5.160 /27
192.168.5.192 /27
192.168.5.224 /27

Now that we have all of our subnets identified, we can determine the subnet in which the IP address of 192.168.5.55 resides.

Since the usable range of IP addresses for the 192.168.5.32 /27 network is 192.168.5.33 – 192.168.5.62 (because 192.168.5.32 is the network address, and 192.168.5.63 is the directed broadcast address), and since 192.168.5.55 is in that range, the subnet to which 192.168.5.55 /27 belongs is:
192.168.5.32 /27

Question #77

Which of the following protocols allow you to view Layer 2 adjacent network devices from a Cisco router or Cisco Catalyst switch command prompt? (Choose 2.)
a. MLP
b. CDP
c. RTP
d. LLDP

Answer: b and d

Both **Cisco Discovery Protocol** (CDP) and **Link Layer Discovery Protocol** (LLDP) allow Cisco routers and Cisco Catalyst switches to view information about Layer 2 adjacent devices (supporting CDP and/or LLDP). However, CDP is Cisco-proprietary, while LLDP is an industry standard (i.e. IEEE 802.1AB).

Question #78

You are working for a company that will be using the 192.168.1.0 /24 private IP address space for IP addressing inside their organization.

They have multiple geographical locations and want to carve up the 192.168.1.0 /24 address space into subnets. Their largest subnet will need 13 hosts.

What subnet mask should you use to accommodate at least 13 hosts per subnet, while maximizing the number of subnets that can be created?

a. 255.255.255.248
b. 255.255.255.224
c. 255.255.255.252
d. 255.255.255.192
e. 255.255.255.240

Answer: e

We can determine the maximum number of hosts allowed in a subnet by raising the number 2 to the power of the number of host bits and then subtracting 2. So, the formula looks like this:

Maximum Number of Hosts per Subnet = $2^h - 2$, where *h* is the number of host bits.

Why are we subtracting two? Well, there are two IP addresses in the subnet that cannot be assigned. These addresses are: (1) the network address, where all of the host bits are set to a 0 and (2) the directed broadcast address, where all of the host bits are set to a 1.

In the actual exam, if you are given scratch paper or access to a note taking application, you might want to write out a table such as the following for your reference:

1 Host Bit: $2^1 - 2 =$ **0**
2 Host Bits: $2^2 - 2 =$ **2**
3 Host Bits: $2^3 - 2 =$ **6**
4 Host Bits: $2^4 - 2 =$ **14**
5 Host Bits: $2^5 - 2 =$ **30**
6 Host Bits: $2^6 - 2 =$ **62**

7 Host Bits: $2^7 - 2 =$ **126**
8 Host Bits: $2^8 - 2 =$ **254**

In this question, we're asked to determine a subnet mask that accommodates at least 13 hosts per subnet. By looking at the reference table we created, we can see that 4 host bits (which support 14 hosts) would work, while 3 host bits (which supports only 6 hosts) would not be enough.

So, we need a subnet with 4 host bits, which are enough host bits to meet the design goal, but not more than we need. Using more host bits than we need would violate the requirement to maximize the number of subnets.

A subnet mask with 4 host bits has 28 network bits (i.e. 32 – 4 = 28), and therefore a 28-bit subnet mask. A 28-bit subnet mask can be written as: **255.255.255.240**

Question #79

A customer is using a Class C network of 192.168.10.0 subnetted with a 28-bit subnet mask. How many assignable addresses are available in each of the subnets?

a. 32
b. 16
c. 30
d. 8
e. 14

Answer: e

An IPv4 address contains a total of 32 bits. Since, in this question, we have 28 subnet bits, the number of host bits is 4 (i.e. 32 – 28 = 4). The number of assignable IP addresses in a subnet can be calculated as follows:

Number of Assignable IP Addresses = $2^h - 2$, where *h* is the number of host bits.

Therefore, in this question, each subnet has 14 assignable IP addresses:

Number of Assignable IP Addresses = $2^4 - 2 = 16 - 2 = 14$

Question #80

Which of the following are configurable versions of Simple Network Management Protocol (SNMP) within Cisco IOS? (Choose 3.)
a. v1

b. v1c
c. v2
d. v2c
e. v3
f. v3c

Answer: a, d, and e

Cisco IOS supports the configuration of SNMP versions 1, 2c, and 3. Versions 1 and 2c are considered to have weak security as compared to version 3. Specifically, while version 3 supports an authentication protocol (i.e. MD5 or SHA), versions 1 and 2c only support community strings. Although there was at one time a version 2 of SNMP, its security was very challenging to configure and support. As a result, the version 1 security solution of community strings was combined with new enhancements in version 2, to create SNMP version 2c.

Question #81

An IP address of 192.168.0.100 /27 belongs to which of the following subnets?

a. 192.168.0.92
b. 192.168.0.128
c. 192.168.0.64
d. 192.168.0.96
e. 192.168.0.32

Answer: d

To determine the subnets created by the 27-bit subnet mask we perform the following steps:

Step #1: Identify the interesting octet (i.e. the octet that contains the first zero in the binary subnet mask).

In this question, we have a 19-bit subnet mask, which is written in binary as:
11111111 11111111 11111111 11100000

The interesting octet is the forth octet, because the forth octet (i.e. 11100000) is the first octet to contain a 0 in the binary.

Step #2: Identify the decimal value in the interesting octet of the subnet mask.

A 27-bit subnet mask can be written in dotted decimal notation as:
255.255.255.224

Since the forth octet is the interesting octet, the decimal value in the interesting octet is 224.

Step #3: Determine the block size by subtracting the decimal value of the interesting octet from 256.

Block Size = 256 – 224 = 32

Step #4: Determine the subnets by counting by the block size in the interesting octet, starting at 0.

Placing a zero in the first interesting octet identifies the first subnet as:
192.168.0.0 /27

We then count by the block size (of 32) in the interesting octet (the forth octet in this question) to determine the remaining subnets:

192.168.0.32 /27
192.168.0.64 /27
192.168.0.96 /27
192.168.0.128 /27
192.168.0.160 /27
192.168.0.192 /27
192.168.0.224 /27

Step #5: Identify the subnet address of the IP address 192.168.0.100 /27.

Looking through the subnets created by the 27-bit subnet mask reveals that the IP address of 192.168.0.100 resides in the **192.168.0.96** subnet.

Question #82

Which of the following is considered to be an unreliable Transport Layer protocol?

a. IP
b. UDP
c. TCP
d. ICMP
e. PPP

Answer: b

IP (Internet Protocol) and **ICMP (Internet Control Message Protocol)** are Network Layer (i.e. Layer 3) protocols.

PPP (Point-to-Point Protocol) is a Data Link Layer (i.e. Layer 2) protocol.

Transmission Control Protocol (TCP) and **User Datagram Protocol (UDP)** are both Transport Layer (i.e. Layer 4) protocols. However, TCP uses acknowledgements to confirm receipt of data, while UDP does not confirm receipt of data. Therefore, TCP is considered to be a reliable, connection-oriented protocol, while UDP is considered to be an unreliable, connectionless protocol.

Question #83

What subnet mask should be used to subnet the 192.168.10.0 network to support the number of subnets and IP addresses per subnet shown in the following topology?

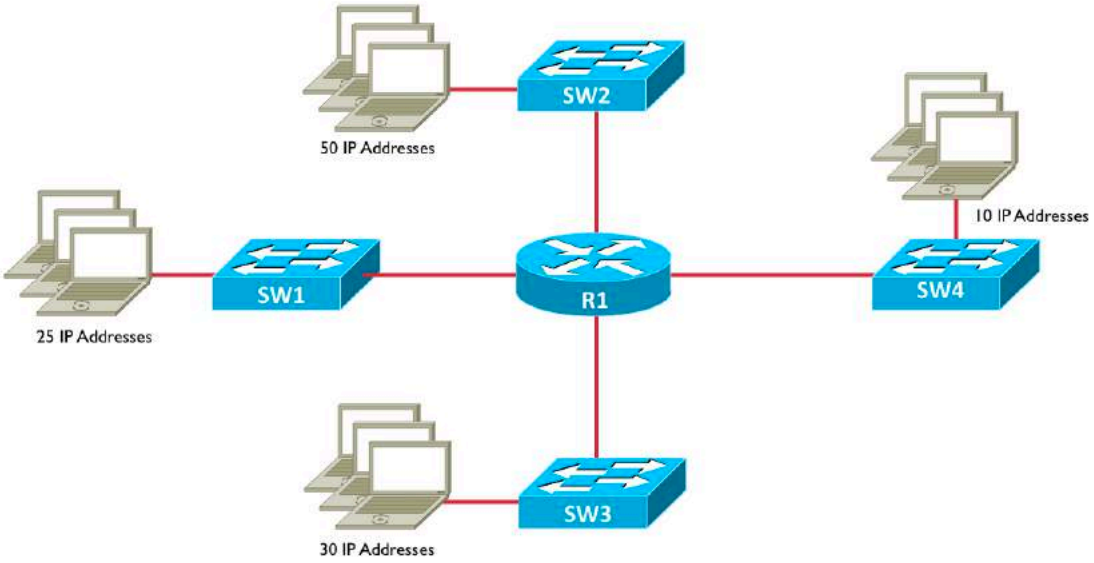

a. 255.255.255.0
b. 255.255.255.128
c. 255.255.255.192
d. 255.255.255.224
e. 255.255.255.240

Answer: c

To meet the design requirements, four subnets must be created, and each subnet must accommodate a maximum of 50 IP addresses.

We can begin by creating a listing of how many subnets are created from different numbers of borrowed bits, using the formula:

Number of Subnets Created = 2^n, where *n* is the number of borrowed bits

1 borrowed bits => 2 subnets

2 borrowed bits => 4 subnets
3 borrowed bits => 8 subnets
4 borrowed bits => 16 subnets
5 borrowed bits => 32 subnets
6 borrowed bits => 64 subnets
7 borrowed bits => 128 subnets

From this, we can see we need at least 2 borrowed bits to accommodate 4 subnets. However, we need to make sure the subnet will accommodate 50 IP addresses. To determine this, we can use the formula:

Number of IP Addresses = 2^h – 2, where *h* is the number of host bits

If we have 2 borrowed bits (i.e. the minimum number of borrowed bits required for 4 subnets), we have 6 host bits (i.e. 8 – 2 = 6). From the above formula, we can determine the number of IP addresses supported by 6 host bits.

Number of IP Addresses = 2^6 – 2 = 62

Since 6 host bits meet our requirement of at least 50 IP addresses per subnet, we can use a 26-bit subnet mask (i.e. 2 bits added to the Class C default mask (also known as the *natural mask*) of 24 bits). A 26-bit subnet mask can be written as:

255.255.255.192

Question #84

How can an IPv6 address of **2200:5678:0001:0000:0000:000A:0000:0001** can be abbreviated?

a. 22:5678:1::A:0:1
b. 2200:5678:1::A:0:1
c. 2200:5678:1::A::1
d. 22:5678:1::a:0:1
e. 2200:5678:1::A:0::1

Answer: b

An IPv6 address is a 128-bit address and can be written in 8 fields, with 4 hexadecimal digits in each field. To make these addresses easier to work with, an IPv6 address can, in some instances, be abbreviated. Following are the abbreviation rules you can use:

- Leading 0s in a field can be omitted.
- Contiguous fields containing all 0s can be written with a double colon.

- The double colon representation of contiguous fields containing all 0s can only be used once per summarized IPv6 address.

In this question, the first field of 2200 does not begin with any 0s. So, it has to be completely written out as **2200**.

The second field of 5678 does not begin with any 0s. Therefore, it has to be completely written out as **5678**.

The third field of 0001 does begin with leading 0s, which can be omitted. This means that the third field can be abbreviated as **1**.

The fourth and fifth fields are contiguous fields containing all 0s. These fields can be collectively written as **::**.

The sixth field of 000A contains leading 0s, which can be omitted, allowing that field to be written as **A**.

The seventh field contains all 0s. Since the double colon has already been used once in this abbreviation, it cannot be used again. However, this all 0s field can be abbreviated as **0**.

The eighth field of 0001 contains leading 0s, which can be omitted, resulting in an abbreviation of **1**.

Combining all of these abbreviations yields the following abbreviated IPv6 address:

2200:5678:1::A:0:1

Question #85

IPv6 unique local addresses are similar to IPv4 private IP addresses, because they cannot be routed over the public Internet. A unique local IPv6 address begins with which of the following patterns?

a. 2000::/3
b. FE80::/10
c. FF02::1:FF
d. FC00::/7
e. FF

Answer: d

IPv6 addresses are 128-bit addresses, commonly written in 8 fields, with 4 hexadecimal digits in each field. There are several categories of IPv6 addresses,

and you can identify the category by how an IPv6 address begins. Following is a listing of common IPv6 address categories, along with their hexadecimal pattern:

- Global unicast addresses begin with **2000::/3**.
- Multicast addresses begin with an **FF**.
- Link local addresses begin with **FE80::/10**.
- Unique local addresses begin with **FC00::/7**.
- The IPv6 loopback address is **::1**.
- An IPv6 unspecified address is written as **::**.
- Solicited-node multicast addresses begin with **FF02::1:FF**.

Question #86

In an attempt to recover a lost password on a Cisco router, you issue a Break during the router's boot sequence. This takes you to the ROM Monitor prompt. From there, you want to set the configuration register such that the router's startup configuration will be ignored the next time it boots. To which of the following values should you set the configuration register?
a. 2102
b. 2142
c. 0x2142
d. 0x2102

Answer: c

The **0x** prefix is a requirement for a configuration register value, because it specifies the number is a hexadecimal value. If you want a router to ignore its startup configuration, you can specify a configuration register value of **0x2142**. However, the configuration register is typically set to a value of **0x2102**, which does not ignore the startup configuration during boot up.

Question #87

You issue the **ping 192.168.1.2** command from a router, and the response displayed on the router is:

```
M.M.M
```

What does this response indicate?

a. The router had to ARP for the MAC address of the next hop IP address.
b. The router is attempting to load balance across two links, and one of the links is not working.
c. The Ping packets needed to be fragmented, but the packets have their DF bit set (which says they cannot be fragmented).
d. The Ping is successful, and the alternating M and dot characters indicate the two directions of the bidirectional communication.

Answer: c

If a router does not know the MAC address of the next hop IP address, it can use Address Resolution Protocol (ARP) to learn the unknown MAC address. This process of sending an ARP Request and receiving an ARP Reply can take time. Therefore, the first one or two Ping packets might time out, resulting in a Ping response of: **`..!!!`** or **`.!!!!`**

If a router is attempting to load balance across two links, and one of the links is not functional, the Ping responses might alternate between a successful response (i.e. an explanation point) and a timeout (i.e. a dot). Therefore, the Ping response might look like the following: **`!.!.!`**

If a router sends a Ping packet whose size is greater than the *Maximum Transmission Unit (MTU)* of another router that resides along the path to the destination, that other router needs to fragment the Ping packet in order for it to be transmitted. However, if the DF bit is set in the Ping packet's header, the router will not be permitted to do the fragmentation. So, the router sends an ICMP message back to the sending router, indicating that fragmentation was needed but not permitted. This results in a Ping response on the originating router of: **`M.M.M`**

If a Ping is successful, the result will be a series of exclamation points: **`!!!!!`**

Question #88

What network architecture layers are combined in a **collapsed core** design?
a. Access and Distribution
b. Distribution and Core
c. Access and Core
d. Access, Distribution, and Core

Answer: b

A Collapsed Core design is sometimes called a *Two Tier design*, where the two tiers are the Access Layer and the Collapsed Core layer. This consolidated Collapsed Core layer combines the functions of the traditional *Distribution* and *Core* layers.

Question #89

What type of Cisco Catalyst Switch port configuration allows a port to be an access port that supports two VLANs, if and only if one of the two VLANs is designated as a voice VLAN?
a. Voice VLAN Access Port
b. Voice Trunk Port

c. Multi-VLAN Access Port
d. 802.1p Access Port

Answer: c

While we typically think of a trunk port as the type of port supporting multiple VLANs, Cisco lets us bend the rules with a voice VLAN. Specifically, we can configure a port as an access port and then configure both a data VLAN and a voice VLAN for that port. This type of port is called a *multi-VLAN access port*.

Question #90

What type of queuing adds a priority queue to CB-WFQ?
a. ECN
b. LLQ
c. WRED
d. WFQ

Answer: b

Low Latency Queuing (LLQ) allows you to place one or more types of priority traffic into a priority queue. Other traffic types can be placed into non-priority queues, which can be given minimum bandwidth guarantees. Note that the priority queue used by LLQ is policed (i.e. has a bandwidth limit).

Question #91

Which of the following are features of Point-to-Point Protocol (PPP)? [Choose 4.]
a. Authentication
b. Encryption
c. Error Detection and Correction
d. Logical Bundling of Multiple Links
e. Compression

Answer: a, c, d, and e

PPP supports authentication, compression, error detection & correction, and the logical bundling of multiple links. However, PPP does not support encryption.

Question #92

A ***Cisco Application Policy Infrastructure Controller - Enterprise Module*** (APIC-EM) uses Northbound APIs to connect to what?
a. Network Devices
b. Peer Controllers
c. Autonomous Controllers
d. Applications

Answer: d

Northbound APIs connect the APIC-EM to Applications. Southbound APIs connect the APIC-EM to Network Devices. Eastbound and Westbound APIs connect a controller to peer controllers.

Question #93

What Cisco technology allows you to interconnect multiple physical switches into a single logical switch?
a. SmartNet
b. Optimum Switching
c. Stackwise
d. Collapsed Core

Answer: c

Cisco Stackwise allows you to interconnect multiple Cisco Catalyst switches (that have Stackwise support) using special interconnect cables. You can then administer the stack of switches as one logical unit.

Question #94

What technology allows an enterprise to more easily change their cloud provider (e.g. change from Microsoft to AWS)?
a. CloudFront
b. Intercloud Exchange
c. MP-BGP
d. APIC-EM

Answer: b

An enterprise can have a connection to an *Intercloud Exchange* provider, which has connections to multiple cloud providers. This allows the enterprise to switch to a different cloud provider while keeping their existing connection to the Intercloud Exchange.

Question #95

A **Unique Local IPv6 Address**, which cannot be routed over the public Internet, begins with what hexadecimal prefix?
a. FE80::/10
b. FC00::/7
c. FF02::1:FF
d. 2000::/3

Answer: b

A *Unique Local IPv6 Address* begins with FC00::/7. These addresses cannot be routed over the public Internet.

Recommended Study Resource

Cisco CCNA R&S (200-125) Complete Video Course

This is a video training series that I personally did for Cisco Press. It covers all of the exam blueprint topics, includes 300+ videos, 25+ hours of instruction, interactive exercises, Glossary Quizzes, Module Quizzes, and more full practice exams.

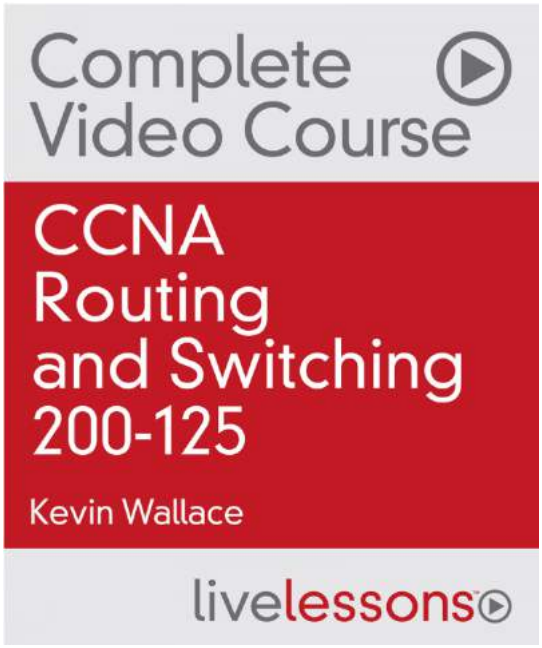

More info available here: **http://kwtrain.com/ccnacourse**

Keep in Touch with Kevin

I am deeply honored that you've trusted me to help you prepare for your exam. If you have found value in this practice exam, please leave a review on Amazon.com.

If you would like to keep in contact with me, or explore the video training I offer, please visit my website:

http://kwtrain.com

You can also follow me on the following social networks:
Twitter: http://twitter.com/kwallaceccie
Facebook: http://facebook.com/kwallaceccie
YouTube: http://youtube.com/user/kwallaceccie
LinkedIn: http://linkedin.com/in/kwallaceccie
Google+: http://google.com/+KevinWallace
Snapchat: kwallaceccie

About the Author

Kevin Wallace, CCIEx2 (R&S and Collaboration) No. 7945, CCSI No. 20061
With Cisco experience dating back to 1989, Kevin has been a Network Design Specialist for the Walt Disney World Resort, an Instructor of Cisco courses for a Cisco Learning Partner (CLP), and a Network Manager for Eastern Kentucky University.
Kevin currently creates video training courses and writes books on networking technologies (**http://kwtrain.com/products**). He lives in central Kentucky with his wife (Vivian) and two daughters (Stacie and Sabrina).